

Contributors

Francis Alt
Arjang A. Assad
Edward Baker
Michael Ball
Maria G. Bardossy
Atul Bhatt
Alfred Blumstein
Lawrence Bodin
Beryl E. Castello
Kenneth Chelst
Geraldine S. Cheok
Michael C. Fu
Bruce L. Golden
Alan J. Goldman
Damon J. Gulczynski
Jeffrey W. Heath
Karla Hoffman
Wolfgang Jank
Anito Joseph
Anthony J. Kearsley
Rudolph P. Lamone
Richard Larson
Jim Lawrence
Feiyue Li
Thomas Magnanti
Kweku-Muata Osei-Bryson
Carter C. Price
Guruprasad Pundoor
S. Raghavan
Randall S. Robinson
Joe Seviek
Douglas R. Shier
Rakesh V. Vohra
Gang Wang
Edward Wasil
Chris Witzgall
Zhanging Zhao

PERSPECTIVES IN OPERATIONS RESEARCH

Papers in Honor of Saul Gass' 80th Birthday

Edited by
Frank B. Alt
Michael C. Fu
Bruce L. Golden



PERSPECTIVES IN OPERATIONS RESEARCH

Papers in Honor of Saul Gass' 80th Birthday

OPERATIONS RESEARCH/COMPUTER SCIENCE INTERFACES SERIES

Professor Ramesh Sharda
Oklahoma State University

Prof. Dr. Stefan Voß
Universität Hamburg

Other published titles in the series:

Greenberg /*A Computer-Assisted Analysis System for Mathematical Programming Models and Solutions: A User's Guide for ANALYZE*

Greenberg / *Modeling by Object-Driven Linear Elemental Relations: A Users Guide for MODLER*

Brown & Scherer / *Intelligent Scheduling Systems*

Nash & Sofer / *The Impact of Emerging Technologies on Computer Science & Operations Research*

Barth / *Logic-Based 0-1 Constraint Programming*

Jones / *Visualization and Optimization*

Barr, Helgason & Kennington / *Interfaces in Computer Science & Operations Research: Advances in Metaheuristics, Optimization, & Stochastic Modeling Technologies*

Ellacott, Mason & Anderson / *Mathematics of Neural Networks: Models, Algorithms & Applications*

Woodruff / *Advances in Computational & Stochastic Optimization, Logic Programming, and Heuristic Search*

Klein / *Scheduling of Resource-Constrained Projects*

Bierwirth / *Adaptive Search and the Management of Logistics Systems*

Laguna & González-Velarde / *Computing Tools for Modeling, Optimization and Simulation*

Stilman / *Linguistic Geometry: From Search to Construction*

Sakawa / *Genetic Algorithms and Fuzzy Multiobjective Optimization*

Ribeiro & Hansen / *Essays and Surveys in Metaheuristics*

Holsapple, Jacob & Rao / *Business Modelling: Multidisciplinary Approaches — Economics, Operational and Information Systems Perspectives*

Sleezer, Wentling & Cude/*Human Resource Development And Information Technology: Making Global Connections*

Voß & Woodruff / *Optimization Software Class Libraries*

Upadhyaya et al / *Mobile Computing: Implementing Pervasive Information and Communications Technologies*

Reeves & Rowe / *Genetic Algorithms—Principles and Perspectives: A Guide to GA Theory*

Bhargava & Ye / *Computational Modeling And Problem Solving In The Networked World: Interfaces in Computer Science & Operations Research*

Woodruff / *Network Interdiction And Stochastic Integer Programming*

Anandalingam & Raghavan / *Telecommunications Network Design And Management*

Laguna & Martí / *Scatter Search: Methodology And Implementations In C*

Gosavi/ *Simulation-Based Optimization: Parametric Optimization Techniques and Reinforcement Learning*

Koutsoukis & Mitra / *Decision Modelling And Information Systems: The Information Value Chain*

Milano / *Constraint And Integer Programming: Toward a Unified Methodology*

Wilson & Nuzzolo / *Schedule-Based Dynamic Transit Modeling: Theory and Applications*

Golden, Raghavan & Wasil / *The Next Wave in Computing, Optimization, And Decision Technologies*

Rego & Alidaee/ *Metaheuristics Optimization via Memory and Evolution: Tabu Search and Scatter Search*

Kitamura & Kuwahara / *Simulation Approaches in Transportation Analysis: Recent Advances and Challenges*

Ibaraki, Nonobe & Yagiura / *Metaheuristics: Progress as Real Problem Solvers*

Golumbic & Hartman / *Graph Theory, Combinatorics, and Algorithms: Interdisciplinary Applications*

Raghavan & Anandalingam / *Telecommunications Planning: Innovations in Pricing, Network Design and Management*

Mattfeld / *The Management of Transshipment Terminals: Decision Support for Terminal Operations in Finished Vehicle Supply Chains*

Alba & Martí/ *Metaheuristic Procedures for Training Neural Networks*

PERSPECTIVES IN OPERATIONS RESEARCH

Papers in Honor of Saul Gass' 80th Birthday

Edited by

FRANCIS B. ALT
University of Maryland

MICHAEL C. FU
University of Maryland

BRUCE L. GOLDEN
University of Maryland



Francis B. Alt
Michael C. Fu
Bruce L. Golden
University of Maryland

Library of Congress Control Number: 2006931932

ISBN-10: 0-387-39933-X (HB) ISBN-10: 0-387-39934-8 (e-book)
ISBN-13: 978-0387-39933-1 (HB) ISBN-13: 978-0387-39934-8 (e-book)

Printed on acid-free paper.

© 2006 by Springer Science+Business Media, LLC
All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer Science+Business Media, LLC, 233 Spring Street, New York, NY 10013, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now know or hereafter developed is forbidden.
The use in this publication of trade names, trademarks, service marks and similar terms, even if the are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

9 8 7 6 5 4 3 2 1

springer.com

Preface

Saul Gass has been a leading contributor to the field of Operations Research for more than 50 years. He has been affiliated with the Robert H. Smith School of Business at the University of Maryland for more than 30 years.

On February 25, 2006, "Operations Research in the 21st Century: A Symposium in Honor of Professor Saul Gass' 80th Birthday," was held on our campus. Opening remarks by Deans Howard Frank and Rudy Lamone were followed by talks by Alfred Blumstein, Karla Hoffman, Richard Larson, Christoph Witzgall, Thomas Magnanti, Rakesh Vohra, and Bruce Golden. The celebration continued into the evening with dinner in the Executive Dining Room of Van Munching Hall, followed by numerous toasts to Saul. It was a special day for all of us who were in attendance, but it was especially memorable for Saul and his family.

This Festschrift companion to the Symposium includes articles from each of the Symposium distinguished speakers plus 16 other articles from friends, colleagues, and several of Saul's former students. The book is divided into three sections. The first section comprises eight articles focusing on the field of Operations Research from a historical or professional perspective. The second section contains nine articles whose theme is optimization and heuristic search, while the third section includes six articles in the general area of modeling and decision making. Collectively, these articles pay tribute to Saul Gass' major interests in the field of Operations Research.

We thank Howard Frank for his sponsorship of the Symposium and dinner, Arjang Assad for suggesting a special day to honor Saul, and G. Anandalingam for his support of our efforts. In addition, we single out Ruth Zuba for her invaluable help in all phases of the development of this volume. Finally, we thank Saul Gass for the outstanding contributions he has made to the field of Operations Research and to the University of Maryland, and for the enormous impact he has had on our lives.

College Park, Maryland
July 2006

Frank Alt
Michael Fu
Bruce Golden

Contents

Photo Gallery 1

Part I History & Perspectives

Reflections on Saul Gass' Influence
Rudolph P. Lamone 19

Four Score Years of Saul I. Gass: Portrait of an OR Professional
Arjang A. Assad 23

In the Beginning: Saul Gass and Other Pioneers
Alfred Blumstein 73

Learning from the Master: Saul Gass, Linear Programming and the OR Profession
Thomas Magnanti 77

Looking Backwards, Looking Forwards: Reflections on Definitions of Operations Research by Morse and Kimball
Richard Larson 99

Ben Franklin: America's First Operations Researcher
Bruce L. Golden 115

Good Management, the Missing XYZ Variables of OR Texts
Kenneth Chelst and Gang Wang 123

The Operations Research Profession: Westward, Look, the Land is Bright
Randall S. Robinson 135

Part II Optimization & Heuristic Search

Choosing a Combinatorial Auction Design: An Illustrated Example
Karla Hoffman 153

Label-Correcting Shortest Path Algorithms Revisited
María G. Bardossy and Douglas R. Shier 179

The Ubiquitous Farkas Lemma
Rakesh V. Vohra 199

Parametric Cardinality Probing in Set Partitioning
Anito Joseph and Edward Baker 211

A Counting Problem in Linear Programming
Jim Lawrence 223

Towards Exposing the Applicability of Gass & Saaty's Parametric Programming Procedure
Kweku-Muata Osei-Bryson 235

The Noisy Euclidean Traveling Salesman Problem: A Computational Analysis
Feiyue Li, Bruce Golden, and Edward Wasil 247

The Close Enough Traveling Salesman Problem: A Discussion of Several Heuristics
Damon J. Gulczynski, Jeffrey W. Heath, and Carter C. Price 271

Twinless Strongly Connected Components
S. Raghavan 285

Part III Modeling & Making Decisions

EOQ Rides Again!
Beryl E. Castello and Alan J. Goldman 307

Federal Express Sort Facility Employee Scheduling Problem
Lawrence Bodin, Zhanging Zhao, Michael Ball, Atul Bhatt, Guruprasad Pundoor, and Joe Seviek 333

Sensitivity Analysis in Monte Carlo Simulation of Stochastic Activity Networks
Michael C. Fu 351

The EM Algorithm, Its Randomized Implementation and Global Optimization: Some Challenges and Opportunities for Operations Research
Wolfgang Jank 367

Recovering Circles and Spheres from Point Data
Christoph Witzgall, Geraldine S. Cheok, and Anthony J. Kearsley 393

Why the New York Yankees Signed Johnny Damon
Lawrence Bodin 415

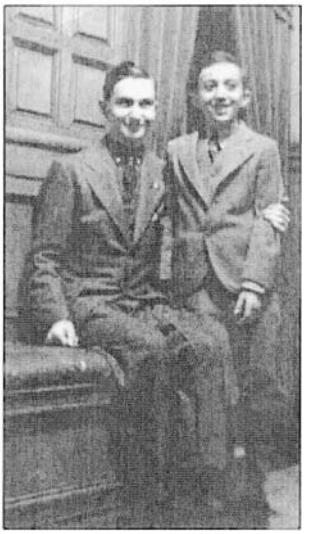
Index 429

Photo Gallery

Family



Parents: Bertha and Louis Gass



With brother Jerry – 1937

The Early Years



Grammar School – 1935



Grammar School – 1936

America's Weapon of Mass Destruction


DRILL WINNERS

Headmaster Robert B. Masterson of Roxbury Memorial High School reads off points to prize drill winners, left to right, Colonel George Epstein, Lieutenant Colonel Bernard Singer, Major Joel Cohen, Major Sumner Goldberg and Major Saul Gass.

Roxbury Memorial High (Boys) High School Military Drill Officers – 1943



After infantry basic training
Boston – June 1944



Postwar R&R boat trip
Danube River, Austria – June 1945

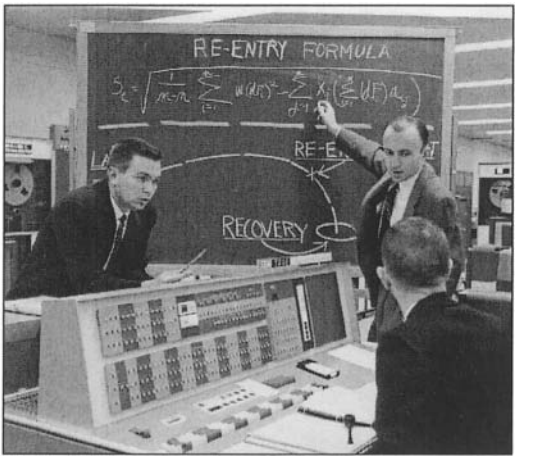
Early Career



First job: Aberdeen Bombing Mission, Los Angeles,
Picnic at Edwards Air Force Base – 1950



IBM ID photo – May 1955

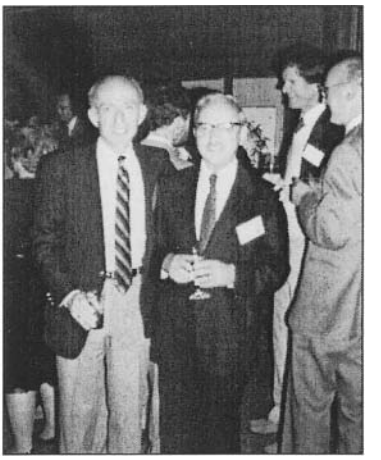


IBM Computer Center, Goddard Space
Flight Center, Greenbelt, MD – 1962

Significant Moments



Jack Borsting presents the ORSA President's gavel –1976



With George Dantzig PhD advisor – 1985



Editors Gass & Harris celebrate Encyclopedia of OR/MS -- 1996

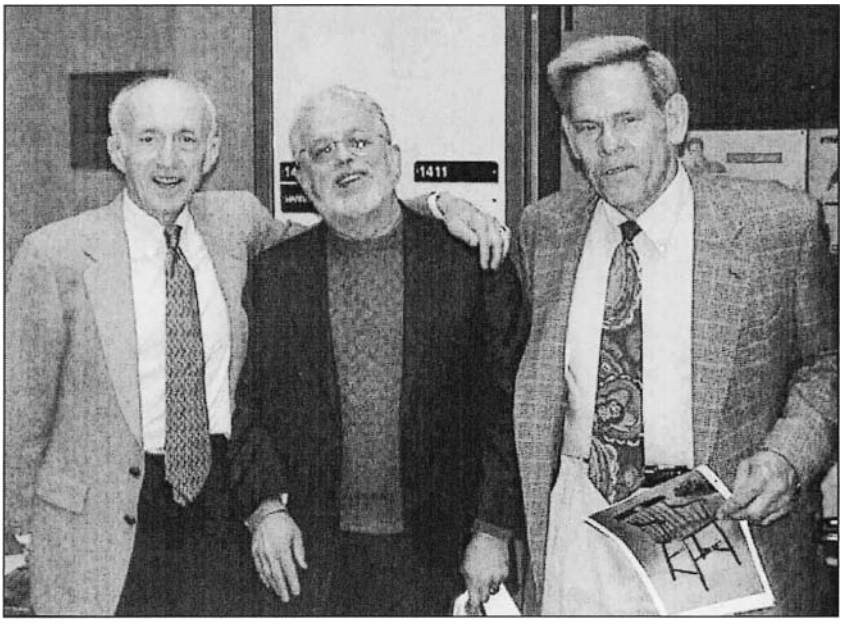


Saul emcees INFORMS Knowledge Bowl – 1996

Business School Colleagues



Westinghouse Professor colleagues – 1990
(Arjang, Saul, Bruce, Larry)



Retirement gang (Saul, Larry, Gus) – 2001

The Athlete



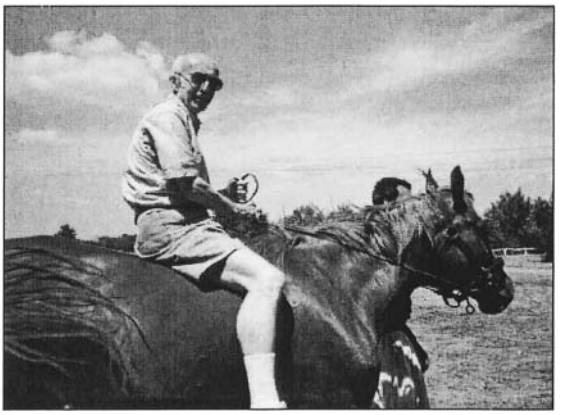
Track Team Practice
Northeastern U. – 1943



10K track meet, ORSA-TIMS
Anaheim – 1991



10K track meet with
Fred Glover, Lisbon – 1993



Brokeback Mountain, Hungary – 1993



ROBERT H. SMITH
SCHOOL OF BUSINESS

Leaders for the Digital Economy

Operations Research in the 21st Century: A Symposium in Honor of Professor Saul Gass on his 80th Birthday

Saturday, February 25, 2006, 1-6pm
Van Munching Hall, Frank Auditorium

Program

1:00-1:30 Opening Remarks
Program Committee members
Dr. Howard Frank, Dean, Smith School
Dr. Rudolph Lamone, former Dean

1:30-2:40 Session 1
Dr. Alfred Blumstein, Carnegie-Mellon University
"*In the Beginning: Saul Gass and Other Pioneers*"
Dr. Karla Hoffman, George Mason University
"*Explaining OR to Management: Letting Illustrations Tell the Story – the Saul Gass Approach*"

2:40-3:00 Break (coffee, tea, refreshments, and snacks)

3:00-4:10 Session 2
Dr. Richard Larson, MIT
"*Looking Backwards, Looking Forwards: Reflections on Definitions of Operations Research by Morse and Kimball*"
Dr. Christoph Witzgall, NIST (retired)
"Recovering Circles and Spheres"

4:10-4:30 Break (coffee, tea, refreshments, and snacks)

4:30-5:40 Session 3
Dr. Thomas Magnanti, MIT (Dean of Engineering)
"Learning from the Master: Saul Gass, Linear Programming, and the OR Profession"
Dr. Rakesh Vohra, Northwestern University
"Predicting the Unpredictable"

5:40-6:00 Closing Remarks

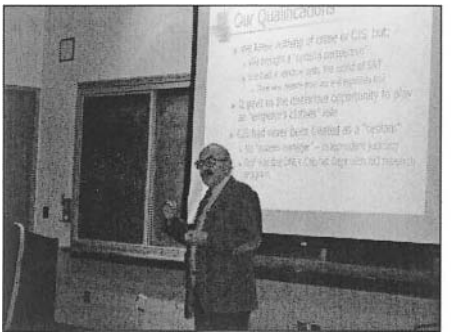
The Symposium



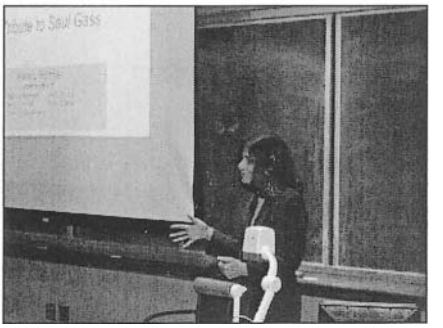
Dean Frank welcomes the attendees



Former Dean Rudy Lamone:
"Top 10 reasons why I hired Saul"



Prof. Al Blumstein-CMU



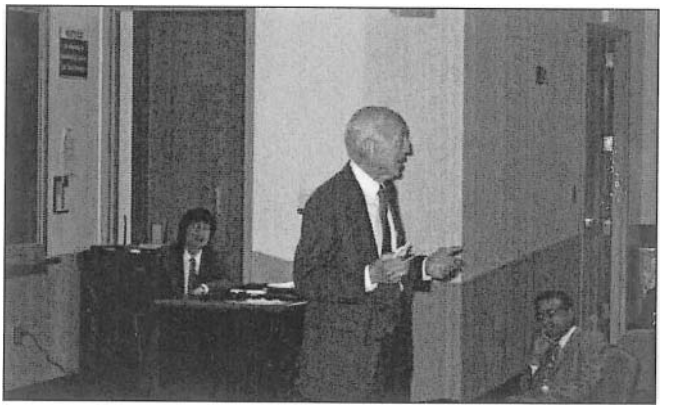
Prof. Karla Hoffman-GMU



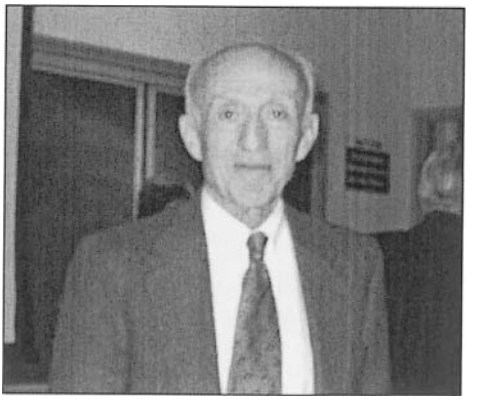
Prof. Bruce Golden



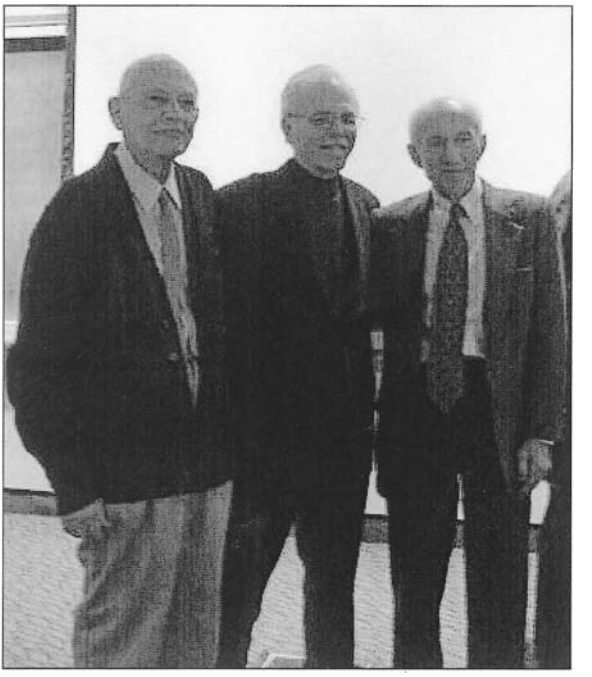
The captive audience



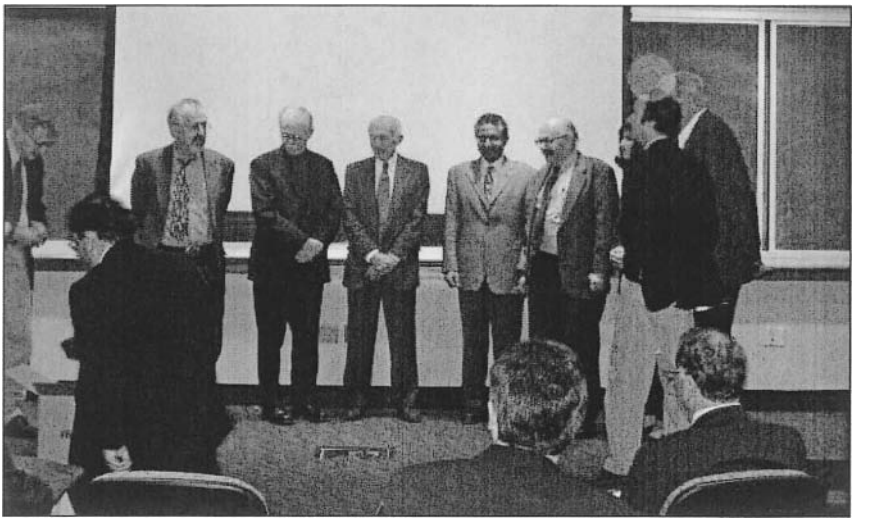
Saul wraps it up!



Phew! That's over!



Alan Goldman, Tom Magnanti, and Saul

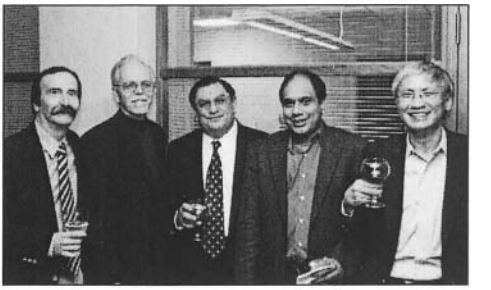


A round of applause for the distinguished speakers

The Festivities Begin

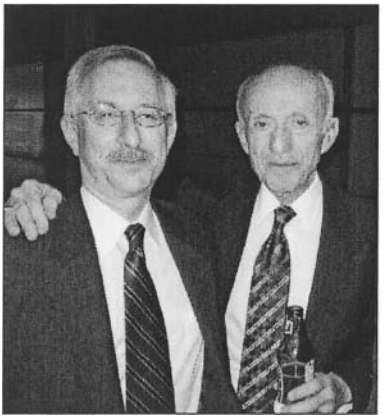
The Reception











Saul with son Ron



Saul with daughter Joyce

The Dinner



Michael welcomes the dinner guests and turns it over to Frank



"I can't believe they're letting me do this!"



Rabbi Israel gives the blessing



The Gass family enjoys dinner



Ron pays tribute to Dad



It's Joyce's turn to pay tribute to Dad



Finally, it's granddaughter Arianna's turn

Further tributes to Saul by



Dean Howard Frank



Former Dean Rudy Lamone



Arjang Assad



Larry Bodin



It's Frank's turn



Why is Trudy laughing?



Bruce enjoys the festivities



Arianna leads singing of "Happy Birthday"

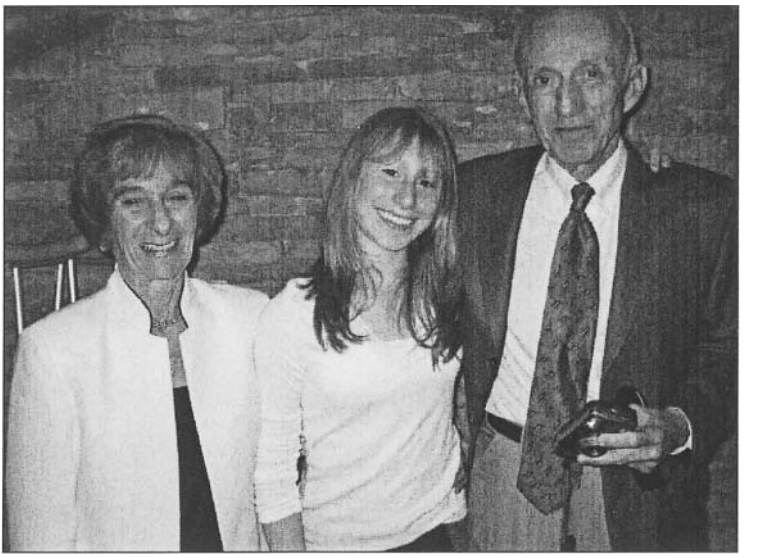


Saul, Joyce, and Arianna after the dinner

Trudy – Saul's Optimal Solution



Marriage photo: Saul and Trudy – August 1946



Trudy, Arianna, and Saul – 2005

Part I

History & Perspectives

Reflections on Saul Gass' Influence

Rudolph P. Lamone (former Dean)

Robert H. Smith School of Business
University of Maryland
College Park, MD 20742
rlamone@rhsmith.umd.edu

Summary. This is an abridged and edited transcription of the author's opening remarks delivered on February 25, 2006.

Key words: Linear programming; operations research.

1 On Retirement

I don't know what to say about Saul, but the word "retirement" is not in his vocabulary. Peter Horner refers to Saul very reverently as "high octane gas." I think that is just a wonderful description of the person we're honoring today. This symposium is really a wonderful tribute to my friend and colleague of over 25 years, Saul Gass.

2 The Book

I believe that the thousands of students who studied Saul's book, *Linear Programming* – not to mention faculty and practitioners around the world – all acknowledge the enormous influence and impact this one man has had on their lives and careers and on the field of operations research. I met Saul first through his book when I was a graduate student at the University of North Carolina (UNC) at Chapel Hill. I was one of eight Ph.D. students enrolled in the first academic program in operations research at UNC. Sometime after finishing the linear programming course, several of us were sitting on a bench on campus on a very hot June Carolina day when someone suggested, "Gee I think we need to get a beer." I said, "Well that's a good idea," and we all reached into our pockets and came up short. We didn't have enough money to buy our beer and I looked at Saul's book and said, "Well, hell, let's go to the bookstore and we'll sell Saul's book." So I sold Saul's book for something like seven bucks, enough to by a cheap six-pack of beer that sure tasted wonderful on that hot day. But I felt badly about it; I genuinely felt badly about it. So when my next scholarship check came, I immediately went back to the bookstore and bought back Saul's book. So, he made royalties on me twice, first with the new book, and then with the used book.

3 Ambassador and Leader

I came to the University of Maryland in 1966 to join a small group of faculty interested in starting an OR program. We introduced several new courses in mathematical programming and recruited several students to the program. The math department became furious. They wanted to know what the hell the business school was doing teaching mathematics. So, that was my first encounter with the problems I was going to have to deal with in eventually becoming chairman of that operations research group.

Long story short, during my tenure as chairman, I was asked to serve as the dean of this new school, now known as the Smith School of Business. And what a dilemma! I had a young faculty, a new program, and no money. Where the hell was I going to find someone to lead this wonderful new program in operations research? In walked Saul Gass, one day, and he didn't need a job, but he was willing to talk. And I put together one of the best packages in the history of this school, a selling package and recruiting package that I could take to Saul Gass to see if I could convince him to take a 20% to 30% reduction in his salary at Mathematika and come to work for the business school at Maryland. Everyone thought I was crazy to think that I could get this giant in the field of operations research to come work for this little school just trying to get off the ground. But God bless him, he said, "I'm going to do this for you."

It's wonderful to see what one person can do to dramatically change an organization, and over time change the make-up of the organization. The kind of excitement, respect, and love that he could build among faculty and students was just absolutely amazing. And I should say that the math department really did forgive me. They said, "Well, if Saul Gass is going to work for the business school, the business school must be doing something right." When they said the business school has no right teaching mathematics, I said, "We don't teach mathematics; we teach mathematical modeling as it applies to solving complex decision problems." I stole that line from Saul, years and years ago.

As the chair of the management science and statistics group, Saul became an ambassador of operations research. He went to our School of Public Affairs and volunteered to teach a course in operations research and mathematical modeling that would focus on complex decision systems in the public sector. So it was Saul, again, who made himself available to faculty in other disciplines, so that he could share with them the power of this field we call *operations research.*

I remember clearly when Saul came to me to talk about the opportunity of recruiting Bruce Golden. And when he told me, I got very excited about this young brilliant guy that Tom Magnanti had said such wonderful things about. And I said to Saul that we have to have this kid; bring him here for faculty interviews. So Saul brings him to my office. The first thing you noticed was this beautiful mustache, but he also had a damn pony tail that went all the way down his back. I almost fell out of my chair. I said to Saul, "What the hell am I going to do with this kid? We'll have to lock him up when my business friends come to this school." Anyway Saul said, "Forget that; we've got to hire this kid." And we did, and it was one of the great early hires that Saul made, and he went on to recruit an outstanding

faculty, guys like Mike Ball, Arjang Assad, and Frank Alt. And they were hired as assistant professors, because we didn't have money to go out and hire senior faculty. So, it was the brilliance of Saul Gass to see in young people emerging in the field of operations research that they could bring really significant value to the program here at Maryland.

4 Conclusions

I am forever grateful to Saul, because he really made me look good, and deans always appreciate faculty members who make them look good. So, Saul, Happy Birthday, my friend, and many more.

Four Score Years of Saul I. Gass: Portrait of an OR Professional

Arjang A. Assad

Robert H. Smith School of Business
University of Maryland
College Park, Maryland 20742
aassad@rhsmith.umd.edu

Summary. Saul Gass has practiced operations research for 55 years and continues to be a vigorous presence in the field. His multifaceted contributions as a scholar, practitioner, involved citizen, world ambassador, and chronicler of operations research are reviewed in this article.

Key Words: Linear programming; multiple criteria; modeling process; history of OR.

1 Introduction

Saul I. Gass turned eighty on February 28, 2006. As a member of Project SCOOP, he was certainly "in just after the beginning" of Operations Research (OR). Saul embraced OR as a profession in its early days. As a member of the Operations Research Society of America since 1954, Saul has been a model citizen who has bridged the academic-practitioner divide. Over a period of nearly 55 years, his involvement in OR as a scholarly and professional endeavor has never slackened and continues in full force to the present day.

Saul's career spans scholarship and what he calls the three P's of OR: practice, process, and professionalism. He has been a leader on all four scores. This article reviews his contributions to date. The first part of this paper is chronological. Sections 2-5 discuss Saul's early career and the road that led him to Project SCOOP, through his career as a practitioner of OR in industry, to his academic career at the University of Maryland. In Sections 6-9, we try to reflect the multiple facets of Saul's career as an OR professional, covering his roles as a scholar-teacher, expositor, practitioner, involved citizen, ambassador, and chronicler of OR. We hope that the rich tapestry that emerges shows why Saul commands the respect of his peers and his profession.

2 Portrait of a Young Mathematician

Saul Irving Gass was born on February 28, 1926 in Chelsea, Massachusetts, to Louis and Bertha Gass, who had emigrated from Russia around 1914. His father

sold life insurance but was a writer as well: He wrote a column in Yiddish in a Boston newspaper and also ran a radio show. Saul was the second of two children: his brother Gerald A. Gass was born October 21, 1922.

Apart from a few months at the Pratt Elementary School in Chelsea, Saul's schooling was done in Boston. He attended the Quincy E. Dickerman Elementary School in Roxbury, the Phillips E. Brooks School for the sixth grade, and continued at the Patrick T. Campbell Junior High School for grades seven through nine. Saul went through high school at the Roxbury Memorial High School for Boys and graduated in 1943. During high school, his favorite subjects were mathematics and physics. He also took navigation and aeronautics in his senior year, and did well in military drill.

Saul started Northeastern University in 1943 and spent a full year in college before joining the army. His army training started on March 17, 1944 at Fort Devens, Massachusetts and continued at Camp Blandig, Florida and Camp Shelby, Mississippi, where he trained as a machine gunner. On January 10, 1945, Saul sailed from New York City to Le Havre, where he camped for a month. In the final weeks of World War II, Saul was part of the 65th division of the U.S. Army. This division continued to move east during April, took Linz on May 4, and stopped on the west bank of the Enns River in Austria on V-E Day, May 8, 1945; the Russian army was on the east bank. Saul was then stationed at Pfaffenhoffen in Germany for the remainder of his stay in Europe. His two years of German in high school came in handy. Saul sailed back home in April 1946 and was discharged on May 23, 1946 at Fort Devens.

After his military discharge, Saul planned to get married and hoped to resume his university studies. He had met his future spouse when he was 15, and had courted her just before going to Camp Shelby. He married Trudy Candler on June 30, 1946 in Los Angeles, and they moved back to Boston. Saul re-enrolled at Northeastern but soon transferred to Boston University (in January 1947) to major in education and mathematics, with the intention of becoming a high school teacher. He graduated with a bachelor's degree in education (with a mathematics major) in June 1949. Saul's interest in mathematics had led him to take extra courses in this subject. Therefore, it only took him an additional summer to earn his masters in mathematics in August 1949. Not having found his semester of student teaching inspiring, Saul decided against a career as a high school teacher and looked for another job.

In November 1949, Saul was offered a job as a mathematician with the U.S. Air Force (as a GS-7) and joined the Aberdeen Bombing Mission (ABM) in Los Angeles. This civilian Air Force group performed ballistics analysis for bombs. It was led by Grace Harris and had eleven members (nine women and two men). According to Saul [88], the main task was to analyze "*photographic plates and high-speed camera film of high-altitude aircraft and bomb drops that took place at Edwards Air Force Base, north of Los Angeles in the desert country.*" Saul continues:

> At ABM, we read the plates and film by eye on a Mann Comparator, recorded the results by hand, and processed the readouts on the Marchant and Monroe

> desk calculators -- the old-fashioned way! I did become deeply involved in bomb ballistic work and was given the task of investigating new machine readers that could automatically record the positions of both the aircraft and bomb images and punch the results on IBM cards [88].

The Los Angeles group sent its results to the Aberdeen Proving Grounds in Aberdeen, Maryland, where bombing tables were developed. Saul visited this location and was offered a position. He declined the offer, but his work at ABM exposed Saul to the need for accuracy, constant checking, and data validation.

3 From Project SCOOP to Project Mercury

Saul was next offered a job as a GS-9 (at $5,060 a year) in the Directorate of Management Analysis of the Air Force in the Pentagon. Saul was a young father when he moved the family of three to Washington D.C. from Los Angeles in the first week of January 1952. Ronald S. Gass was born on June 3, 1951, to be followed by Joyce A. Gass (born June 22, 1955).

At the Pentagon, Saul joined Project Scoop (Scientific Computation of Optimal Programs). This Pentagon-based research program of the U.S. Air Force was formed in June 1947, and its official designation of Project SCOOP came in October 1948. This is where applied linear programming started. Saul has aptly dubbed it "*the first linear-programming shoppe*" and stressed its historical significance:

> All of us in OR are indebted to project SCOOP. The linear-programming model, the simplex method, the first computer-based solution of LP problems, much of the theory of linear and mathematical programming, the basic computational theory of linear programming, and the extension of LP to industry and business all stemmed, wholly or in part, from the research and developments of Project SCOOP. [81]

When Saul arrived at the Pentagon, the Directorate was headed by the economist Marshall Wood, and its scientist was George Dantzig. The main objective of Project SCOOP was to plan the requirements for air force programs. As Wood and Dantzig explained it in 1949:

> Programming, or program planning, may be defined as the construction of a schedule of actions by means of which an economy, organization or other complex of activities may move from one defined state to another, or from a defined state towards some specifically defined objective. Such a schedule implies, and should explicitly prescribe, the resources and the goods and services utilized, consumed, or produced in the accomplishment of the programmed actions. [150]

An example of an organization requiring such planning was the Air Force. A typical programming exercise was to construct a time-phased plan of requirements of materials for supporting a specific war scenario. Within Project SCOOP, the word "programming" was used in the specific military sense; computer programs were barely known and called codes at that time. As Dantzig has put it:

> The military refer to their various plans or proposed schedules of training, logistical supply, and deployment of combat units as a *program*. When I first analyzed the Air Force planning problem and saw that it could be formulated as a system of linear inequalities, I called my paper *Programming in a Linear Structure*. [12]

At the core of the Project was Dantzig and Wood's approach to modeling the economy or organization based on Dantzig's mathematical statement of the LP problem. The model used a triangular or rectangular technology matrix to specify the requirements and their interrelationships. This extended the Leontief input-output model from the triangular case (where no optimization was required) to the rectangular case where one could optimize an objective function using the LP structure. With the formulation of these planning models, the members of Project SCOOP came to realize the power of the LP model. While Dantzig, Alex Orden, and others were developing the key algorithmic procedures for the simplex method, the computational challenges of the task also came into clearer focus. With keen foresight, Wood and Dantzig identified the promise of the new technology:

> To compute programs rapidly with such a mathematical model, it is proposed that all necessary information and instructions be systematically classified and stored on magnetized tapes in the "memory" of a large scale digital electronic computer. It will then be possible, we believe, through the use of mathematical techniques now being developed to determine the program which will maximize the accomplishment of our objectives within those stated resource limitations. [150]

In the Mathematical Formulation Branch, Saul worked on the formulation and solution of Air Force problems and also developed and tested new procedures for solving LP structures. He recalls his entry into this dynamic and heady research environment.

> I was assigned to the Mathematical Formulation Branch. Walter Jacobs, a mathematician, was branch chief. He introduced me to linear programming by suggesting that I read reprints of Dantzig's three seminal papers .... Even though I was a fairly recent mathematics graduate, the concepts and ideas described in these papers were new to me and rather complex. What does the uninitiated make of such things as zero-sum games and the solving of hundreds of equations in hundreds of variables, especially in pre-computer days? Fortunately, I had a course in numerical calculus and knew something about Gaussian elimination and how to solve (3x3) systems of equations. [88]

The role of Project Scoop in advancing the use of computers is an important point made by Saul.

> Project SCOOP was responsible for much of the federal government's early involvement in computers, especially the efforts of the National Bureau of Standards (NBS). NBS received over $1 million from the Air Force, and used these funds to build the Standards Eastern Automatic Computer, the SEAC. [81]

The SEAC machine was located at the National Bureau of Standards (NBS) in Washington, D.C. Saul drove the problems from the Pentagon to the NBS campus on Van Ness Street. Early computational tests by Alex Orden on the NBS SEAC compared the simplex method with other approaches (relaxation and fictitious play), and typical results are reported in [81]. Project SCOOP also led to the installation of the second production unit of the UNIVAC machine in April 1952, formally turned over to the U.S. Air Force in June 1952. The simplex code for this machine was written by the Air Force's Mathematical Computation branch led by Emil Schell. The UNIVAC could handle LP problems of dimensions up to 250x500. Saul solved LP problems on this computer and describes it as follows:

> The UNIVAC had more than 5,000 vacuum tubes and could do about 2.000 additions or subtractions per second. It had an internal acoustical mercury-delay line memory of 1,000 12–character words... Its external memory consisted of 8 magnetic tapes that could read or write at the rate of 1,000 words a second. The UNIVAC, although a clunker by today's standards, was [a] great improvement over desk calculators. It was always exciting (and chilling) to walk into the special air-conditioned, room-sized cabinet that held the mercury delay-line memory tubes. [88]

George Dantzig left the Pentagon in June 1952 for the RAND Corporation. By 1955, Project SCOOP was starting to wind down, and research funds were being cut in the government. But Project SCOOP had assembled a remarkable network of researchers.

In addition to Marshall Wood and George Dantzig, its members included Saul Gass, Murray Geisler, Leon Goldstein, Walter Jacobs, Julian L. Holley, George O'Brien, Alex Orden, and Emil D. Schell. The summer students of the group included Phillip Wolfe from Princeton (Albert Tucker's student) and Tom Saaty, who was finishing his Ph.D. at Yale. This group also worked closely with the Washington-based national Applied Mathematics Laboratories that included Sam Alexander, John Curtiss, Alan Hoffman, Henry Antosiewicz, Peter Henrici, John Todd, and Olga Taussky-Todd. Also associated with the group were T. S. Motzkin and George Forsythe from the Institute for Numerical Analysis, the west-coast research arm of NBS, and Princeton's John von Neumann, Albert Tucker and his students Harold Kuhn and David Gale, Abraham Charnes and William Cooper from Carnegie-Mellon University, and Isidor Heller from Washington University. Listing this remarkable group, Saul remarks, "*What in the hell was I doing amongst*

that bunch of heavies?" [65]. In fact, this group exposed Saul to the wave front of operations research just after the beginning. Project SCOOP also ran two symposia of great historical importance on linear programming in 1951 and 1955. Saul attended the second seminar and gave a paper on finding first feasible solutions in LP [24].

Saul left the Project in May 1955 to join IBM as an Applied Science Representative. The job advertisement for this position required a degree in mathematics or engineering and exposure to "automated computing equipment, or system design and methods." He was hired along with other new sales trainees who were new college graduates and went to the standard three-week sales training class in Endicott, New York, where he sang IBM songs from the IBM songbook! Saul was assigned to the Washington commercial sales office located at 1111 Connecticut Avenue. His job was to help the salesman selling and installing IBM computers. The IBM 701-704 series of machines were just out and, later, Saul was also trained to program the IBM 650.

Saul's next employer was CEIR, the Corporation for Economic and Industrial Research, a consulting services company at which Saul had helped install an IBM 650 machine. William Orchard-Hayes was an early hire of this firm. Saul was approached by Jack Moshman to build up the OR group at CEIR. He joined the firm in 1959 as Director of the Operations Research Branch. However, his tenure in this position was cut short by the expansion of the space program, which led to an offer for Saul to return to IBM.

Saul rejoined IBM in 1960 as Manager of the Simulation Group of the Project Mercury Man-in-Space Program. He was responsible for the development of a full range of real-time simulation procedures used to validate the computational and data flow equipment system that IBM developed for Project Mercury. The key task for IBM was to calculate the orbit of the space capsule based on radar telemetry data collected from various tracking stations across the globe. This data was processed at two IBM 7090 computers located at the Goddard Space Center. IBM had to conduct the necessary analysis and develop the computer programs, run a duplexed computing center, and operate an engineering and communications subsystem that enabled the flight controllers to monitor all phases of a Project Mercury mission. Saul's initial assignment at Goddard was to dry run the computer programs that computed the orbit with simulated data, which he describes as follows:

> We simulated radar data from the world-wide tracking stations and ran the programs in real-time by entering the timed data into teletype machines connected to the computers... By this time, IBM was also given the responsibility of analyzing lift-off radar data to predict whether the space capsule would go into a successful orbit.... We simulated that phase, plus predicting when to fire the retro-rockets to bring the capsule back to the earth and the splash point. Our computer-based system was the first real-time decision-making system with a man-in-the-loop. [97]

The first U.S. manned-capsule sub-orbital flight occurred on May 5, 1961 with Alan Shepard. Just a few days before (on May 1), Saul was appointed manager of IBM's Project Mercury. Saul's recollection of this event conveys the atmosphere:

> I recall the scene just before his [Shepard's] lift-off: the now spruced-up Goddard computer room with its duplexed (A and B) computers, the side-by-side plotboards, that would, hopefully, trace calculated launch track over the already inked-in nominal track, and the output console with its switch that enabled the computed output to come from either the A or B computer. Systems operations manager Al Pietrasanta and I manned the switching console. The crowds of NASA and IBM VIPs that gathered were kept back by a set of stanchions and ropes. We felt like gladiators who would be fed to the lions if something went wrong. All went well... [85]

Saul also went to Cape Canaveral to watch the launches for all the manned orbital flights. There, he supervised a team of engineers responsible for data transmission from the Cape to Goddard and the running of the control center charts as well as launch and orbital plot boards. From the VIP grandstand at Cape Canaveral, Saul watched John Glenn's liftoff on February 20, 1962 in the first U.S. manned orbital flight. He then rushed inside Mercury Control center to watch the tracking plotboards. The success of Glenn's historic flight brought top management attention to Project Mercury within IBM. The chief scientist for IBM paid Saul and his team a visit and questioned them on system reliability and testing. A week after the flight, Saul briefed the board of IBM directors in New York and received a standing ovation.

The computer-based activities of Project Mercury, which engaged a team of over 100 scientists, programmers, and engineers, paved the way for future manned-space projects. It also foreshadowed the critical role of real-time computing in making manned space flight a reality [28]. It was therefore not only one of the largest projects Saul had managed within industry, it was also a highly intense assignment in uncharted territories. Project Mercury was also a great practicum for project management.

> We all learned from Project Mercury. For me, it had to do with the integration of people, computers, programs, and real-world necessities. I learned the importance of bringing control to an ever-changing set of tasks and the need to impose a rigorous verification and validation process. I learned how pressure, contractual responsibilities, and finances can be honed to meet a project's goals given that those involved communicate, cooperate, and compromise in a manner that does not seriously distort their objectives.... I had to negotiate what it meant to turn over a real-time, man-rated system that had never been developed. How does one demonstrate to tough-minded NASA managers and engineers that one met system specifications when such specifications, although set down in the past, had to be constantly changed to meet the realities of the present? [85]

As Project Mercury came to a close, the NASA space program moved to Houston, Texas. Saul had a principal role in preparing IBM's proposal for developing the Houston Real-Time Computing Center, but did not want to move to Houston himself.

4 Back to School and Return to OR Practice

Ever since his days at SCOOP, Saul had shown a continuing interest in taking courses related to his areas of interest. In 1953-54, he took the two semester course that Albert Tucker and Harold Kuhn taught at the American University on Thursday nights. Kuhn and Tucker were involved with a research project at George Washington University. One or the other would come down to Washington, D.C. for this purpose and teach the night course. Later, Saul formally enrolled in the doctoral program in mathematics at American University and took the following classes: "Linear Programming and Game Theory" from Alex Orden; "Methods of Operations Research" from Joe McCloskey (known for his early work in the history of OR); "Linear Programming" from Alan J. Hoffman; and "Numerical Analysis" from Peter Henrici. Saul also took two computer courses from the NBS staff.

In September 1963, Saul decided to take advantage of IBM's resident graduate fellowship program that allowed IBM employees to go back to school on a two-year leave with full pay. The IBM fellowship allowed Saul to go to the school of his choice, and American University would have been a convenient choice. However, Saul had also maintained contact with George Dantzig, who had joined Berkeley in 1960 to head its OR department. Saul chose Berkeley for his doctoral studies, and the Gass family drove to California in August 1963.

At Berkeley, Dantzig taught the linear programming course using notes that formed the basis of his famous text, *Linear Programming and Extensions*, which came out later in 1963. Because of Saul's substantial background in LP, he was not allowed to take this course for credit, but he audited it (in fact, Saul never took any courses for credit with his mentor.) There is an amusing story about how uneasy his classmates felt when they found out that Saul had already written the first text on LP!

His doctoral course work at Berkeley included network flows, discrete programming, theory of probability and statistics (I and II), mathematical economics (I and II), inventory theory, nonlinear programming, applied stochastic processes, and advanced statistical inference. In addition, he had to pass minor examinations in economics and probability/statistics, and two languages (French and German). Saul recalls the networks class given by Bob Oliver as one of his best courses at Berkeley. His best instructor was Elizabeth Scott, who taught probability and statistics. The Dantzig and Gass families socialized and often went to dinner together. He also socialized with Bob Oliver and Bill Jewell [99].

For his oral defense, Saul had to present a paper from outside his field. He was given a paper on busy periods in queueing written by Takacs. At the end of the presentation, he was asked only one question: Dantzig asked, "What's a

convolution?" When Saul started to look for a dissertation topic, Dantzig suggested that he contact Roy Harvey at Esso. Harvey had a large-scale LP problem, for which Saul devised a novel decomposition scheme and algorithm. This was the dualplex algorithm, which constituted his Ph.D. dissertation [29]. Saul completed his Ph.D. in summer 1965. He was one of the earliest doctoral students of George Dantzig. Before him, Richard Cottle had completed his Ph.D. in 1964. Other students of Dantzig who were Saul's contemporaries include Earl Bell, Mostafa el-Agizy, Ellis Johnson, Stepan Karamardian, and Richard van Slyke, all of whom earned their doctoral degrees in 1965.

Saul returned to IBM in the summer of 1965. IBM had already formed its Federal Systems Division with offices in Gaithersburg, Maryland. For the next decade, he was involved in projects. Saul was manager of Federal Civil Programs and responsible for applying information retrieval and other data procedures, advanced graphics techniques, and data analysis to urban problems. While most of his work at IBM did not have a heavy dose of OR or LP modeling, Saul did get a chance to apply OR thinking to urban problems as a full-time member of the Science and Technology Task Force of the President's Commission on Law Enforcement.

The Commission was created by President Lyndon Johnson in 1965, partly in reaction to the issue of "crime in the streets" that Barry Goldwater had raised in the 1964 election campaign. The Commission was mainly comprised of lawyers and sociologists. The Task Force was formed to augment the work of the Commission by bringing scientific thinking to bear on crime. The Task Force was led by Al Blumstein (see [1]), who recruited Saul to join in 1966. Other recruits were Richard Larson (who was completing his undergraduate degree in Electrical Engineering), Ron Christensen (a physicist and lawyer), the consultant Sue Johnson, the statistician Joe Navarro, and Jean Taylor. The Task Force was based at the Institute for Defense Analyses (IDA), where Al Blumstein and Jean Taylor worked at the time. Saul was responsible for developing the Task Force's approach to how science and technology can best serve police operations.

From 1969-70, Saul was Senior Vice-President of World Systems Laboratories, Inc. This was a Washington-based consulting firm with five key principals. He then joined Mathematica, the well-known OR and economics consulting firm headquartered at Princeton, New Jersey. Tibor Fabian was President, and Harold Kuhn and Oscar Morgenstern were on the board. Saul headed the Bethesda office of the firm and worked on several government projects. These included the development of an educational student aid model for the U.S. Department of Education; the establishment and analysis of an educational data bank for the Appalachian Regional Commission; the development for the Corporation for Public Broadcasting of a procedure to evaluate the effectiveness of a series of telecasts on environmental issues; consulting to the systems group of the Chief of Naval Operations; the development of a simulation model of the dispatch/patrol functions of the Washington Metropolitan Police Department; the development of operational planning materials for the National Center for Educational Statistics common core data system; and principal investigator on the NSF project to evaluate policy-related research in police protection. He also organized an

unclassified symposium for the CIA that focused on techniques for analyzing intelligence information.

One of the projects Saul undertook at Mathematica was a contract from the Environmental Protection Agency to conduct a survey of modeling in the non-military governmental area. This resulted in the volume *A Guide to Models in Governmental Planning and Operations,* which Saul edited along with Roger L. Sisson [124]. This book was privately published by its editors in 1975. A total of 2,000 copies were printed and distributed out of Saul's basement. The name of the publisher – Sauger Books – indicates this upon closer inspection [97]. Saul's chapter in this volume devoted 45 pages to a review of modeling efforts in law enforcement and criminal justice and lists 103 references. Among the models reviewed were patrol beat design, the police emergency response system, and court models [34].

5 Academic Home Found

Long before starting his professorial career, Saul had revealed his academic bent. In addition to doing research and writing the first text on LP, he also taught the subject regularly at the US Department of Agriculture (USDA), American University, and George Washington University. In 1973 and 1974, he taught an evening course in operations research for the business department of the University of Maryland. In 1973, the business administration department became the College of Business and Management. Rudy P. Lamone was appointed dean. Dean Lamone, who had received a Ph.D. in OR from the University of North Carolina, was interested in building a high-quality OR department. He persuaded Saul to join the University of Maryland in September 1975 to become the chair of the Management Science and Statistics Department. Saul was to spend the next 26 years at the university. As he puts it in an interview: "I had found a home." [97]

Saul was Professor and Chairman of the Faculty of Management Science and Statistics from 1975 to 1979. The faculty of this department already included Gus Widhelm and Stan Fromovitz in OR. Saul lost no time in building up the department. He hired Larry Bodin at the full professor rank and Bruce Golden as a fresh assistant professor in 1976. In the next two years, he recruited Frank Alt, Mike Ball, and Arjang Assad as assistant professors. While Saul stepped down as department chair in 1979, he remained closely involved with its development and growth for the next 22 years. During these years, he taught LP and OR subjects at doctoral, MBA/MS, and undergraduate levels. He also supervised doctoral and masters students and sat on numerous thesis and dissertation committees. Saul was the dissertation advisor of eight students, which we list in chronological order: Stephen Shao and Jeffrey Sohl (1983), Rakesh Vohra (1985), Noel Bryson (1988), Hiren Trivedi (1990), Anito Joseph and Pablo Zafra (1993), and Pallabi Guha Roy (1999).

As a respected citizen of the University of Maryland, Saul was asked to participate in important committees, especially when sensitive issues needed to be tackled. Two of his contributions are still in effect at the Robert H. Smith School

of Business: As the chairperson for the committee that designed the faculty pre-tenure and post-tenure reviews, Saul prepared the "Gass Report," which continues to govern the review process at the Smith School. Also, as a repeat member of the annual merit review committee, Saul suggested a framework for ranking the faculty reports and is reputed to have advocated the use of AHP for this task!

Saul garnered many university honors in the course of his academic career. He was designated a University of Maryland Distinguished Scholar-Teacher in 1998. He held the Westinghouse Professorship during 1983-1992, and was appointed Dean's Lifetime Achievement Professor in 2000. In July 2001, he was appointed Professor Emeritus.

Saul's professional honors also make for a long list. He served as the 25th President of the Operations Research Society of America (ORSA) in 1976-77 and was elected an INFORMS Fellow in 2002. In 1991, he was awarded the Kimball Medal for service to ORSA and the profession, followed by the INFORMS Expository Writing Award in 1997. Saul received the 1996 Jacinto Steinhardt Memorial Award of the Military Operations Research Society (MORS) for outstanding contributions to military operations research.

Saul served as President of Omega Rho, the international operations research honor society in 1985-1986, Vice President for international activities of the Institute for Operations Research and the Management Sciences (INFORMS), and Vice President for the North American Operations Research Region of the IFORS Administrative Committee. He was general chairman of the 1988 TIMS/ORSA meeting held in Washington. Saul was invited to deliver the plenary address at the May 1996 INFORMS Washington, DC conference and the San Francisco TIMS/ORSA meeting in 1984. In 1994, Saul gave the third E. Leonard Arnoff Memorial Lecture at the University of Cincinnati.

While at the University of Maryland, Saul maintained a close relationship with the National Institute of Standards and Technology (NIST – formerly the National Bureau of Standards or NBS). Aside from several consulting projects, Saul organized a number of conferences through NIST. The most recent such conference, which he organized and co-chaired with Al Jones, was a workshop on Supply Chain Management practice and research co-sponsored by the NIST, the National Science Foundation, and the Robert H. Smith School of Business, University of Maryland. Selected papers from this conference were collected in a special issue of *Information Systems Frontiers* [135].

6 The OR Scholar

In this section, we focus on Saul's contributions to OR methodology and applications. We review the methodological contributions under two broad categories: the theory of LP and its extensions and decision making tools. We also describe Saul's work as a builder and user of OR models developed for specific applications. His work on the modeling process will be covered in the next section.

6.1 Linear Programming and Extensions

Saul's first major and lasting contribution to the theory of linear programming was his work on the parametric objective function with Thomas Saaty [121, 122, 144]. The idea arose within Project SCOOP in 1952. Walter Jacobs introduced Saul to this problem in the context of production smoothing. In production planning problems, one faces the dual objectives of minimizing the monthly fluctuations in production *and* the inventory carrying costs. By attaching weights to the two objectives, one can express this as a single objective LP, where the key parameter reflects the ratio of the cost of a unit increase in output to the cost of carrying one unit of inventory (see [51, pp. 353-358]).

Saul first solved this transformed parametric problem by hand on some test problems using a modified simplex tableau. When Thomas Saaty joined SCOOP in the summer of 1952, he and Saul worked out the details with some help from Leon Goldstein and Alan Hoffman. This resulted in three well-known papers by Gass and Saaty that address the parametric problem

$$\begin{aligned} &\text{Min } (c + \lambda d)\, x \\ &\text{s.t.} \quad Ax = b, \\ &\qquad x \geq 0, \end{aligned}$$

where c and d reflect the cost vectors of the two conflicting objectives and λ is the parameter to be varied. The first paper [144] described the case where some components of the new costs $c + \lambda d$ are linear functions of λ and described how each basic optimal solution remains optimal for a range of values that defines a closed interval of values for λ. Moreover, by tracking these contiguous intervals, a finite number of optimal solutions can be obtained to account for all possible values of λ. Gass and Saaty [121] specified the algorithmic steps of the parametric problem that is described in Saul's [26] and other LP texts. Gass and Saaty also published a companion paper [122] that considered parametric programming when the costs depended on two independent parameters λ_1 and λ_2. For a historical view of the development and impact of parametric programming and Saul's seminal role, see Gal's historical reviews [20, 21].

The conceptual link between parametric programming and Multi-Objective Linear Programming (MOLP) problems was present in Saul's mind. Recalling how Walter Jacobs gave him the task of solving the parametric version, Saul writes:

> This led to the parametric programming algorithm in which we balanced off two competing linear objectives by the single parameter. We recognized that the scheme could be generalized to more objectives and described a multi-parameter approach where two objectives meant one parameter, three objectives, two parameters and so on. [83]

Saul's interest in goal programming dates back to the mid 1970s. He traces its development to his consulting days at Mathematica, where he encountered goal

programming features in most planning models for the government [83] and used this model in his work on personnel planning models [9, 67]. Such models are typically large-scale goal programming problems involving thousands of deviation variables. Saul has addressed the problem of setting the objective function weights for these variables in a consistent way [53, 57]. Also related to this area is Saul's critique of preemptive goal programming [57].

In a string of papers, Saul and Moshe Dror have proposed an interactive approach to MOLP [14, 15, 16, 107]. In a more recent paper with Pallabi Roy, Saul proposes the use of a compromise hypersphere to rank efficient extreme point solutions to an MOLP [120]. In a vector maximization problem with q objectives, the utopian solution that simultaneously maximizes all q objectives is typically unavailable. There are various ways of defining a *compromise solution* that attempts to strike a balance among the q objectives. The approach proposed by Gass and Roy is to find an annulus of minimum width that encloses the efficient solutions in the q-dimensional objective space. We discuss this below, when we review Saul's work on fitting circles (or spheres) to a collection of points [130].

Encounters with Degeneracy: Among the topics of interest to Saul within LP are degeneracy and cycling. Describing his "encounters with degeneracy," Saul [71] traces the roots of this interest to his days at Project SCOOP and provides a proof for an unpublished result due to Leon Goldstein. The result states that a basic feasible solution with a single degeneracy cannot cause a cycle and appeared as an exercise in the first edition of Saul's text [26] (problem 2, p. 70). Saul made the useful distinction between *classical cycling* and *computer cycling* [40]. The former is what we usually see in textbooks of LP and refers to cycling when the computations of the simplex algorithm are carried out accurately without round-off errors. Computer cycling refers to cycling encountered when solving the LP with a computer system, and hence is "*a function of the mathematical programming system being used*" [40]. Neither concept logically implies the other: A problem can exhibit classical cycling but not computer cycling. In fact, recently Gass and Vinjamuri [128] performed a test of 11 LP problems that cycle classically but are solved successfully by three popular simplex solvers. Saul also studied the effect of degeneracy in solving transportation problems [63]. It is worth noting that Dantzig [11] and Magnanti and Orlin [140] have shown how parametric programming can be used to avoid cycling, thus bringing together two strands in Saul's research program.

Other early work of Saul's on linear programming includes his short note on an initial, feasible solution for a linear program [24], and a transportation-based algorithm for meeting scheduled manhour requirements for a set of projects [25]. Saul's next major algorithmic work in linear optimization was in the area of large-scale systems and constituted his doctoral dissertation.

The Dualplex Method: Saul introduced and studied the Dualplex algorithm in his doctoral dissertation [29, 32]. The problem was motivated by an application with a staircase structure in the technology matrix. The problem may be conceptualized as involving K stages, coupled by a set of "tie-in" constraints. In a decomposition

approach, it is natural to handle the coupling constraints in such a way as to fully exploit the block structure of the individual stages. Saul used this approach focusing on the dual problem for handling the complicating variables. The following account, which is necessarily very brief, is meant to outline the approach.

Assume that the LP problem can be put into the form:

$$\begin{aligned} &\text{Min} && c\,x + d\,y \\ \text{s.t.} \quad & && w + A\,x + B\,y = b \qquad (1) \\ & && w, x, y \geq 0 \end{aligned}$$

where w is the vector of basic variables, and the non-basic variables (x, y) are partitioned in such a way that the activities corresponding to the x variables appear in only one stage, while those for y appear in more than one stage and serve to couple the various stages. The matrix A therefore exhibits block-diagonal structure: it consists of the submatrices A_1 , A_2 , ..., A_K on the diagonal and zeros elsewhere. Correspondingly, we partition x by stage as $x = (x_1 , x_2 , \ldots, x_K)$. We rewrite constraint (1) as follows for stage k ($k = 1, \ldots, K$).

$$w_k + A_k\, x_k = b_k - B_k\, y^0 \ . \qquad (2)$$

Now suppose that the dual problem

$$\text{Min } \pi\, b \quad \text{s.t.} \quad \pi\, B \geq d \ , \ \pi \geq 0 \qquad (3)$$

is solved to obtain the multipliers $\pi^0 = (\pi_1{}^0 , \pi_2{}^0 , \ldots, \pi_K{}^0)$ and an associated set of values y^0. We construct the solution $(w^0, x^0 = 0, y^0)$ to the original problem and test it for optimality by computing the reduced costs $c_k - \pi_k{}^0 A_k$. If this vector has any strictly positive components, then the optimality condition does not hold and the solution can be improved by introducing a variable with positive reduced cost from stage k into the basis. A key attractive feature of the dualplex method is that up to K such variables (one for each stage) can be pivoted into the basis at the same time. Once this is accomplished, the form given in (1) can be recovered with new values for the quantities A_k, B, b, c_k, and d, and the dual system (3) can be resolved again. For this approach to be effective, the previous basis for (3) should be exploited in full. This is where most of the technical details of the procedure have to be worked out. These are presented in [29], together with proof of the correctness of the overall algorithm.

Apart from the proceedings article [32], Saul did not publish anything on the dualplex method for over 20 years. He returned to this algorithm in his joint work with doctoral students at the University of Maryland. The algorithm is applied to the knapsack problem with generalized upper bounding (GUB) constraints and the piecewise linear approximation to the separable convex minimization problem [123]. In [4], the method is proposed for solving discrete stochastic linear programs.

Fitting Circles to Data Points: How does one fit a circle to a given set of points in the plane? This simple geometric question led to a collaboration among Saul, Christoph Witzgall, and Howard Harary [130]. To clearly define the optimization problem, one must specify a measure of fit. The criterion Saul and his coworkers use for fit is a modified Chebychev minimax criterion. To define it, suppose that n points P_i with coordinates (x_i , y_i), $i = 1, \ldots, n$ are fixed on the plane. Consider a circle of radius r_0 with its center located at (x_0 , y_0). Let r_i denote the distance from the center to the point P_i with a radial line. If $r_i = r_0$, then P_i lies on the circle constructed and a perfect fit obtains. Otherwise, we consider the difference of the squared radial distances $| r_i^2 - r_0^2 |$ as the error term. The objective is therefore

$$\text{Min Max}_{i=1, \ldots, n} | r_i^2 - r_0^2 | ,$$

or

$$\text{Min Max}_{i=1, \ldots, n} | (x_i - x_0)^2 + (y_i - y_0)^2 - r_0^2 | .$$

Since the (x_i , y_i) are fixed for $i = 1, \ldots, n$, the decision variables are (x_0 , y_0) and r_0. These are the variables for which the outer minimization problem is solved. This problem can be formulated as a linear program [129]. Extensive computational experience with this problem is reported in [130], and these results show that the squared formulation above provides a very good approximation to the minimax criterion without squares, that is Min Max $| r_i - r_0 |$. The paper by Witzgall in this volume provides further details and developments of this approximation problem.

As described in [130], this problem arose from a manufacturing setting where coordinate measuring machines calibrate circular features of manufactured parts. Once the center of the fitted circle is determined, one can use it to define two concentric circles that are the inscribed and circumscribed circles for the n points. This defines an annulus of minimum width that contains all n points and is checked against tolerances in the manufacturing setting. Saul has adapted this basic idea to his research on MOLP [120]. The basic idea of a minimum-width annulus is used, now extended from circles in the plane to hyperspheres in $\mathbf{R}^q$. The given points correspond to vector-valued objectives of the available efficient points. Gass and Roy [120] describe how to use the minimum-width annulus to rank efficient points in the MOLP.

Modified Fictitious Play: Early in his career, Saul was exposed to the elegant results that established the equivalence of matrix games and LP. He has since expressed his appreciation for this result (see for example, p. 60 of [97]). He has also described how Alex Orden used the fictitious play method as one of the solution methods for LP in his computational work on the SEAC [81, 97]. Saul returned to this early interest in his work with Pablo Zafra [131,132]. The method of fictitious play was proposed by Brown [3] as an iterative method where the players take turns and at each play, each player chooses "*the optimum pure strategy against the mixture represented by all the opponent's past plays*" [3]. The convergence of this method, while guaranteed theoretically, can be extremely slow.

To speed up convergence, Gass and Zafra [131] modify the fictitious play method by using a simple restart method based on the existing lower and upper bounds on the value of the game. The modified fictitious play (MFP) procedure is then tested to assess its value as a computational procedure for finding initial solutions to LP problems. This MFP start is tested in combination with simplex-based or interior point LP solvers and found to be useful for certain types of linear programs. Gass, Zafra, and Qiu [132] show that MFP can achieve the same accuracy as regular fictitious play with 40-fold reductions in the number of iterations and significant savings in computational times.

In several publications, Saul and his coworkers Anito Joseph and K. Osei-Bryson have studied integer programming problems and proposed bounds, heuristics, and problem generation issues. This stream of research has resulted in several publications [136, 137, 138, 139] that detail the procedures and provide computational experience.

6.2 The Analytic Hierarchy Process (AHP)

Saul has maintained an interest in AHP [18, 143] since its inception. As he mentions in [95], he taught the subject at the University of Maryland since the early 1980s [2] and was the first author to discuss it in a text (chapter 24 of [50]). His expository article on AHP co-authored with Ernie Forman [18] covers the basics and describes 26 applications of AHP in different domains. In [95], Saul reviews the arguments of critics who fault AHP for violating the axioms of multiattribute utility theory (MAUT) and concludes that AHP is an alternative approach to MAUT that offers its own advantages.

Saul has also proposed the use of AHP in several new areas. In large-scale goal programming models for personnel planning, the decision maker must assign weights to thousands of deviation variables in the objective that measure the differences from the desired target values. In one realistic example, 13,900 such variables are present. Saul proposes the use of AHP to determine the priorities and weights in such models [53]. With Sara Torrence, Saul [127] has described the use of the AHP to rate the complexity of conferences for the Public Affairs Division of the National Institute of Science and Technology (NIST), which plans and coordinates approximately 100 conferences a year. The complexity of each conference is determined based on two AHP hierarchies used to provide ratings for both the time required and the level of expertise needed from the staff assigned to the conference. Saul has also proposed using AHP to provide a numerical rating for model accreditation [70].

Saul's most interesting methodological contribution to the AHP appears in his work with Tamás Rapcsák. Consider the pairwise comparison $n \times n$ matrix $A = (a_{ij})$ as constructed in Saaty's AHP and denote its maximum eigenvalue by λ_{max} [143]. It is well known that the vector $\mathbf{w}$ of weights or priorities of the alternatives in Saaty's AHP is the normalized right eigenvector of A corresponding to the largest eigenvalue λ_{max}, so that $A\ \mathbf{w} = \lambda_{max}\ \mathbf{w}$ and the sum of components of $\mathbf{w}$ is unity. Gass and Rapcsák [118, 119] propose a different derivation of the weights that makes use of the *singular value decomposition* (SVD) of A. If the $n \times$

n matrix A has rank k, the singular value representation is given by $A = U\ D\ V^T$, where D is a $k \times k$ diagonal matrix and U and V are matrices that satisfy $U^TU = I$ and $V\ V^T = I$. The diagonal entries of D, which must be positive are called singular values; we arrange them so that $\alpha_1 \geq \alpha_2 \geq \ldots \geq \alpha_k$. Let $\mathbf{u}_i$ and $\mathbf{v}_i$ be the columns of U and V for $i = 1,\ldots, k$. The main result from singular value decomposition is that for any $h \leq k$, the matrix

$$A_{[h]} = \sum_{i=1}^{h} \alpha_i \quad u_i \quad v_i^T \qquad (4)$$

constructed from the h largest singular values of A is the best rank h approximation to the original matrix A in the sense of minimizing the Frobenius (matrix) norm $\| A - G \|$ among all matrices G of rank h or less. (The Frobenius norm of a matrix is simply the Euclidean norm of the vector defined by the n^2 entries of the matrix.)

If the matrix A is *consistent* (i.e., $a_{ik} = a_{ij}\ a_{jk}$ for all i, j, k), then it is of rank 1 and its SVD consists of only the first term or diad of (4). In AHP, however, the pairwise comparison matrix A is typically not consistent. Gass and Rapcsák [119] argue that the use of the matrix $A_{[1]}$ is still justified as the best rank-one approximation to A and provide a simple expression for the weights $\mathbf{w}$ in terms of the singular vectors $\mathbf{u}_1$ and $\mathbf{v}_1$ corresponding to the largest singular value α_1 and prove that these weights are identical to the standard AHP weights if A is consistent. In addition, they provide a measure of consistency based on the Frobenius norm. In this way, they provide an alternative approach to the computation of priorities with certain appealing theoretical underpinnings, but acknowledge that further investigation of the relative merits of this approach is warranted.

Saul has also investigated the occurrence of intransitive relationships in pairwise comparison matrices [84]. An intransitive relation occurs when we have $A > B > C > A$ where $A > B$ means that item A is preferred to B. Given a pairwise comparison matrix $A = (a_{ij})$, we can construct a binary $n \times n$ preference matrix $P = (p_{ij})$ with zeros on the diagonal, where $p_{ij} = 1$ if item i is preferred to item j, in which case $p_{ji} = 0$. The preference graph corresponding to P is a directed graph with an arc from i to j if $p_{ij} = 1$.

Clearly, an intransitive relation among three items corresponds to a cycle of length three or a *triad.* It can be shown that triads are present if any cycles of higher order exist. Thus, searching for triads as the elemental representatives of intransitivity is appropriate. In [84], Saul reviews the known results for the number of triads in preference graphs and proposes using a transshipment algorithm to identify triads.

One might ask how frequently intransitive relations occur in the context of the AHP. The empirical study of Gass and Standard [125] sheds some light on this question. They consider 384 instances of positive reciprocal comparison matrices taken from real-world AHP studies and consider the distribution of the numbers in the basic 1-9 comparison scale. They also investigate the incidence of intransitive relations in such matrices.

6.3 OR Applications

In this section, we review Saul's work as a builder and user of OR models developed for public or governmental planning and analysis.

Design of Patrol Beats: One of the earliest applied modeling studies conducted by Saul was his work on the design of patrol beats: Given k patrol units to be assigned during a patrol shift, the problem is to design k patrol beats so as to equalize workload [30]. The area assigned to each patrol beat must have desirable geometric properties: it must be contiguous and avoid shapes (such as long and narrow) that would make patrolling difficult or inefficient. As such, the problem has some similarities to the well-known political districting problem. The computation of the workload for a given area requires historical data on the incidence of various types of incidents by census tract. The problem is solved with a heuristic developed for political districting that uses a transportation subproblem. The computations of the design study were carried out using 1966 data for the City of Cleveland, which had 58 beats and 205 census tracts.

Traffic Safety: Together with Nancy David and Robert Cronin, Saul [105] conducted a simulation study on the role played by citizens band (CB) radio in traffic safety for the National Highway Traffic Safety Administration (NHTSA). The main task of the simulation model is to compare the response times for the case where the highway patrol system has CB installed to the case when CBs are not used. The simulation model tracks the movement of vehicles and the occurrence of incidents along a single road link with two-way traffic. It is assumed that a vehicle with CB can report an incident to a patrol car with CB, or to the highway patrol base station. As an example of the summary statistics produced, the model indicated that the time to accident notification with the use of CB is 47% of the time without CB, and that the time to actual response to the accident with CB was 75% of the time to response without CB. Similar percentages were developed for other incidents that compromise traffic safety (e.g., roadway hazards).

In another project for the NHTSA, Gass, David, and Levy [106] describe a system flow model that captures the common structure of NHTSA traffic safety demonstration projects. This common structure involves the flow of people through a network of stages where the stages correspond to activities of a traffic safety program. For example, the flow might track convicted DWI offenders and a program might be a special treatment for a subset of such offenders. The stages can be modeled as queueing subsystems with sojourn times and delays. The model tracks the flows between successive stages, exit flows, and recidivism. The authors describe the model details and illustrate the structure with a case study on DWI probation.

Personnel Planning Models: Another class of models that can also be viewed as flow models are personnel planning systems. The basic personnel planning problem is to "*determine the* number *of personnel and their skills that best meets*

the future operational requirements of an enterprise" [67]. In his review paper [67], Saul describes the basic structure of these models as integrating a Markov transition submodel with a network (or LP) optimization model. Saul's work in this area has focused on military manpower planning. In 1984, Saul was engaged as a consultant by Sigma Systems to extend and implement the Army Manpower Long Range Planning System (MLRPS). This model was originally designed in 1982-83 and moved to the Army computer in Washington, D.C. in August 1983.

U.S. Army personnel models are "*designed to determine by grade, skill, and/or years-of-service, the number of soldiers (or officers the Army can put in the field... so as to meet manpower goal over 7- to 20-year planning horizons.*" [104]. The flow model uses a Markov chain to project the flow of personnel accounting for transitions due to promotions, separations, skill migrations, and so forth. This is coupled with an optimization model that analyzes how the Army can best achieve a future force structure starting from its present state. The optimization goal programming (GP) model is driven by the goals the user inputs for skill-grade combinations by time period. The model allows the user to address such policy questions as "*what accessions, promotion and separation rates will best transform the Army's present force structure to the one that will be required 20 years from now?*" [104]

Personnel planning models are generally large-scale [49]. A typical model may have 7 grade classifications and 33 skill identifiers, which gives 7 x 33 = 231 grade-skill combinations to track over T years. If T = 10, the linear goal programming formulation of the MLRPS has 9,060 equations and 28,730 variables, of which 13,900 are deviation variables associated with the target goals. Saul's work with the MLRPS model led him to investigate procedures for setting the objective function weights for large-scale GP models [53] and exploiting the network structure within large GPs [79]. Saul has also studied the Accession Supply Costing and Requirements (ASCAR) model that determines the recruiting needs of an all volunteer Armed Forces to meet target goals [9].

In reviewing Saul's work on modeling, it is also important to mention his edited volumes where one can find comprehensive reviews of models in a given domain. The Gass and Sisson edited volume [124] is an early and important example of this category. This volume was prepared for the Environmental Protection Agency to review governmental, urban, and public modeling efforts. Environmental topics such as waste management and air pollution are discussed. The book has a much broader range, however, covering such areas as justice, education, health, and transportation planning. Saul's work on expert systems in emergency management [102, 103] is another example.

7 The Expositor of OR

Many individuals in the field of OR first came to know Saul through his linear programming text. Others may have first read one of Saul's articles in *Interfaces*. Today, we recognize Saul as one of the master expositors of OR and the recipient

of the INFORMS Expository Writing Award for 1997. How did this dimension of Saul's career develop? To find the roots, *cherchez Project SCOOP*!

Saul wrote his first expository paper on LP at the request of Walter Jacobs at Project SCOOP. Jacobs asked Saul to prepare a non-technical pamphlet on LP aimed at an Air Force audience. The result was a guide entitled, "The Application of Linear Programming Techniques to Air Force Problems," published as an internal USAF paper in December 1954 [23]. This guide is 27 pages long and makes use of Air Force applications. For example, the transportation problem is introduced as shipping units of equipment required for the B-47 from three source Air Force bases to five destinations. Saul also describes a production smoothing problem and points out the conflicting objectives of minimizing the fluctuations in output and the excess monthly surpluses. The guide ends with a discussion of the contracts awards problem that arises when the Air Force procures items from civilian sources. Each bidder offers the unit at a known price and must observe minimum and maximum limits on its production quantity. The Air Force must award contracts in such a way as to minimize the total cost. This problem reappears in Saul's texts [26, 50].

The origin of Saul's *Linear Programming* text goes back to the introductory LP course George Dantzig taught at the US Department of Agriculture (USDA) Graduate School in 1950. When Dantzig left for RAND, George O'Brien – a part-time researcher at Project SCOOP – taught the course in Fall 1952. O'Brien, in turn, asked Saul to teach the course when he left the Washington area for Detroit. He then prepared a set of notes for the USDA course that he also taught to his coworkers at the Pentagon. This led to the notion of writing a text on LP. The material in Gass' text was originally prepared for a semester-long course with 16 lectures, each 2.5 hours in duration. In 1954-55, Saul sent a couple of chapters to a few publishers, until McGraw-Hill signed a contract to publish the text. He finished writing the text during his first year at IBM (1955-56). During this time, Saul went to the Library of Congress on Saturdays to do his writing. Appropriately, the dedication to Trudy Gass mentions "lost weekends."

Published in 1958, this was the first text on LP. In fact, the only book-length accounts of LP preceding Saul's text were the monographs by Charnes and Cooper [7] and Vajda [148]. The first edition of *Linear Programming* was translated into Russian in 1961 and was the first book on the subject in the Russian language. Later, the text was also translated into Spanish, Polish, Czechoslovakian, Japanese, and Greek, providing the first text on LP in these languages as well. Subsequent editions appeared in 1964, 1969, 1975, and 1985. The fifth edition of this text [51] was also reprinted as a Dover paperback in 2003.

The first edition of the text comprised 12 chapters. The first chapter introduced LP and gave three examples: the transportation problem, activity-analysis, and the diet problem. Chapter 2 provided basic mathematical background on linear algebra and convexity. Chapters 3-9 developed the theory and algorithmic procedures for LP. The last three chapters covered the transportation problem, general LP applications, and the relation between LP and the theory of matrix games. The book also included a bibliography of LP applications organized by application

domain and based on the Riley-Gass bibliography [141]. A comparison between the first and fifth editions is given in Table 1.

Table 1. Comparison of the contents of the first and fifth editions of *Linear Programming.*

	First Edition	*Fifth Edition*
Publication year	1958	1995
Number of pages	223	532
Length (in pages) of chapter on LP applications (Ch. 11)	32	77
Number of references	102	736
Number of works in the applications bibliography	91	433
Number of exercises	43	380

Saul also published another major book and work of reference in May 1958. This was the compendium entitled *Linear Programming and Associated Techniques: A Comprehensive Bibliography on Linear, Nonlinear, and Dynamic Programming*, published by the Operations Research Office (ORO) at the Johns Hopkins University. This bibliography cited over 1000 items that included articles, books, monographs, conference proceedings, and theses. The cut off was set for June 1957 (Saul's LP text is cited as *in press*!). The book was divided into three parts. Part I provided the introduction and Part II covered the general theory. Part III of the bibliography, which was devoted to applications, accounted for 290 pages of the total 613 pages of the book.

Saul's co-author, Vera Riley, was Staff Bibliographer at ORO. While she had no training in OR or mathematics, she located the key references and was adept at reading a document and abstracting its main contributions. To expedite the work, Riley and Gass used the abstract provided by the work cited. However, when a cited work did not have an abstract that could be readily used (as with books), Saul wrote the annotation for the bibliography [99].

In 1961, Saul published a survey of recent developments in LP as a chapter in the prestigious series *Advances in Computers* [27]. The topics reviewed in this chapter include decomposition methods, LP applications, stochastic LP, integer programming, and nonlinear programming. Most of the research and references cited in this work were published during 1957-1960. Thus, in a sense, the work updated the Riley-Gass compendium [141] in the domain of LP. Of special historical interest is the section on LP codes and procedures, which comprises close to 40 of the 83 pages of the chapter. In this section, the LP codes are arranged by computing machine, and then by problem type. The reader will find many codes where the problem size is constrained by $m \leq 100$ or $m\,n \leq 1500$.

Among Saul's books, the primer entitled *An Illustrated Guide to Linear Programming* occupies a unique position [31]. The novelty lies in the illustrations and the lively and humorous expositional style. Not counting graphs and tables, there are 25 illustrations involving graphic characters conjured by the illustrator

William F. McWilliam to bring the concepts to life. Previously, Williams' primer on game theory [149] had made use of illustrations. But in the opinion of this reader, the illustrations in [31] are more elaborate and support the discussion in the text more effectively.

The first half of the *Illustrated Guide* is devoted to problem formulation. The reader encounters the classical LP problems: the transportation problem and traveling salesman problems, the caterer and trim problems, personnel assignment and activity analysis. The caterer problem, for example, is described as a management science consulting assignment addressing the issue of dirty napkins for the Mad Hatter's tea parties. The consultant's report appears in a different typescript under the title: *An Analytical Analysis of Interactive Activities as related to the Economics of Functional Gatherings: A preliminary linear programming model of a tea-party subsystem.* Clearly, Saul was writing as one who knew the consulting profession all too well!

Despite the introductory nature of the text, Saul takes the time to explain the historical origins of some of the military applications. A historical footnote explains that the caterer problem disguised the procurement of repaired aircraft engines. Another footnote recounts the origin of the contract-awards problem, which goes back to Saul's Project SCOOP days:

> The relationship between the contract-awards problem and the transportation problem was first discovered and exploited by mathematicians of the U.S. National Bureau of Standards (NBS), working with personnel from the Philadelphia quartermaster depot. Prior to the advent of linear programming, each problem was solved by submitting the bids to a series of evaluations conducted by different analysts. When no change could be found, the successful bidders would be announced, and everyone would hope for the best. When the first operational tests were conducted using the NBS computer, the SEAC, the quartermaster group continued to solve the problems with their analysts in order to build up the necessary confidence in the linear-programming approach. The computer solutions were always better than, or at least as good as, the quartermaster solutions. [31, pp. 107-108].

In 1985, Saul published a new textbook entitled *Decision Making, Models and Algorithms* [50]. The objective of this text was to expose the mathematically inclined undergraduate student to the decision sciences as used in OR/MS. Of the 22 chapters of this text, 16 are taken from the *Illustrated Guide.* Thus, this new text is partly a reincarnation of Saul's 1970 classic, complete with the illustrations that enlivened the *Guide*. However, *Decision Making* also has much new material, including chapters on modeling, decision trees, and the AHP. Of special pedagogical interest is the material collected in the five chapters that close the five parts of the book. These chapters reflect Saul's skills of exposition and his passion for explaining the methods and practice in an accessible way.

Saul's major book-length work in the 1990s is the *Encyclopedia of Operations Research and Management Science,* which he edited with Carl M. Harris [110, 111]. Saul was approached by the publisher to embark upon this project in 1994.

Saul invited Carl Harris to join him, and it took them two years to complete the *Encyclopedia*; the first edition was published in 1996.

The second edition [111], which appeared in 2001, is over 900 pages long and lists 1136 entries. The editors commissioned over 200 expository articles that run 3-6 pages in length. One of the noteworthy features of this work is its list of contributors. The editors succeeded in soliciting contributions from leading authorities in OR and management science. Having read a large number of the longer entries, I believe that the *Encyclopedia* is also an expository triumph: special care is taken to make the articles accessible and tutorial in nature. The book is dedicated to Carl Harris and Hugh Miser, *in memoriam*.

This section has mainly focused on Saul's book-length publications. Saul is also well-known for his expository or review articles. These works are not discussed in this paper but the following examples may be cited: AHP and rating [2, 18, 87, 95]; LP and extensions [52, 63, 71]. Modeling and applications [47, 59, 67, 68], and teaching OR [112, 2, 80].

8 Manager of the Modeling Process

8.1 Scrutinizing the Modeling Process

For close to thirty years, Saul has studied not just the technical contents of models, but the "*total environment of decision-making with models,*" which he has chosen to call "*managing the modeling process*" [56, 66]. Saul's interest in the area was an outgrowth of a series of studies sponsored by the National Institute of Standards and Technology (NIST), The Department of Energy (DOE), and the U.S. General Accounting Office (GAO). Saul also drew upon his experience with OR studies and projects in the public sector during his consulting years.

In the early 1970s, Saul became involved in the evaluation of police protection models as a principal investigator of the National Science Foundation grant to Mathematica. His task was to evaluate 50 models based on the documentation provided and the accompanying research papers. Saul was faced with the challenge of developing a framework for evaluation. Reflecting on the project, he writes:

> Our first problem was to determine what was meant by evaluation and how you do it. The literature gave little guidance, so we developed our own evaluative process. I quickly learned that analysts do not document, cannot, or will not write well, do not state their modeling assumptions, are unclear as to their data sources, maybe perform sensitivity analyses but do not tell you, and so on.... For me, the outcome was to start thinking about the problem of what we really mean by good modeling practices and implementation, and how do you evaluate model-based projects. [56]

Saul's next investigation of the modeling process was to help the U.S. General Accounting Office (GAO) develop a procedure for the evaluation of complex

models. The historical background on the events that led to this effort is provided in [73] and [81] and is briefly summarized here.

In the mid-1970s, Federal agencies were using computer-based policy models to support the executive branch in addressing policy and legislative issues in such areas as energy and welfare. Congress wanted to know how these models worked and directed the U.S. Government Accounting Office (GAO) to evaluate the Project Independence Evaluation System (PIES) model developed by the Federal Energy Agency. In 1976, GAO published its review of PIES [145], marking the first federal effort in model evaluation. Soon afterwards, a provocative GAO report [146] suggested that model development activities can be improved in a host of different ways. In the face of such concerns, the question of model evaluation arose naturally: How does one conduct an evaluation of a complex computer-based policy model? Moreover, the evaluation of PIES caused GAO to recognize the need for a more formal process of model evaluation. It is important to recall that PIES was the first experience of GAO in assessing complex policy models and that in 1975-1976, the prior literature on this subject was "*basically non-existent.*" [56].

In April 1977, Saul organized a workshop at NBS [39] to expose model developers and users to the GAO report [146] and to elicit reactions from this group. The main controversial topic was GAO's gated approach with five phases, each of which could potentially arrest further development of the model. Also discussed was Saul's questionnaire-based approach to model assessment [35, 36]. The discussions at the workshop reinforced Saul's belief in the need for a life-cycle approach to computer model development and documentation. Saul's framework defined thirteen phases starting from embryonic initiation, through development and validation, to implementation, maintenance, evaluation, and documentation [37, 38]. Along with Bruce Thompson, Saul also continued his involvement in shaping the framework GAO was preparing to guide model evaluation. The final document appeared in 1979 as the publication *Guidelines for Model Evaluation* [147]. As Saul [81] put it, "*the importance of this document is that if project personnel were constrained to furnish the evaluative information suggested by the 'Guidelines', the success rate of our modeling projects would increase greatly.*" The GAO Guidelines identified five major criteria for model evaluation: Documentation, validity (theoretical, data, and operational), computer model verification, maintainability (updating and review), and usability. To bring its message to the attention of OR analysts, Gass and Thompson published a summary of the *Guidelines* in *Operation Research* [126].

Saul's perspective on managing the modeling process has drawn from his long involvement with the assessment of energy models. He has cited the case of DOE's Energy Information Administration (EIA) on several occasions for having "*put into place a quality assurance program that, except for nuclear and some defense activities, is probably the most advanced in the world*" [81]. Saul was an early contributor to the EIA quality assurance and evaluation program. In the mid-1970s, he was funded by EIA through NBS, along with parallel efforts at MIT (Edward Kuh and David Wood), Berkeley (David Freeman), and Stanford. Saul's work in this area resulted in his first publications on the subject [35, 36], followed by a review of different assessment frameworks for energy models [41, 43, 44].

8.2 Evaluation of Complex Models

In one of his first publications on the modeling process, Saul defines model assessment or evaluation as "*the process by which interested parties, who were not involved in the model's origins, development and implementation, can assess the model's results in terms of its structure and data inputs so as to determine, with some level of confidence, whether or not the results can be used in decision making*" [35]. This process is necessary because (a) the ultimate decision makers are typically far removed from the modeling process, (b) the applicability of the model to a new domain of use must be assessed, and (c) because the complex interactions between the model assumptions, inputs, structure, and results are best understood through a formal independent evaluation process. Saul states the basic philosophy behind model assessment and its relevance to the analyst in the following terms:

> In the development of any model, the model developers should assume that their model will be subjected to an independent assessment, and thus, it behooves them to impose the explicit discipline of an assessment methodology and related project management controls to ensure that acceptable and correct professional procedures are imposed on all aspects of the model structure and use. [19]
>
> The OR analyst should be particularly concerned with the model evaluation process because the outcome of an evaluation is really a recommendation to the decision maker whether or not to use the output of the analyst's model. [46]

Saul's approach to model evaluation appears in a string of publications [35, 36, 114, 126, 42, 43, 47] culminating in the comprehensive account provided in the *Operations Research* feature article [46]. Saul [43] reviews several frameworks for the assessment of energy policy models. His framework in this paper builds on and synthesizes previous proposals for evaluating computer-based and policy models. He adopts the distinction made by Fishman and Kiviat [17] between model verification and validity. Briefly stated, for computer-based models, "*the process of demonstrating that the computer program 'runs as intended' is termed verification.*" Model validation "*attempts to establish how closely the model mirrors the perceived reality of the model user/developer team.*" [46]

Verification must ensure that the computer program describes the model accurately, has no errors of logic, and that it runs reliably. As such, verification is chiefly the concern of programmers and must precede the validation process. Following the GAO framework [126] that Saul helped develop, validation comprises three major components: technical, operational, and dynamic validity. "*Technical validity requires the identification of all model assumptions, including those dealing with data requirements and sources, and their divergences from perceived reality.*" It can further be subdivided into model validity, data validity, mathematical, and predictive validity. The last element examines errors reflecting

the difference between actual and predicted outcomes. Operational validity assesses "*the importance of errors found under technical validity*" and includes sensitivity analysis. Finally, dynamic validity examines "*how the model will be maintained during its life cycle so that it will continue to be an acceptable representation of the real system.*" [46]

Saul also uses the assessment framework to propose a questionnaire for the utility evaluation of a complex model [36]. The questionnaire uses thirteen criteria to assess model utility as follows:

1. Computer program documentation
2. Model documentation
3. Computer program consistency and accuracy
4. Overall computer program verification
5. Mathematical and logical description
6. Technical validity validation
7. Operational validity
8. Dynamic validity
9. Training
10. Dissemination
11. Usability
12. Program efficiency
13. Overall model

In another study related to the utility of complex models, Saul [142] sought expert opinion on how the utility of large-scale models can be improved by surveying. The study surveyed 39 modelers with diverse affiliations and application domains to obtain their reactions to a list of 18 key ideas for model improvement (e.g., model user training, verification and validation plan, and a model ongoing review panel).

The preceding brief account suggests that model evaluation and assessment is a multi-faceted topic with fluid boundaries and subject to different interpretations. Saul's path through this maze is marked by a level-headed consideration of past frameworks, avoidance of rigid requirements, and pragmatism. Based on [58, 60] and other sources cited above, I would summarize Saul's overall view of model assessment as follows:

- there is no single "correct" assessment process that works for all models
- the assessment process can and should be formally developed
- the assessment process should be informed by the valuable past experience with complex models, their use, and implementation
- the assessment framework should remain attentive to the intended use of the model and what is practical or helpful to the ultimate users
- effective project management requires that a life-cycle view of model development and assessment be adopted
- assessment should not be an afterthought; it is expensive and ineffective to conduct *ex post*

- ongoing evaluation is a mark of professional quality and results in an improved development process overall
- the OR analyst should maintain a keen interest in the processes of evaluation

Based on his long-standing scrutiny of complex models, Saul is emphatic on the importance of model evaluation. Here is how he brings the point home.

> From our perspective, good modeling practice assumes that someone, someday, will knock at your door and shout: "Open up, it's the modeling police force! We're here to take your model down to Headquarters, question it to see what makes it tick, and plug it into the lie detector machine to determine if it tells the truth. You have three minutes to call your analyst, programmer, and lawyer". ... If you were the poor soul behind the door, what would you wish you had done so that the evaluation of your model comes out excellent? [81]

8.3 Model Documentation

The preceding account of model evaluation indicates the pivotal role played by model documentation in the process of independent or third-party assessment of policy models. Often, the analyst conducting such an assessment must often rely on the documentation provided for the model as the primary source of information. Moreover, "*it is clear that the lack of documentation hinders both the dissemination and use of models, especially models for analyzing policy decisions*" [113]. This makes model documentation a critical element of model usability. As Saul puts it:

> The evaluator must review the model documentation that will be made available to the decision makers, computer system operators, analysts, other possible users, and interested researchers. This review must establish the value of every type of documentation in terms of their specific purposes. The understanding of a complex model rests on the clarity of its documentation and its associated training program. [48]

> Documentation of a model is extremely important. A management plan for the production of documentation is an essential component of the management plan for a model... This must be understood by sponsors and developers and should be addressed (and funded) in contractual arrangements. [113]

In 1979, Saul and his coworkers conducted an assessment of a major modeling effort at the Energy Information Administration (EIA) of DOE. They evaluated a large-scale model called MOGSMS that provided forecasts of oil and gas supply and proposed guidelines to DOE on the organization of model documents [113]. Their hierarchical view of documentation has four levels and calls for twelve

specific documents. The levels are (1) rote operation of the model, (2) model use, (3) model maintenance, and (4) model assessment. The relevant documents for each level are described and placed on a four-shelf "model bookcase" for visual effect.

In 1979, Saul prepared a report for the Institute for Computer Sciences and Technology of the National Bureau of Standards in which he proposed guidelines for the documentation of computer-based decision models of DOE. This report prescribed a total of 13 documents that corresponded to different phases of the life cycle of the model [38]. Saul has also summarized his guidelines for model documentation in a paper in *Interfaces* [48]. The outcome of the documentation process is four manuals directed at the analyst, the user, the programmer, and the manager. For each of these, Saul suggests a detailed table of contents.

Saul summarized the guiding theme of his approach to documentation as follows:

> The documentation of computer-based models should provide specific and detailed information organized and presented in a manner that will satisfy the needs of each segment of the model's audience. The audience consists of the model's sponsors and users (possibly non-technically oriented), the model's analysts, programmers, and computer operators; other users, analysts, programmers and computer operators; and independent model evaluators. [48]

8.4 Model Credibility

Suppose that a set of procedures for model assessment is agreed upon and followed. What is the ultimate result of following such procedures? Gass and Parikh [116, 117] discuss the notions of baseline analysis and credible analysis for complex policy models. Ultimately, a judgment is made on "*the measure of confidence or credibility a user should give to the model's results*" [46]. Saul reminds us of the subjective element in the notion of model confidence:

> We emphasize that model confidence is not an attribute of a model, but is an information-based opinion or judgment of a particular user for a given decision environment,... The level of confidence may vary from user to user because of differences in application requirements, and subjective judgmental preferences... Model confidence is subjective, each decision maker internalizes the available information by means of an imprecise algorithm for evaluating model confidence. [46]

While there is no precise procedure for determining this level of confidence, Gass and Joel [114] propose an approach and provide a set of criteria that an independent observer can use to rate (score) the model under evaluation. Building on this work, Saul suggests a modified approach for rating model accreditation based on the AHP. The criteria for the AHP are: specifications, verification, validation,

pedigree, configuration management, usability, and documentation, each with its own list of subcriteria.

Together with several coworkers from the U.S. GAO, Saul also adapted the model assessment framework to evaluate simulation models and tested the approach on three Army simulation models [19]. The three simulation models were used in two antiaircraft defense systems—the portable surface-to-air Stinger missile and the division air defense gun (DIVAD). The objective of this study "*was to demonstrate that it is possible to systematically collect and analyze information about a simulation that would permit an assessment of the credibility of that simulation to be made*" and to identify the steps DOD had taken to ensure this credibility [19].

For each of the three simulation models evaluated (ADAGE, Carmonette, and COMO III), Fossett et al. [19] describe the type and key features of the model as well as its intended use. They assess the credibility of each model based on theory, model design, input data, correspondence to the real world, organizational support structures, documentation for users, and reports of strengths and weaknesses. Fosset et al. [19] acknowledge that "*of all OR methodologies, simulation has always been in the forefront of verification and validation research and stressed the necessity of applying verification and validation tests.*" The assessment framework they provide suggests that OR analysts may have focused on just a subset of the full list of assessment components that relate to validation and verification (e.g., input data quality, statistical representation, and sensitivity). One contribution of this study is to bring the challenges of assessment into sharper focus for a class of models that form an integral part of the OR toolkit.

Based on his assessment work on simulation models, Saul was invited to participate in the model validation, verification, and accreditation (VV&A) working group of MORS- the Military OR Society. A brief account of model accreditation and the relevant activities of this group is given in [81]. Other model assessment studies in which Saul participated include the evaluation of the Dynamic General Equilibrium Model, a dynamic model of the U.S. economy used to understand the impact of mobilization on the economy [6]. Finally, Saul has also participated in the development of guidelines for emergency management models [5].

9 The OR Professional, Statesman, and Ambassador

In 1996, in the course of a discussion of academics and practice as the two sides of OR, Saul summarized his career as follows.

> My career spans both sides of the OR Equation. For 25 years I worked as a mathematician and an OR analyst, and I directed OR for a couple of consulting firms; I had the good fortune to work in some important areas: Linear programming, the first man-in-space program, criminal justice, energy modeling and manpower systems. I have been there. For the last 20 years, I have been an academic. I have been there, too! [78].

If we add 10 years to the numbers cited, it remains an accurate description of Saul's 55-year long career in OR. Saul's experience base and his blend of academics and practice place him in a select group among the leading members of our profession. As a scholar and teacher, Saul has published extensively. His publications include six books, 80 journal articles, nine book chapters, at least 12 proceedings articles, 4 edited volumes and reports, and a dozen other occasional pieces. The counts given are based on the references given in this article and should be viewed as lower bounds. While reasonably comprehensive, this list leaves out some items and does not account for multiple editions, translations, and so forth.

Saul has written about the "three P's of OR: Practice, Process, and Professionalism" [58]. Remarkably, another facet of Saul's contributions emerges when one considers his involvement in the profession as an active and vocal citizen. This section focuses on this facet and its relation to the three P's of OR.

9.1 The OR Professional as Concerned Citizen

Saul's care and concern for the well-being of the OR profession has been a theme expressed in numerous articles, plenary addresses, position statements and other outlets. Chief among these is his collection of "Model World" articles in *Interfaces*. In these articles, Saul expressed his views on the full range of issues in our profession in an accessible and direct style. Twelve such articles have appeared during the period 1990-2005 addressing such issues as models and the modeling process [62, 64, 65, 112], the profession and its ethics [77, 82, 115], the history of OR [65, 89, 101], the meaning of such rankings as the *Business Week* survey of business schools [87], and the publication practices of OR/MS journals [78].

In addition to "Model World," Saul has expressed his views on the profession in various plenary addresses, invited articles, or special feature articles. A review of his writings on these occasions shows that Saul returns to several key themes. I take these to be Saul's professional *credo*. Using Saul's article for the President's Symposium in *Operations Research* [54] as a starting point, I will try to summarize these themes.

OR: its past, present and future. Saul has presented his view of "*what OR has done and should do*" in [54]. Saul states that "*OR has arrived,*" in the sense that it has successfully solved a wide variety of operational problems that the OR pioneers took as their challenge. While OR should take pride in its accomplishments, the operational outlook is too "*restrictive a view of what OR can and should be.*" In the future, OR will have to tackle complex organizational problems, meet the challenge of decision-making in real time, and contribute to policy analysis [82].

Professional identity and the OR process: On several different occasions, Saul has sounded a note of alarm on how the status of OR as a profession runs the risk of being diluted as diverse professions adopt and use OR techniques [54, 64]. For Saul, this risk is heightened by the identification of OR by its methods alone and calls for a shift in perspective.

> Unless we bring structure to our profession, that is develop and promulgate guidelines of technical conduct — what may be termed the process of Operations Research — we will soon not have much to claim as our specific professional domain. In the eyes of technical compatriots and those in management who support us, we must be something other than methods. [58]

> There has not been any improvement in what I term "managing the modeling process." I think it is a serious concern; one, that if not corrected, will see the profession of OR/MS being subsumed by the other disciplines that are managed better (or at least appear to be better managed). [56]

Professional Quality: As a profession, OR must not only develop "*high standards, guidelines and management procedures,*" but also demonstrate to management that OR professionals can manage a complex project and complete it "*to specifications, on time and within budget.*" Such questions as "*how does the total modeling effort fit into the life-cycle of the project?*" must be addressed [54]. For Saul, the broad issues of model evaluation discussed in Section 8 are an integral part of this integration.

The Science of Modeling: Saul has challenged the OR profession to develop *a science of modeling.* [54, 58]. He commends computer science for making the study of what it does an object of inquiry and believes that OR should do the same. "*Unlike those working in the field of computer science, we have not attracted the behavioral and psychological researchers to the study of the practice and implementation of our professional endeavors.*" [66]

Ethics: Saul has articulated his position on ethics in MS/OR clearly:

> It is essential for the future well-being of the MS/OR profession that its ethical concerns and problems be investigated and discussed in a more demanding fashion, especially by our professional societies. The problems will not go away... MS/OR needs a code of ethics and standards for professional practice. [77]

> The reader can rightly ask, "Do modelers need a code of ethics?" My answer is a resounding "Yes!" ... That most important ethical concern [that we as OR analysts face] is adhering to proper professional practices in the development and reporting of our work. [72]

Saul's interest in professional ethics and code of conduct in OR is of long standing. He took a keen interest in the *Guidelines for Professional Practice* that appeared in *Operations Research* in September 1971 [8], participated in the debate surrounding the *Guidelines* [33], and reviewed its history for the next generations of OR professionals [66, 69]. Saul supports the original 1971 *Guidelines* and has regularly cited this effort in his writings on ethics [66, 76]. He suggests that the guidelines are most appropriate when combined with the framework he offered in

[56]. In two articles [66, 77], Saul reviews other examples of codes of professional conduct and gives an interesting account of his presentation to (and interactions with) the doctoral students at ORSA doctoral colloquia on the subject of professional ethics.

In his historical review, Saul reminds us that the *Guidelines* were written in the aftermath of a controversial study of the antiballistic missile system (ABM). Because the techniques of the study drew upon military OR, the Operations Research Society of America was drawn in. This brings up the thorny issue of the appropriate code of conduct for analysts when they find their models entangled within an adversarial process. Saul reminds us that "*models have a life of their own and are subject to use beyond the original intent; they are becoming tools of advocacy*" [66]. Saul presents two interesting case studies of this situation in [66, 69, 73].

9.2 Ambassador to the OR World

Globalization may be very much in vogue nowadays, but for Saul it has been second-nature to view MS/OR as a "*problem-solving discipline that crosses geographical boundaries*" [62] . His extensive international travels have served to build personal and professional relations, for sure, but Saul has also actively sought to learn about OR practice in other countries and has chosen to spend his sabbaticals outside the United States. A brief review of his international contacts will indicate Saul's passion for worldwide travel.

In 1977, under the US/USSR Academy of Sciences Exchange Program, Saul spent one month in the Soviet Union visiting scientific institutions in Moscow, Kiev, Tbilisi, and Novosibirsk. This led to his participation in other US-USSR exchanges. One such exchange was a workshop on large-scale optimization, in which researchers from the Central Economic Mathematical Institute of the USSR Academy of Sciences (Moscow) participated. Saul organized this workshop in College Park in January 1990. In June 1990, Saul led the U.S. delegation that visited Moscow. In May 1993, he led a delegation of 15 Operations Research professionals to Russia and Hungary under the auspices of the Citizen Ambassador Program.

In 1981, Saul was invited to present a seminar on Operations Research in the U.S. for the International Institute of Technology in Tokyo. In 1983, he lectured in England on model validation and assessment under the auspices of the British Science and Engineering Council and the Operational Research Society. In 1988, Saul spent his sabbatical at the University of Anadolu, Eskisehir, Turkey and lectured at universities in Ankara and Istanbul. In 1985, the Beijing University of Iron and Steel Technology and the Chinese Academy of Sciences invited Saul to visit China, where he lectured in Beijing, Xian, and Shanghai. In 1986, he was an invited lecturer at universities in Brisbane, Australia and Auckland, New Zealand. In 1988, he was a plenary speaker at the First International Symposium on the Analytic Hierarchy Process, Tianjin, China.

Saul spent two months at the Computer and Automation Research Institute, Hungarian Academy of Sciences as a Fulbright Scholar for 1995-96. He was an

invited speaker at the First (1995) International Conference on Multi-Objective Programming and Goal Programming (MOPGP) held in Southhampton, England [79]. He presented the first Abraham Charnes Lecture at the second MOPGP conference held in Torremolinos, Spain in May 1996 [83].

In 2001, Saul was selected for a five-year term as a Fulbright Senior Specialist by the Fulbright Program administered by the Council for International Exchange of Scholars and has visited several countries as part of this program. More recently, he was selected as the IFORS Distinguished Lecturer plenary speaker for the joint IFORS/EURO conference held in Reykjavik, Iceland, July 2-5, 2006.

In his capacity as an observer of the international practice of OR, Saul has discussed OR technology transfer to developing countries and the practice of OR around the world. In 1987, Saul was the keynote speaker at the 11th Triennial Conference of the International Federation of Operational Research Societies' (IFORS), held in Buenos Aires. At this meeting, Saul addressed MS/OR technology transfer and stated:

> I believe that each country must develop its own approach to the practice of MS/OR, a practice that fits within and is part of the country's cultural and managerial decision-making framework. A country's unique ethos must define its approach to the practice of MS/OR. [62]

In September 1990, Saul returned to this theme when he delivered the opening plenary address at the national meeting of the British Operational Research Society in Bangor, Wales [66]. He reviewed some of the reasons why projects involving OR technology transfer failed in developing countries. Commenting on the challenges faced by OR analysts in these projects, Saul wrote:

> Operations researchers in developing countries have a most difficult task. Their management wants results now; they want pay-offs similar to those acclaimed in the developed countries. But the successes in the developed countries were not easy to come by; they took 40 years of research and application, and an evolving management style before they bore fruit.... Operations researchers in developing countries know the science of OR, but they must develop their own culture of OR that can grow and bear fruit in their soil. The developing countries will benefit if they do diverge from the way OR is practiced in the developed countries. I think they can do a more efficient job of it by applying the discipline of a modeling life-cycle approach to their OR projects. [66]

9.3 The Chronicler of OR

In his reflections on the 50th anniversary of the founding of ORSA [89], Saul recalls the 71 persons who attended the founding meeting of ORSA on May 26 and 27, 1952 at the Arden House, Harriman, New York. Saul calls this group, together with a dozen pioneers who were not founding members, the first generation of OR professionals. Thirteen members of this group served as ORSA Presidents through

1974. This places Saul, who was President of ORSA in 1976, in the second generation. He remarks:

> One might say that starting in 1975, the administrative and organizational aspects of ORSA were managed by a new, second generation of OR professionals... My emphasis on the generational aspects of OR lets me segue into my main theme: the training and experiences of my generation of OR professionals were shaped by the early decades of remarkable scientific and management advances by the now somewhat forgotten first generation. [89]

The remarkable (some would call it explosive) achievements of these pioneers of OR is what Saul has tried to convey in his various historical writings. He believes that the insights of these founding OR scholars and professionals continue to be instructive today. This is one reason why I have included Saul's historical writings as part of his expository *opus*.

Saul's historical papers fall into two groups. The first group has to do with key periods of historical and scientific interest in which Saul was involved personally. Project SCOOP and the origins of LP is covered in [61, 65, 81, 88]. Saul's Project Mercury experience appears in [27, 85]. The Washington OR scene is described in [81]. Other works in this category are Saul's appreciations of George Dantzig [93, 96] and the memorial article Saul co-authored with Don Gross on their friend Carl Harris [109]. Saul has also written short biographies of George Dantzig, A. W. Tucker, and John von Neumann for the Hall of Fame section of the *International Transactions in Operational Research* [93, 94, 98].

The second group consists of historical pieces that cover specific OR/MS pioneers or a given period of the OR/MS profession. Saul recalls and reflects on the founding of ORSA [89], reviews the history of the definition of OR [101], and the early history of LP [61] as a response to the article by Robert Dorfman on the same subject [13].

The work of Saul with Susan Garille on the diet problem [22] offers an interesting blend of history and practice; one might refer to this work as old wine in a new bottle! Garille and Gass consider the original diet problem posed by Stigler in 1945: Given 77 foods and nine nutrients (including calories), find the diet that meets the recommended dietary allowances (RDAs) for the nutrients at minimum cost. Garille and Gass recount how Stigler's diet cost $39.93 using 1939 data, as compared to the optimal solution obtained with the simplex method in 1947 at a cost of $39.69. They go on to discuss the subsequent history for the diet problem and solve an updated version using 1998 prices and RDAs.

In 2002, Saul published several articles to commemorate the 50th anniversary of the founding of ORSA [88, 89, 90, 91, 92]. One of these was a timeline of key events in OR history [91]. It listed 241 events that captured the development of the field. Together, Saul and Arjang Assad expanded this list into a book with 417 annotated entries published in 2005 as *An Annotated Timeline of Operations Research: An Informal History* [100]. Saul dedicated the book to his grandchild, Arianna Gass.

Saul's humor and flair for showmanship have made OR history fun at national meetings of our profession. Starting in 1990, he organized and ran three Knowledge Bowls held in conjunction with national meetings of ORSA, TIMS and INFORMS. In these contests, two teams of prominent OR/MS citizens competed to answer questions that tested their OR/MS "cultural literacy." I was a judge at two of these events and witnessed the pleasure Saul took in preparing and pitching the questions. Sample questions appear in [92].

10 Conclusion

This article has reviewed Saul's career to date as an OR scholar, teacher, and professional. In the introduction to this paper, I indicated that Saul has been a leader in SPPP – Scholarship and the three P's of OR. On all four scores, through his extensive writings and his personal example, he has given us much to cherish and reflect upon.

Saul has taken his own advice on full documentation to heart by making a rich record of his career available to us. From this record, we can see the roots of much of Saul's research interests extending back to his first 10-15 years in OR. His interests in LP, multi-objective optimization, military OR, public sector planning, and policy models all go back to this period. From Saul's accounts of the early days of OR, it is evident that Saul's professional identity is nurtured and fortified by his remarkable exposure to the first generation of OR pioneers. To his credit, his awareness of the past has never compromised his forward-looking perspective. He has taken up new interests himself and challenged the OR profession to shake off restrictive outlooks, stake out new areas of interest, and establish standards of practice [54, 60, 82]. Nearly two decades ago, Saul was a member of the Committee on the Next Decade in OR [10]. In November 2006, at the INFORMS National meeting at Pittsburgh, Saul will chair a cluster on "The Great Problems of OR," in which prominent citizens of the OR community will try to delineate the challenges for the next generation of OR researchers. This is just one example of Saul's continuing alertness to the future course of OR.

As a committed and vocal citizen of the OR community, Saul has never shied away from expressing his opinion on key issues of the OR profession. (As his colleague, I can attest that he has done the same at the university.) When he is moved to state his opinion, he does so directly and forcefully, clearly stating his credentials and relevant background as needed. At times, Saul has gone against the grain of majority opinion [74, 75, 33, 112]. In such cases, his conduct sets an example of how professional debate can be carried out with equanimity and integrity. He is a capable organizer of professional gatherings and events and the perfect master of ceremonies. Finally, Saul has a strong sense of humor, which finds its way into his writings and enlivens his speeches.

I would like to close with a personal remark. I was hired by Saul 28 years ago and have been his colleague ever since. My appreciation of Saul's role as an OR professional has been monotonically increasing during this period. I remember clearly how, in the course of a brief exchange in the parking lot, I approached him

with the idea of expanding the timeline [91] into an annotated book. That was in Winter 2002, and we thought we could knock it off in less than a year. The book [100] took us close to 26 months to finish. During this period, as Saul's co-author, I had an opportunity to work closely with him and learn more about him. I remain grateful for this.

Five years ago, I was pleased to read Peter Horner's tribute to Saul Gass on the occasion of his 75th birthday [134]. Given Saul's "high-octane" energy level, the accomplishments listed in that tribute are now badly out of date! I suspect the same will be true of this article before too long. Nonetheless, in keeping with the occasion feted by this volume, I offer this review as a small tribute to a great colleague.

Acknowledgments

This article makes liberal use of the writings of Saul Gass. I am indebted to Saul for making some of his earlier papers available to me and for answering certain queries I put to him. I acknowledge my debt to the excellent interview with Saul published by MORS [97]. Given Saul's superior powers of exposition, I have chosen to quote him directly throughout the paper. Where references to his works are less specific, it is to avoid the excessive clutter that referencing by source and page would have caused. I thank Bruce Golden for inviting me to contribute to this volume and for his patience, flexibility, and helpful comments as this article grew much beyond its original scope and length. I am grateful to Ruth Zuba for her valuable suggestions and assistance in the final stages of the manuscript preparation.

References

1. A. Blumstein. Crime Modeling. *Operations Research*, 50:16-24, 2002.
2. L. Bodin and S.I. Gass. On Teaching the Analytic Hierarchy Process. *Computers & Operations Research*, 30:1487-1497, 2003.
3. G.W. Brown. Iterative Solution of Games by Fictitious Play. In *Activity Analysis of Production and Allocation,* pages 374-376. Tjalling C. Koopmans, ed., John Wiley & Sons, New York, 1951.
4. N.A. Bryson and S.I. Gass. Solving Discrete Stochastic Linear Programs with Simple Recourse by the Dualplex Algorithm. *Computers & Operations Research*, 21:11-17, 1994.
5. R.E. Chapman, S.I. Gass, J.J. Filliben, and C.M. Harris. Evaluating Emergency Management Models and Databases: A Suggested Approach. National Institute of Standards and Technology, NBSIR 88-3826, Gaithersburg, Maryland, 1988.
6. R.E. Chapman, C.M. Harris, and S.I. Gass. Analyzing the Economic Impacts of a Military Mobilization. In *Lecture Notes in Economics and Mathematical Systems, #332: Cost Analysis Applications of Economics and*

Operations Research, pages 353-386. T. R. Gulledge, Jr. and L. A. Litteral, eds., Springer-Verlag, New York, 1989.

7. A. Charnes and W.W. Cooper, *An Introduction to Linear Programming*, John Wiley & Sons, New York, 1953.
8. T.E. Caywood, H.M. Berger, J.H. Engel, J.F. Magee, H. Miser, and R.M. Thrall. Guidelines for the Practice of Operations Research. *Operations Research*, 19:1123-1258, 1971.
9. R.W. Collins, S.I. Gass, and E.E. Rosendahl. The ASCAR Model for Evaluating Military Manpower Policy. *Interfaces,* 13(3):44-53, 1983.
10. CONDOR (Committee on the Next Decade in OR). Operations Research: The Next Decade. *Operations Research*, 36:619-637, 1988.
11. G.B. Dantzig. Making Progress During a Stall in the Simplex Algorithm. *Linear Algebra and Its Applications*, 114/115:251-259, 1989.
12. G.B. Dantzig. Linear Programming. *Operations Research*, 50:42-47, 2002.
13. R. Dorfman. The Discovery of Linear Programming. *IEEE Annals of the History of Computing*, 6:283-295, 1984.
14. M. Dror and S.I. Gass. Interactive Scheme for MOLP Given Two Partial Orders. *Applied Mathematics and Computation*, 24:195-209, 1987.
15. M. Dror and S.I. Gass. A Case for Interactive Multiobjective Linear Programming - MOLP. In *Functional Analysis, Optimization and Mathematical Economics,* Memorial volume for Prof. L. V. Kantorovich, pages 222-234. L. Leifman, ed., American Mathematical Society, Oxford University Press, 1990.
16. M. Dror, S.I. Gass, and A. Yellin. Experiments with an Interactive Procedure for MOLP Given Weak Orders on Variables and Objectives. *European Journal of Operational Research*, 34:78-85, 1988.
17. G.S. Fishman and P. J. Kiviat. *Digital Computer Simulation: Statistical Considerations,* Rand Report RM- 5387-PR, The RAND Corporation, Santa Monica, California, 1967.
18. E.H. Forman and S.I. Gass. The Analytic Hierarchy Process: An Exposition. *Operations Research*, 49:469-486, 2001.
19. C.A. Fossett, D. Harrison, H. Weintrob, and S.I. Gass. An Assessment Procedure for Simulation Models: A Case Study. *Operations Research*, 39:710-723, 1991.
20. T. Gal. A "Historiogramme" of Parametric Programming. *Journal of the Operational Research Society,* 31:449-451, 1980.
21. T. Gal. A Note on the History of Parametric Programming. *Journal of the Operational Research Society,* 34:162-163, 1983.
22. S.G. Garille and S.I. Gass. Stigler's Diet Problem Revisited. *Operations Research*, 49:1-13, 2001.
23. S.I. Gass. The Application of Linear Programming Techniques to Air Force Problems. Paper AFAMA 3-4-52, Air Force, Directorate of Management Analysis Service, Headquarters USAF, December 1954.
24. S.I. Gass. A First Feasible Solution to the Linear Programming Problem. In *Proceedings of the Second Symposium in Linear Programming,* pages 495-508. H. Antosiewicz, ed., January 1955.

25. S.I. Gass. On the Distribution of Manhours to Meet Scheduled Requirements. *Naval Research Logistics Quarterly*, 4:17-25, 1957.
26. S.I. Gass. *Linear Programming: Methods and Applications*, McGraw-Hill Book Company, 1958. (Fourth edition, 1975.)
27. S.I. Gass. Recent Developments in Linear Programming. In *Advances in Computers*, Vol. 2, pages 295-377, (Chapter 4). Franz L. Alt, ed., Academic Press, New York, 1961.
28. S.I. Gass. The Role of Digital Computers in Project Mercury. In *Computers: Key to Total Systems Control, Proceedings of the Eastern Joint Computer Conference*, pages 33-46. Macmillan Company, New York, Vol. 20, 1961.
29. S.I. Gass. The Dualplex Method of Large-Scale Linear Programs. ORC Report 66-15, Operations Research Center, U. of Calif., Berkeley, 1966.
30. S.I. Gass. On the Division of Police Districts into Patrol Beats. *Proceedings 1968 ACM National Conference*, 459-473, 1968.
31. S.I. Gass. *An Illustrated Guide to Linear Programming*, McGraw-Hill Book Company, 1970. (Reprinted by Dover Press, 1990.)
32. S.I. Gass. The Dualplex Method Applied to Special Linear Programs. In *Proceedings of the International Federation of Information Processing Societies,* 1971, pp. 1317-1323, North-Holland, 1972.
33. S.I. Gass. Reactions to the Guidelines for the Practice of Operations Research, *Operations Research*, 20:224-225, 1972.
34. S.I. Gass. Models in Law Enforcement and Criminal Justice. In *A Guide to Models in Governmental Planning and Operations,* pages 231-275, (Chapter 8). S. I. Gass and R. L. Sisson, eds. Sauger Books, 1975.
35. S.I. Gass. Evaluation of Complex Models. *Computers and Operations Research*, 4:27-35, 1977.
36. S.I. Gass. A Procedure for the Evaluation of Complex Models. In *Proceedings of the First International Conference in Mathematical Modeling*, University of Missouri, 1977.
37. S.I. Gass. Computer Model Documentation. In *Proceedings 1978 Winter Simulation Conference (Miami Beach, Florida)*, pages 281-287. IEEE, Piscataway, New Jersey, December 1978.
38. S.I. Gass. *Computer Model Documentation: A Review and an Approach*, National Bureau of Standards Special Publication 500-39, U.S. GPO Stock No. 033-003-02020-6, Washington, D.C., February 1979.
39. S.I. Gass, editor. *Utility and Use of Large-Scale Mathematical Models*, Proceedings of a Workshop, National Bureau of Standards Special Publication 534, U.S. GPO Stock No. 003-003-02060-5, Washington, D.C., May 1979.
40. S.I. Gass. Comments on the Possibility of Cycling with the Simplex Algorithm. *Operations Research*, 27:848-852, 1979.
41. S.I. Gass. *Validation and Assessment Issues of Energy Models*, Proceedings of a Workshop, National Bureau of Standards Special Publication 569, U. S. GPO Stock No. 033-003-02155-5, Washington, D.C., 1980.

42. S.I. Gass. Assessing Ways to Improve the Utility of Large-Scale Models. In *Validation and Assessment Issues of Energy Models,* U.S. GPO, Washington, D.C., 1980.
43. S.I. Gass. Validation and Assessment Issues of Energy Models. In *Energy Policy Planning*, pages 421-441. B. A. Bayraktar, A. Laughton, E. Cherniaavsky, and L. E. Ruff, eds., Plenum Press, New York, 1981.
44. S.I. Gass, editor. *Validation and Assessment of Energy Models*, Proceedings of a Symposium, National Bureau of Standards Special Publication 616, 1981.
45. S.I. Gass, editor. *Operations Research: Mathematics and Models*, American Mathematical Society, 1981.
46. S.I. Gass. Decision-Aiding Models: Validation, Assessment, and Related Issues for Policy Analysis (feature article). *Operations Research*, 31:603-631, 1983.
47. S.I. Gass. What is a Computer-Based Mathematical Model? *Mathematical Modelling*, 4:467-471, 1983.
48. S.I. Gass. Documenting a Computer-Based Model. *Interfaces*, 14(3):84-93, 1984.
49. S.I. Gass. On the Development of Large-Scale Personnel Planning Models. In *Proceedings of the 11th IFIP Conference on System Modeling and Optimization*, pages 743-754. P. Throft-Christensen, ed., Springer-Verlag, New York, 1984.
50. S.I. Gass. *Decision Making, Models and Algorithms: A First Course*, Wiley-Interscience, New York, 1985.
51. S.I. Gass. *Linear Programming: Methods and Applications*, 5th edition, McGraw-Hill Book Company, 1985.
52. S.I. Gass. On the Solution of Linear-Programming Problems with Free (Unrestricted) Variables. *Computers and Operations Research,* 12:265-271, 1985.
53. S.I. Gass. A Process for Determining priorities and weights for Large-Scale Linear Goal Programs. *Journal of the Operational Research Society,* 37:779-785, 1986.
54. S.I. Gass. President's Symposium: A Perspective on the Future of Operations Research. *Operations Research*, 35:320-321, 1987.
55. S.I. Gass. On Artificial Variables and the Simplex Method. *Journal of Information and Optimization Sciences*, 8:77-79, 1987.
56. S.I. Gass. Managing the Modeling Process: A Personal Reflection. *European Journal of Operational Research,* 31:1-8, 1987.
57. S.I. Gass. The Setting of Weights in Linear Goal-Programming Problems. *Computers & Operations Research*, 14:227-229, 1987.
58. S.I. Gass. Operations Research - Supporting Decisions Around the World. In *Operational Research '87, Proceedings of the 11th International Conference on Operational Research (IFORS),* Plenary paper, pages 3-18. G. K. Rand, ed., North-Holland, New York, 1988.
59. S.I. Gass. A Model is a Model is a Model is a Model. *Interfaces,* 19(3):58-60, 1989.

60. S.I. Gass. The Current Status of Operations Research and a Way to the Future. *The Journal of the Washington Academy of Sciences*, 79:60-69, 1989.
61. S.I. Gass. Comments on the History of Linear Programming. *IEEE Annals of the History of Computing*, 11:147-151, 1989.
62. S.I. Gass. Model World: Have Model, Will Travel. *Interfaces*, 20(2):67-71, 1990.
63. S.I. Gass. On Solving the Transportation Problem, *Journal of the Operational Research Society*, 41:291-297, 1990.
64. S.I. Gass. Model World: Danger, Beware the User as Modeler. *Interfaces*, 20(3):60-64, 1990.
65. S.I. Gass. Model World: In the Beginning There was Linear Programming. *Interfaces*, 20(4):128-132, 1990.
66. S.I. Gass. The Many Faces of OR. *Journal of the Operational Research Society*, 42:3-15, 1991.
67. S.I. Gass. Military Manpower Planning Models. *Computers & Operations Research*, 18:65-73, 1991.
68. S.I. Gass. OR in the Real World: How things Go Wrong. *Computers & Operations Research*, 18:629-632, 1991.
69. S.I. Gass. Model World: Models at the OK Corral. *Interfaces*, 21(6):80-86, 1991.
70. S.I. Gass. Model Accreditation: A Rationale and Process for Determining a Numerical Rating. *European Journal of Operational Research*, 66:250-258, 1993.
71. S.I. Gass. Encounters with Degeneracy: A Personal View. *Annals of Operations Research*, 47:335-342, 1993.
72. S.I. Gass. Ethical Concerns and Ethical Answers. In *Ethics in Modeling*, pages 207-225. William A. Wallace, editor, Elsevier, Tarrytown, New York, 1994.
73. S.I. Gass. Public Sector Analysis and Operations Research/Management Science. In the *Handbook of OR: Operations Research and the Public Sector*, pages 23-46 (Ch. 2). A. Barnett, S. M. Pollock, and M. Rothkopf, eds., North-Holland, New York, 1994.
74. S.I. Gass. Not This Merger Proposal. *OR/MS Today*, 44-46, Feb. 1994.
75. S.I. Gass. Don't Merge, Restructure. *OR/MS Today*, 69-70, April 1994.
76. S.I. Gass. On Ethics in Operational Research. *Journal of the Operational Research Society*, 45:965-966, 1994.
77. S.I. Gass. Model World: Ethics in the Not So Real MS/OR World. *Interfaces*, 24(6):74-78, 1994.
78. S.I. Gass. Model World : On Academics, Applications, and Publications. *Interfaces*, 26(6):105-111, 1996.
79. S.I. Gass. Goal Programming in Networks. In *Proceedings of First Conference in Multi-objective Programming and Goal Programming*, pages 212-234. M. Tamiz, editor, Springer-Verlag, Berlin, 1996.
80. S.I. Gass. On Copying A Compact Disk to Cassette Tape: An Integer-Programming Approach. *Mathematics Magazine*, 69:57-61, 1996.

81. S.I. Gass. The Washington Operations Research Connection: The Rest of the Story. *Socio-Economic Planning Sciences,* 31:245-255, 1997.
82. S.I. Gass. Model World: OR is the Bridge to the 21st Century. *Interfaces,* 27(6):65-68, 1997.
83. S.I. Gass. On the Education of a Multi-Criteria Researcher: A Personal View. Abraham Charnes Distinguished Lecture, pp. 5-26 In *Lecture Notes in Economics and Mathematical Systems, #332: Advances in Multiple Objective and Goal Programming,* Proceedings of the Second International Conference on Multi-Objective Programming and Goal Programming, Torremolinos, Spain, May 16-18, 1996, *Rafael Caballero, Francisco Ruiz, and Ralph E. Steuer, eds.* Springer-Verlag, Berlin, 1997.
84. S.I. Gass. Tournaments, Transitivity and Pairwise Comparison Matrices. *Journal of the Operational Research Society,* 49:616-624, 1998.
85. S.I. Gass. Project Mercury Man-in-Space Real-Time Computer System: "You Have a Go, at Least Seven Orbits!" *IEEE Annals of the History of Computing,* 21:37-48, 1999.
86. S.I. Gass. Ethics in Journal Refereeing. Letter to the Editor, *Communications of the ACM,* 42:11, 1999.
87. S.I. Gass. Model World: When is a Number a Number? *Interfaces,* 31(5):93-103, 2001.
88. S.I. Gass. The First Linear Programming Shoppe. *Operations Research,* 50:61-68, 2002.
89. S.I. Gass. Model World: Reflections on the 50th Anniversary of the Founding of ORSA. *Interfaces,* 32(4):46-51, 2002.
90. S.I. Gass. OR's Top 25. *ORMS TODAY,* 29:42-45, 2002.
91. S.I. Gass. Great Moments in HistORy. *ORMS TODAY,* 29:31-37, 2002.
92. S.I. Gass. Not a Trivial Matter. *ORMS TODAY,* 29, 5, 46-48, October 2002.
93. S.I. Gass. IFORS' Operational Research Hall of Fame: George B. Dantzig. *International Transactions in Operational Research,* 10:191-193, 2003.
94. S.I. Gass. IFORS' Operational Research Hall of Fame: Albert William Tucker. *International Transactions in Operational Research,* 11:239-242, 2004.
95. S.I. Gass. Model World: The Great Debate: MAUT versus AHP. *Interfaces,* 35:308-312, 2005.
96. S.I. Gass. In Memoriam: The Life and Times of the Father of Linear Programming. *ORMS Today,* 32:40-48, 2005
97. S.I. Gass. Military Operations Research Society (MORS) Oral History Project Interview of Saul I. Gass. Interviewed by G. Visco and B. Sheldon, *Military Operations Research,* 10:39-62, 2005.
98. S.I. Gass. IFORS' Operational Research Hall of Fame: John von Neumann, *International Transactions in Operational Research,* 13:85-90, 2006.
99. S.I. Gass. Private communication, April 2006.
100. S.I. Gass and A.A. Assad. *An Annotated Timeline of Operations Research: An Informal History,* Springer + Business Media, New York, 2005.

101. S.I. Gass and A.A. Assad. Model World: Tales from the Timeline: The definition of Operations Research and the Origins of Monte Carlo Simulation. *Interfaces*, 35(5):429-435, 2005.
102. S.I. Gass, S. Bhasker, and R. E. Chapman. *Expert Systems and Emergency Management: An Annotated Bibliography*, NBS Special Publication 728, National Bureau of Standards, Washington, DC, November 1986.
103. S.I. Gass and R.E. Chapman. *Theory and Application of Expert Systems in Emergency Management Operations,* NBS Special Publication 717, National Bureau of Standards, Washington, DC, November 1986.
104. S.I. Gass, R.W. Collins, C.W. Meinhardt, D.M. Lemon, and M.D. Gillette. Army Manpower Long Range Planning System. *Operations Research*, 36:5-17, 1988.
105. S.I. Gass, N.A. David, and R.H. Cronin. To CB or Not to CB: A Simulation Model for Analyzing the Role of Citizens Band Radio in Traffic Safety. *Computers and Operations Research*, 6:99-111, 1979.
106. S.I. Gass, N.A. David, and P. Levy. An Application of System Flow Models to the Analysis of Highway Safety Demonstration Projects. *Accident Analysis and Prevention*, 9:285-302, 1977.
107. S.I. Gass, and M. Dror. An Interactive Approach to Multiple-Objective Linear Programming Involving Key Decision Variables. *Large Scale Systems*, 5:95-103, 1983.
108. S.I. Gass, H.J. Greenberg, K.L. Hoffman , and R.W. Langley; eds.; *Impacts of Microcomputers on Operations Research*, North-Holland, 1986.
109. S.I. Gass, and D. Gross. In Memoriam: Carl M. Harris, 1940-2000. *INFORMS Journal on Computing*, 12:257-260, 2000.
110. S.I. Gass and C.M. Harris, eds., *Encyclopedia of Operations Research and Management Science*, Kluwer Academic Publishers, 1996.
111. S.I. Gass and C.M. Harris, eds., *Encyclopedia of Operations Research and Management Science*, 2nd edition, Kluwer Academic Publishers, 2001.
112. S.I. Gass, D.S. Hirshfeld and E.A. Wasil. Model World: The Spreadsheeting of OR/MS. *Interfaces*, 30(5):72-81, 2000.
113. S.I. Gass, K.L. Hoffman, R.H. F. Jackson, L.S. Joel, and P.B. Saunders. Documentation for a Model: A Hierarchical Approach. *Communications of the ACM*, 24:728-733, 1981.
114. S.I. Gass and L.S. Joel. Concepts of Model Confidence. *Computers and Operations Research*, 8:341-346, 1981.
115. S.I. Gass, S.Nahmias, and C.M. Harris. Model World: The Academic Midlife Crisis. *Interfaces*, 27(5):54-57, 1997.
116. S.I. Gass and S.C. Parikh. Credible Baseline Analysis for Multi-Model Public Policy Studies. In *Energy Models and Studies,* pages 85-97, (Chapter 4). B. Lev, ed., North-Holland, New York, 1983.
117. S.I. Gass and S.C. Parikh. On the Determinants of Credible Analysis. In *Mathematical Modeling in Science and Technology,* Proceedings of the Fourth International Conference, Zurich, Switzerland, pages 975-980. X. J. R. Avula, ed., Pergamon Press, 1983.

118. S.I. Gass and T. Rapcsák. A Note on Synthesizing Group Decisions. *Decisions Support Systems*, 22:59-63, 1998.
119. S.I. Gass and T. Rapcsák. Singular Value Decomposition in AHP. *European Journal of Operational Research*, 154:573-584, 2004.
120. S.I. Gass and P.G. Roy. The Compromise Hypersphere for Multiobjective Linear Programming. *European Journal of Operational Research*, 144:459-479, 2003.
121. S.I. Gass and T.L. Saaty. The Computational Algorithm for the Parametric Objective Function. *Naval Research Logistics Quarterly*, 2:39-45, 1955.
122. S.I. Gass and T.L. Saaty. Parametric Objective Function (Part 2)—Generalization. *Journal of the Operations Research Society of America*, 3:395-401, 1955.
123. S.I. Gass and S.P. Shao, Jr. On the Solution of Special Generalized Upper-Bounded Problems: The LP/GUB Knapsack Problem and the λ-Form Separable Convex Objective Function Problem. *Mathematical Programming Study*, 24:104-115, 1985.
124. S.I. Gass and R.L. Sisson, eds., *A Guide to Models in Governmental Planning and Operations*. Sauger Books, Potomac, Maryland, 1975.
125. S.I. Gass and S.M. Standard. Characteristics of Positive Reciprocal Matrices in the Analytic Hierarchy Process. *Journal of Operational Research Society*, 53:1385-1389, 2002.
126. S.I. Gass and B.W. Thompson. Guidelines for Model Evaluation: An Abridged Version of the U.S. General Accounting Office Exposure Draft. *Operations Research*, 28:431-439, 1980.
127. S.I. Gass and S.R. Torrence. On the Development and Validation of Multicriteria Ratings: A Case Study. *Socio-Economic Planning Sciences*, 25:133-142, 1991.
128. S.I. Gass and S. Vinjamuri. Cycling in Linear Programming Problems. *Computers & Operations Research*, 31:303-311, 2004.
129. S.I. Gass and C. Witzgall. On an Approximate Minimax Circle Closest to a Set of Points. *Computers & Operations Research*, 31:637-643, 2004.
130. S.I. Gass, C. Witzgall, and H.H. Harary. Fitting Circles and Spheres to Coordinate Measuring Machine Data. *International Journal of Flexible Manufacturing*, 10:5-25, 1998.
131. S.I. Gass and P.M.R. Zafra. Modified Fictitious Play for Solving Matrix games and Linear-Programming Problems. *Computers & Operations Research*, 22:893-903, 1995.
132. S.I. Gass, P.M.R. Zafra and Z. Qiu. Modified Fictitious Play. *Naval Research Logistics*, 43:955-970, 1996.
133. S.E. Hauser, G.R. Thoma, and S.I. Gass. Factors Affecting the Conversion Rate of Bound Volumes to Electronic Form. In *Proceedings of the Electronic Imaging East '89 Conference*, pages 1041-1046. Boston, Mass., October 2-5, 1989.
134. P.R. Horner. High Octane Gass. *ORMS TODAY*, 28:6, 2001.
135. A. Jones and S.I. Gass, eds. Special Issue : Supply Chain Management. *Information Systems Frontiers*, 3:397-492, 2001.

136. A. Joseph and S.I. Gass. A Framework for Constructing General Integer Problems with Well-Determined Duality Gaps. *European Journal of Operational Research*, 136:81-94, 2002.
137. A. Joseph, S.I. Gass, and N. Bryson. A Computational Study of an Objective Hyperplane Search Heuristic for the General Integer Linear Programming Problem. *Mathematical and Computer Modelling*, 25:63-76, 1997.
138. A. Joseph, S.I. Gass, and N. Bryson. An Objective Hyperplane Search Procedure for Solving the General All-Integer Linear Programming (ILP) Problem. *European Journal of Operational Research,* 104:601-614, 1998.
139. A. Joseph, S.I. Gass, and N.A. Bryson. Nearness and Bound Relationships Between an Integer-Programming Problem and its Relaxed Linear Programming Problem. *Journal of Optimization Theory and Applications (JOTA)*, 98:55-63, 1998.
140. T.L. Magnanti and J.B. Orlin. Parametric Linear Programming and Anti-Cycling Pivoting Rules. *Mathematical Programming*, 41:317-325, 1988.
141. V. Riley and S.I. Gass. *Linear Programming and Related Techniques; A Comprehensive Bibliography*, Operations Research Office, Johns Hopkins Press, Chevy Chase, Maryland, 1958.
142. P.F. Roth, S.I. Gass, and A.J. Lemoine, Some Considerations for Improving Federal Modeling. *Proceedings of the 10th Winter Simulation Conference*, 1:211-218, 1978.
143. T. Saaty. *The Analytic Hierarchy Process*, McGraw-Hill Book Company, New York, 1980.
144. T. Saaty and S.I. Gass. The Parametric Objective Function - Part I. *Journal of the Operations Research Society of America*, 2:316-319, 1954.
145. U.S. GAO, *Review of the 1974 Project Independence Evaluation System.* OPA-76-20, Washington, D.C., April 21, 1976.
146. U.S. GAO, *Ways to Improve Management of Federally Funded Computerized Models.* LCD-75-111, Washington, D.C., August 13, 1976.
147. U.S. GAO, *Guidelines for Model Evaluation.* PAD-79-17, Washington, D.C., January 1979.
148. S. Vajda. *The Theory of Games and Linear Programming*, Methuen & Co., London, 1956.
149. J.D. Williams. *The Compleat Strategist*, McGraw-Hill, New York, 1954.
150. M.K. Wood and G.B. Dantzig. Programming of Interdependent Activities, I, general discussion. *Econometrica*, 17:193-199, 1949. Also in *Activity Analysis of Production and Allocation,* pages 15-18. Tjalling C. Koopmans, ed., John Wiley & Sons, New *Activity Analysis of Production and Allocation*, John Wiley & Sons, New York, 1951.

Appendix: Reminiscences of Saul Gass

On the occasion of the 80th birthday of Saul Gass, I invited a group of his friends and colleagues to contribute brief reminiscences of their interactions with him. My objective was to add first-hand and personal observations of Saul to my overview article in this volume. I was very grateful to receive the contributions that follow. For each piece, I have indicated the author and the date on which I received the contribution. Except for minor edits, I have reproduced the original writings. I extend my sincere thanks to all contributors who responded to my request.

Arjang Assad

Al Blumstein

Alfred Blumstein is University Professor and the J. Erik Jonsson Professor of Urban Systems and Operations Research at the Heinz School, Carnegie Mellon University
2/19/06

I got to know Saul when I arrived in Washington in 1961. Saul was then already a key player in the OR community there. As I recall, Saul was working for a very strong consulting firm named CEIR. We must have met at a meeting of the local Washington Operations Research Council, or WORC. There were lots of members and lots of good folks who were serious about OR, especially as it applied to government problems – of course, this was Washington. Most of us were working with the military, some in DoD, some in the think tanks like IDA, RAND, ORO, CNA, and many in the various consulting firms around before the term "Beltway Bandits" became known – mostly because the Beltway wasn't built until the late 1960s. In 1965, Lyndon Johnson's President's Crime Commission decided to establish a Task Force on Science and Technology – presumably because science and technology was sending a man to the moon, so it certainly must be able to solve the urban crime problem. Somehow, I got recruited to lead that effort, and Saul was a natural candidate to join, and so I called him to see if he could become available. He was working for IBM at the time, and he was able to arrange for part-time participation without leaving IBM.

The project he took on was finding out how to stem the then-growing rate of auto theft, and he discovered that a very large fraction of stolen cars had their keys left in them. That led to a variety of proposals to get people to remove their keys from the cars. Our favorite was to have the key spring-loaded so that it would pop out into the driver's hand as soon as he turned the car off. The auto manufacturers claimed that the keys would too often land on the floor, so they settled for just making a lot of noise whenever the car was turned off and the keys were left in the ignition. And that led to a whole variety of noisy sounds in cars whenever the driver and other occupants weren't doing what they were supposed to – keys, doors, seatbelts. So you can blame Saul Gass for all that noise.

Don Gross
Donald Gross is Research Professor at the Department of Systems Engineering & Operations Research, George Mason University
2/12/06

I have known Saul since the 1960s when I first came to the DC area. While most of our interaction has been professionally through ORSA and now INFORMS, my wife and I have been at several social events with Saul and Trudy. About two years ago, we were at a wedding where a lot of dancing took place. It was a pleasure to watch Saul and Trudy on the dance floor. They outperformed even the youngsters (20 and 30 somethings). My advice to Saul is to compete for the "Dancing with the Stars" TV program. If he *optimizes* the choreography, I think he could win hands down!

Tom Magnanti
Thomas L. Magnanti is the Dean of the School of Engineering and Institute Professor at MIT
2/20/06

I first encountered Saul Gass, like I imagine many others, through his pioneering textbook, *Linear Programming: Methods and Applications.* At the time I was an undergraduate student studying chemical engineering and his book opened my eyes to an entirely new world – optimization and operations research – that has consumed my professional life ever since. It is fair to say that I am an operations researcher because of Saul Gass.

Later, my first professional research grant was for a DOT program called TARP, the Transportation Advanced Research Program. Other grantees included such luminaries as Saul and George Dantzig. At the end of a day of a grant review meeting held in Boston, my wife and I drove Saul, George and Robert Dorfman to dinner in our aging Oldsmobile, equipped with very poor shock absorbers. Upon hitting one of Boston's finest potholes, Saul, George, and Bob all levitated, seemingly hitting their heads on the car roof. Later that night my wife said she feared that we might have put an end to the field of linear programming.

Saul and my professional lives have been interwoven ever since. As but a few examples: (i) we both served as President of the Operations Research Society of America (ORSA), (ii) several years ago, I contributed an article to the *Encyclopedia of Operations Research and Management Science* that he so expertly edited (including careful and thoughtful editing of my modest contribution), (iii) he invited me to participate in one of the famous Knowledge Bowls that he has organized for our professional meetings (my team was pretty dismal!), (iv) on a couple of occasions, we participated in celebratory sessions that Saul organized in honor of our common thesis advisor, the great George Dantzig (and so, by virtue of our common thesis advisor, Saul and I are academic brothers), and (v) as the current President of the International Federation of Operational Research Societies (IFORS), I had the privilege recently of inviting him to be the IFORS

Distinguished Lecturer at the 2006 EURO meeting. The later of these examples provides just one indication of Saul's professional preeminence. Items (i)-(iv) provide a glimpse into the enormous impact that he has had in serving the profession with singular creativity, flair, and distinction. Indeed, he has always seemed to be everywhere; I marvel at his endless energy and at his steadfast commitment to the profession — unlike any other.

Heiner Müller-Merbach
Heiner Müller-Merbach is Professor Emeritus at the Technische Universität Kaiserslautern in Germany
2/20/06

When I was a student in the late 1950s, we used Saul's textbook, which was published prior to George Dantzig's textbook. During 1963-64, I spent a post-doctoral year at the Operations Research Center of the University of California-Berkeley. The center was at the "Richmond Field Station," far away from the campus. George Dantzig was the director. He told me to share the office with Saul. I was quite surprised to meet him working for his Ph.D. - since he was so far ahead of me in linear programming. It was a great year.

After my return to Germany Saul and I met on many occasions. Saul and Trudy visited us once or twice. It may have been sometime around 1980 that they came from Russia, almost starving. We organized some good stuff to feed them to gain some energy again, which he depends on for his 10k races. I also remember that Saul attended the 1981 IFORS Conference in Hamburg. At the general assembly, Saul - on behalf of the US delegation- nominated me for the next IFORS President. The other candidate was my friend Hans-Jurgen Zimmermann, also from Germany; he was supported by several European OR societies. The voting process in the following year was in my favor, and I became President of IFORS for the three year period 1983 through 1985. In 1984, we had the triennial IFORS Conference in Washington DC, where we celebrated the 25th anniversary of IFORS. [*Dr. Muller-Merbach has also written an appreciation entitled "Saul Gass: 80 Years" for the German OR newsletter.*]

Graham Rand
Graham Rand is Senior Lecturer in Operational Research at Lancaster University and Managing Editor, International Transactions on Operational Research
2/17/06

It is a great pleasure and an honour to be asked to contribute some reflections on my association with Saul on the occasion of his 80th birthday. It is hard to believe that he has reached that milestone. I remember the years when he was competing in the Saul Gass 10k runs at conferences: even then I was never fit enough to

follow in his footsteps. Nevertheless, I have been greatly privileged to count Saul as a friend, whom it has been a delight to meet at INFORMS and IFORS meetings over the years. The pleasure has frequently been enhanced by Trudy's company. I recall a lovely dinner in San Antonio on the evening of Bush Junior's first election "victory". The British contingent were joined by John Little and Saul. Saul kept disappearing to catch up on the latest news: well at least we knew the result was going to be close!

Saul is, of course, a very welcome visitor to the UK. I was involved a little in the organisation of the Operational Research Society's annual conference held in Bangor, Wales in 1990, when Saul gave the opening address on "The many faces of OR". The Welsh members present were probably not too impressed when the Chairman welcomed him to England. Most of our professional collaboration has concerned the history of OR. In my role as the organiser of the IFORS' Operational Research Hall of Fame, Saul has given invaluable help, suggesting possible citation authors, and contributing citations himself for George Dantzig, Albert Tucker, and John von Neumann. His contributions, as would be expected, were interesting and delivered on time. Although I made a small contribution to his encyclopedia, and have fielded several questions about British OR history, the balance of debt is greatly in my favour. From my perspective, this particular Anglo-American relationship is very special.

Rick Vohra
Rakesh V. Vohra is John L. and Helen Kellogg Professor of Managerial Economics and Decision Science at Northwestern University
2/11/06

I have Margaret Thatcher to thank for becoming a Ph.D. student of Saul's. In the spring of 1980 while she ran for government, I campaigned for Labor and pondered what to do upon graduation. The fall of that year found me, a graduate student at the LSE, learning economics in the pubs and cafes. By December, Thatcher's program to shrink Government was in full swing. One consequence was deep cuts in University spending. My tutor, Ailsa Land, called me into her office and told me, because of the cuts, if I wanted an academic career, I should "go west." She pointed me to College Park and urged me to take up with Saul Gass.

In the fall of 1981 I joined the Mathematics department at the University. It was, I think, in the spring of 1982 that I made my first expedition (it took 15 minutes) from the Glenn Martin building (where Mathematics was housed) to Tydings Hall (where the Business school was housed) to meet Saul. To get there, one had to traverse the mall, which was criss-crossed by a series of concrete paths placed with no particular concern for symmetry or order. Legend has it that the mall was designed with no walkways but a great expanse of green. The planners would wait to see what paths were beaten into the grass to decide where to locate concrete paths.

Saul's office was dark, at least that is what I recall. He was surrounded by piles of books and papers. From the perspective of a regular patron at the local pizza shop, he seemed awfully emaciated to me; I did not know then that he ran a lot. I told him of my interest in OR and something of my interests in using it to house the homeless, clothe the unclothed, and shod the shoeless. I'm pretty sure it was drivel. He was kind enough not to throw me out of the office and asked instead about courses I had taken and told me what courses I should take.

My next meeting with Saul was in his LP class. Saul taught three things. Some I absorbed right away, others have taken time to sink in. First, formulate everything as an LP. Second, pass to the dual. Now this is not always helpful, but when it is, it is. I tell my graduate students the same thing now. The third thing Saul taught, was the notion of the right level of generality. LP is useful not because it is the most general of all classes of optimization problems. It is useful because it is just general enough to model a rich enough variety of situations and structured enough that one can say useful things about the phenomena being modeled. Saul was by no means an obvious match for an advisor. He knew and took an interest in the activities of the basketball coach Lefty Driesell; I did not. When the urge to exercise would come upon me, I would lie down, Saul would not. Saul was interested in real models that would support real decisions. I wanted to prove theorems. Nevertheless, it worked out!

Chris Witzgall
Christoph Witzgall has retired as Research Mathematician from the National Institute of Standards and Technology (NIST)
3/15/06

Traveling with the Gass family in Russia 1990, we visited an art museum. Afterwards, gleaming with a mathematician-father's joy, Saul recounted to whoever would listen how his son, the lawyer, had recognized a Möbius band among the decorations and pointed it out to him.

Larry Bodin
***Laurence D. Bodin is Professor Emeritus of Management Science at the Robert H. Smith of Business of the University of Maryland professor* 2/9/06**

Saul Gass introduced me to Operations Research in 1962 through his book, *Linear Programming*, the first edition. I saw Saul's book in the bookstore and convinced the chair of the Mathematics Department at Northeastern University to offer me a reading class. Nobody in the Mathematics Department at Northeastern University knew anything about Operations Research in 1962.

I met Saul for the first time at IBM in Bethesda in 1964. We were introduced by Ed Brown, a mutual colleague. Saul started his Ph.D. program in Operations Research at the University of California, Berkeley in 1963, and I became a student

in the same program in 1964. Our paths intertwined for the next 10 years – graduate school, IBM, consulting companies and a faculty position at a university. I joined Saul Gass in the Business School at the University of Maryland in 1976. We had many rewarding experiences together. We introduced AHP into the graduate and undergraduate programs in the mid-90s and co-authored two papers on the teaching of the AHP, began writing the "definitive OR text for MBA students" (parts were written and used in the classroom, but the book was never completed). I fondly remember looking forward to Saul's presentations at the College's faculty assembly. His use of Dilbert to emphasize a point was terrific. Saul and I retired on July 1, 2001 and are now roommates in Van Munching Hall. Carol and I look forward to our continued association with Saul and Trudy Gass.

Ed Baker
Ed Baker is Professor of Management Science at the School of Business, University of Miami
2/10/06

One of things that has always impressed me about Saul Gass has been his openness and willingness to talk with and mentor students. In November of 1976, ORSA-TIMS held its fall meeting in Miami Beach. I was a second year doctoral student at the University of Maryland, and Saul Gass was the department chair. Saul was also the President of ORSA that year and, as a result, he had a large suite overlooking the ocean. He also, as I came to learn, would be hosting a cocktail party of ORSA-TIMS officials in his suite.

I was thrilled to be attending my first ORSA/TIMS conference and looked forward to hearing a number of papers. What a wonderful surprise it was for me to also receive an invitation to come to a reception in Saul's suite. I went to the event not knowing what to expect, but found Saul and Trudy, as gracious as ever, talking to a small group of people. Saul welcomed me and proceeded to introduce me to the group that included David Hertz, Art Geoffrion, and Jack Borsting, all ORSA luminaries. I was duly impressed with the company, and particularly with Saul's kindness and gracious nature in including me in the group.

I am sure that no one in that group remembers my introduction, but the meeting proved fortuitous for me. In 1984, I would become a colleague of David Hertz and Jack Borsting at the Business School at the University of Miami. Jack was serving as dean of the School of Business, and David was a University Professor. As it happened, I came up for promotion under Jack's deanship, and I know that he called upon Art Geoffrion, David Hertz, and Saul Gass to write letters reviewing my work. I don't know what they said, but promotion came through. I am certainly indebted to each of them, but I am especially grateful to Saul Gass for giving me that introduction.

In the Beginning: Saul Gass and Other Pioneers

Alfred Blumstein

H. John Heinz III School of Public Policy & Management
Carnegie Mellon University
Pittsburgh, PA 15213
ab0q@andrew.cmu.edu

Summary. This paper is an abridged and edited transcription of the author's lecture given on February 25th, 2006, in honor of Saul Gass.

Key words: Missionary; President's Crime Commission; system; flow chart.

What we've heard so far today is recent history about Saul. I would like to go back a few years to a period when I got to know Saul as essentially one of the missionaries bringing OR to a new field.

I made a dichotomy in an article in *OR/MS Today,* back in the old days when Armand Weiss was editor of that publication. The argument was that OR consists both of methodologists, who are obviously the more numerous, and also a population of missionaries, whose basic mission is bringing OR to the not-yet-quantified fields. Our mission is to transform those heathens or the unenlightened into people who can deal with, and operate quantitatively in, the field in which we believe. That concept seems to make a lot of sense in terms of getting more rigorous analysis associated with their work.

In most cases, these heathens get transformed into believers and into users. Many fields of engineering have adopted and incorporated as part of their field methods like linear programming and optimization, which hadn't been there before. So, we bring the OR techniques, and they become appropriated and a part of that field.

Some of the early missionaries were the ones that began Operations Research with that bizarre name, because it was research after all. They had to do research, because they were academics, and it was on operations, including military operations. Therefore, we got the name Operations Research that nobody in the world understands until they have been through OR. It's been one of the burdens we carry in trying to communicate with a broader world.

I have had many interactions with the press, and they frequently ask, "What are you a professor of?" I say, "Operations Research." They say, "What is that?" I try to give them an explanation, and they ask, "Can I call you a criminologist?" I respond by saying, "Ok, but,..." Anyway, that was certainly the start.

Some of the very early work that was important was in the field of transportation. Bob Herman and Les Edie made some major contributions in that area. More recently, missionary activities certainly include the continuing efforts in military OR, as well as the ongoing research in transportation, pertaining not

only to highways, but also to air, rail, auto, and so on. We have also seen a lot of effort in the area of health care.

I want to talk about the missionaries that started bringing OR into the criminal justice system, which is probably the most heathen of all. Perhaps it's also the least able and willing to pick up on this kind of quantification, in large part because it's dominated by lawyers, and lawyers are the antithesis, typically, of systems people. Lawyers deal with an event at a time, a case at a time, and that's the way they learn; that's the way they operate. So it's a very tough community to totally transform, but some of us have tried for a long time.

This work started with the creation of the President's Crime Commission in the mid 1960s. For some reason, they established a task force on science and technology. I presume they did so, because if science and technology could get a man to the moon, it ought to be able to solve the urban crime problem. Therefore, we needed a task force to accomplish this.

Somehow I got recruited to lead that task force and brought together what was then, and especially in retrospect, a rather exciting team. Dick Larson had just graduated with a brand-new bachelor's degree in electrical engineering from MIT. Dick was clearly an important part of the group. Ron Christensen brought an intriguing collection of skills. He was then a Ph.D. student in nuclear physics at Berkeley, but also had a law degree from Harvard. He is now applying a wide variety of statistical methodology to a variety of areas. Sue Johnson was a private consultant. Of course, there was Saul Gass, who was then at IBM. He had previously been with a very well regarded consulting firm in Washington called CEIR. At that point (this was about 1965), Saul was already a very significant factor in OR, including the Washington Operations Research Council that went through a whole variety of name changes as the field changed. Saul was a very important participant there, as well as on the task force.

Saul took as his theme the issue of auto theft. Cars were being stolen at an increasing rate. Saul looked into the problem and found that a major factor was that the cars being stolen had the keys left in them. This gave rise to a whole variety of intriguing proposals that we explored to encourage people to remove their keys from their cars. For example, one proposal would have the keys spring loaded, so that as soon as the ignition switch went into the off position, the key would come out and pop into the driver's hands. The auto companies didn't like that, because people were going to complain that the keys fell on the floor, and the driver would have to search for them. The resolution was a buzzer, so all the strange sounds you hear in a modern car regarding unbuckled seatbelts, open doors, and keys left in the car, you can blame on Saul. That was a major transformation of the relationship between the automobile industry and their drivers. There were a number of other proposals, mostly in information systems, technology, and others.

What did we bring to bear to the criminal justice system? Although we knew nothing about crime or the criminal justice system, we brought a "systems perspective" that really was a very different perspective from that characterizing the lawyers who were the major part of the group. We had a window into the world of science and technology, but it was then very remote from law and especially the Department of Justice. At that time, it was the only Cabinet

department that didn't have an Assistant Secretary for Science and Technology, and it barely does today.

Although the systems perspective was very different, it provided the kind of opportunity that really doesn't occur very often. The opportunity was very exciting, because it gave us a chance to play an "emperor's clothes" role. Since we knew nothing about the system, we were in a great position to ask the people who said they knew what the evidence was, and the evidence was astonishingly thin on much of what was believed at the time. The criminal justice system had never been treated as a system in large part, because there was no system manager. Virtually every other system we deal with has a CEO or chief of staff, or someone of that sort, but the criminal justice system has, right in the middle, a judiciary which doesn't have control over the other parts. So there is no system manager, and so each part is seen as independent and autonomous. It's a very complicated, strange system. One of the lasting contributions we made was an embryonic version of what is now very common and very widely used as a system, an embryonic version of the now classic criminal justice system (CJS) flow chart.

Over the years that initial flow diagram has been extensively modified to incorporate feedback and other factors. Much more has happened in the last 40 years since Saul became a crime stopper. We have richer models, and we have fuller data. It's still a very political environment, as indicated at the outset. Dick Larson has made some major contributions to the field of Operations Research in general, and with regard to the criminal justice system in particular, in modeling police patrol. Ron Christensen developed a new statistical modeling methodology. Sue Johnson became a New York court administrator.

What about Saul? Saul was and is a missionary! Saul Gass has made an impressive career for his work in linear programming as a teacher and an educator, and in general as an explainer of OR methodology. We have all learned that being a missionary can be very rewarding. I suggest that all of you try it. Hopefully, you will like it. Thank you!

Learning from the Master: Saul Gass, Linear Programming, and the OR Profession

Thomas L. Magnanti

Dean, School of Engineering
Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139-4307
magnanti@mit.edu

Summary. This paper is an edited transcription of the author's lecture given on February 25th, 2006, in honor of Saul Gass.

Key words: Scholar, professional activist, community builder, author, house of optimization.

1 Introduction

As I look at today's program, I see names of the speakers and their organizations. But for two of us, the program says something different. For Chris Witzgall, it says "NIST, retired," and for Tom Magnanti it says "MIT, Dean of Engineering." I wonder if the organizers are trying to tell me something. This reminds me of a story.

A fellow was traveling through the Far East Islands and comes across an island inhabited by cannibals. Climbing onto the shore, he sees and then enters a butcher shop. The butcher sells brains identified by their source. The traveler looks on top of a case with brains and sees three signs that read, "Student's brains $2.00/lb," "faculty brains $10.00/lb," and then "dean's brains $75.00/lb." The guy exclaims, "My gosh, those deans must be smart!" to which the butcher quickly responds, "Do you know how many deans it takes to get a pound of brains?" So put the remainder of this talk in context. Unlike our previous speakers, I will not speak about specific content related to Saul's research. Rather, I will talk about Saul's professional accomplishments generally, touching a bit about what he has been doing all these years.

I start with the following question, "Why am I here?" This is a darn good question, by the way. Everyone else that is speaking today has conducted research with Saul and published with him. Not me. We did publish something together (but not on research), and I'll come back to that a little bit later. I doubt Saul recalls what that was, but I'll tell you all soon.

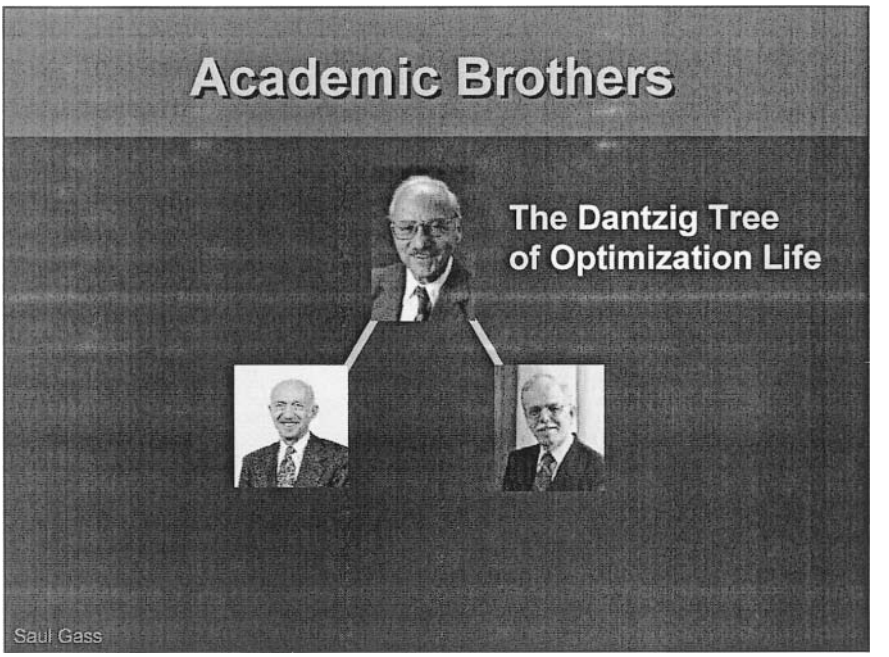


Fig. 1.

Saul and I are academic brothers. Figure 1 shows our father and the tree of life. It's the Dantzig tree of (optimization) life, and that's two of his academic children, Saul and me. Given that we are brothers, you could say that perhaps I should tell stories about dad. After all, that's what brothers do: tell stories about dad. Or, maybe I should tell you about how my older brother taunted me and tormented me in my youth. I could do that as well. Let me tell one family story. Although I've told this before, and Saul's heard it before, I'll tell it to the rest of you.

My first research contract ever (it was when I was a junior faculty member at MIT) was for an agency called the Transportation Advance Research Program, which was part of DOT. DOT decided at one point that it was going to fund a small number (I think there were only about 4 or 5) of projects on advanced research in transportation. Among those involved were George Dantzig and Saul. We held a project review meeting in Cambridge, MA. After the meeting, we decided to go to dinner and George said, "Why don't we invite my dear old friend Bob Dorfman to join us?" Off we went in a car that I'd driven across the country from Stanford, where I did my Ph.D. It was a really old car with terrible shock absorbers. As we were riding to dinner with my wife sitting beside me and Bob Dorfman, George Dantzig, and Saul Gass in the back seat, we hit a pothole, one of Boston's finest. I swear our backseat occupants levitated and hit their heads on the roof of the car. Later that night, my wife and I were discussing the incident, and here is where our two accounts vary somewhat. I claim she said, "I fear that today

you've just put an end to the world of optimization." Her version is, "I fear you've just destroyed the brain trust of America." Either version is pretty good. That's my family story.

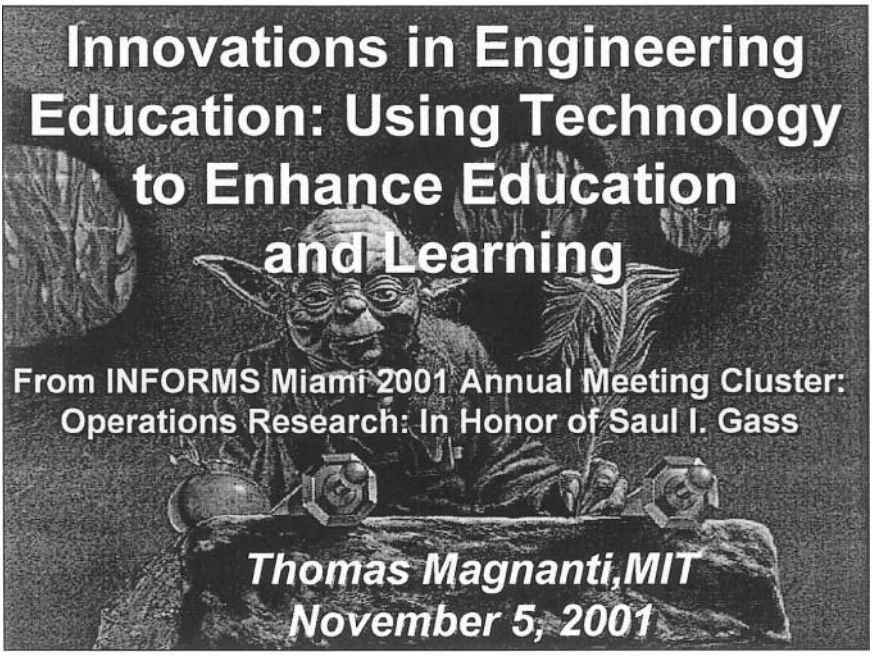


Fig. 2.

Figure 2 shows the theme of a talk that I presented in 2001. That was five years ago on the occasion of Saul's 75th birthday. One thing that I do now is give talks like this. I seem to go from one event to another, giving talks about friends and colleagues; it's part of being a dean. Just shortly before the talk represented in Figure 2, Saul won a distinguished award at the University of Maryland for contributions to teaching and education.

Let me share with you a bit of the citation. It stated that the University of Maryland named him one of five Distinguished Scholars/Teachers. The announcement for the award cited his constant interest in learning new ideas and approaches; he uses a PC software package as introductory state-of-the-art software in keeping with his commitment to enriching education. At that occasion, I thought it might be best to talk about education, and so I spoke in Saul's honor about education broadly. I also said Happy Birthday to you on that occasion, Saul, and I'll say Happy Birthday to you again today.

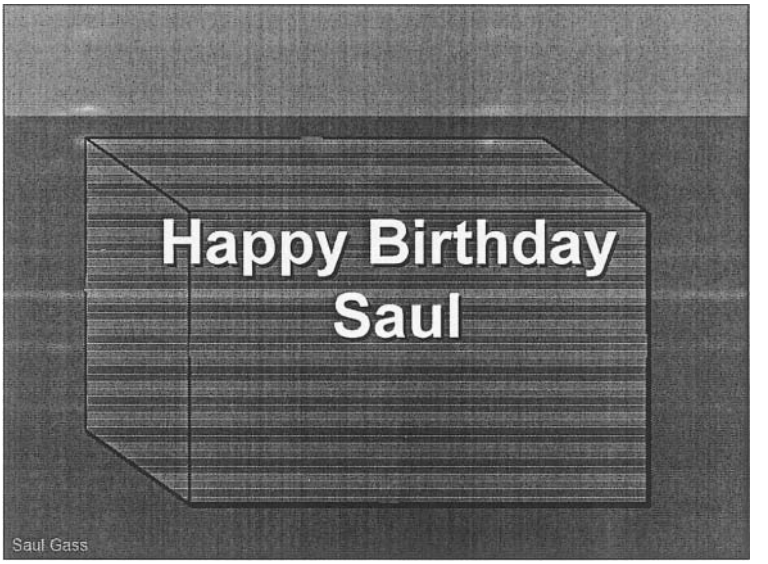


Fig. 3.

Today's talk is something of a reprise of that talk (at least in spirit), though I am not going to talk about education. I will talk about something else.

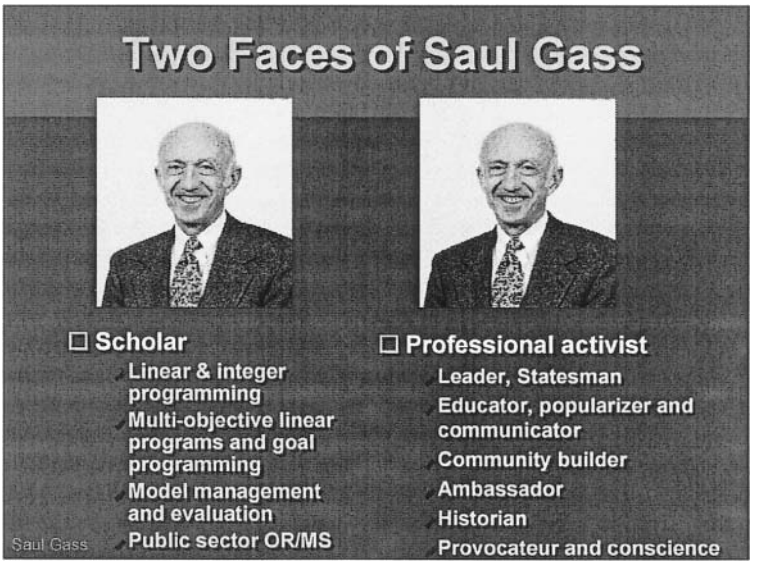


Fig. 4.

Figure 4 shows two faces of Saul Gass. (My wife told me that I've got to pick out the 10 points that are different in the two pictures of him.)

1. The first face represents Saul as a scholar of optimization (in particular, linear and integer programming, multi-objective linear programming, goal programming), model management and evaluation, and public sector OR/MS. I'm not sure if that's how Saul would categorize his scholarly work, but this is at least one cut at it. For this occasion, Arjang Assad will also be preparing an account of Saul's intellectual contributions. Saul as a scholar is just one face of Saul.
2. The other face of Saul is what I'll call the "professional activist." I say that in the most positive way. He is a leader and statesman; we've heard some of those words previously today. He is an educator/popularizer (if there is such a word) and a communicator. He does this brilliantly. In addition, he is a community builder, an ambassador, an historian, and a provocateur. The last, I think, is in some ways as important as all the rest.

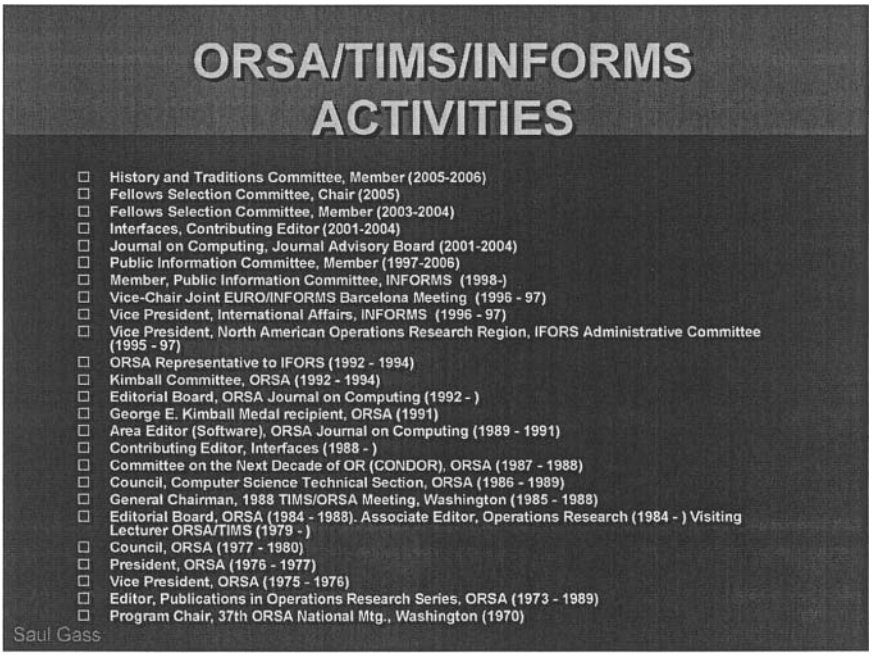


Fig. 5.

Figure 5 is a very "busy" chart, and the important thing for my purpose today is that it is busy. This is a list of some of the things (and I've probably missed several) that Saul has done over the years just for ORSA, TIMS, and INFORMS.

He was the 25th President of ORSA, which is really a nice number, isn't it? He's been involved in a variety of editorial activities and a variety of committees. For example, he's served as Vice President of International Affairs for INFORMS.

Most recently, he has headed the committee "responsible" for the appointment of INFORMS Fellows. By any account, that's a remarkable list, and that's just his activities on behalf of ORSA, TIMS, and INFORMS.



Fig. 6.

Figure 6 is a very long list of Saul's activities for the profession in general. An important aspect of this list is simply the extent and reach, not necessarily the specifics, of the things he has done for IFORS, Omega Rho, WORMSC, MPS, the Winter Simulation Conference, SIAM, and ACM. His imprint, as reflected by his contributions, is everywhere in the mathematical sciences broadly construed. He's done this in a wide variety of ways: serving in leadership capacities, serving in editorial capacities, and even being a trustee on a few occasions. Saul is clearly someone remarkably committed to the profession and remarkably active professionally.

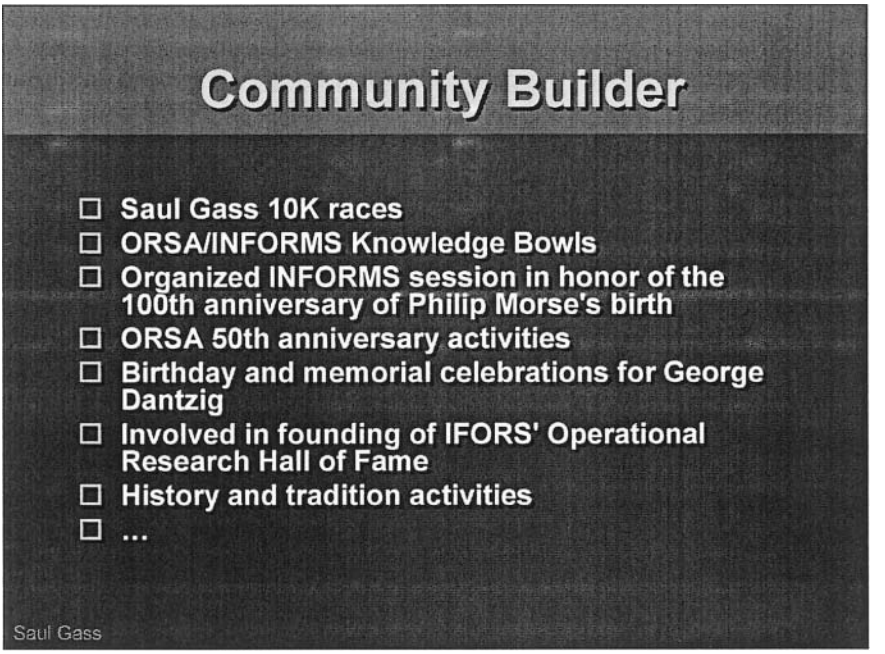


Fig. 7.

Figure 7 reveals some of the activities that we all know quite well and demonstrate Saul's talent as a community builder: the famous 10K race that Dick Larson previously commented about today and that we can all sign up for at the next INFORMS meeting (ha, ha); and the Knowledge Bowls. He shamed me into participating in one of these. I think some of you in the audience today were there on that occasion. Saul was the moderator, and our team was dismal. It was great fun to be embarrassed in public, and Saul did that to us. He was actively involved in ORSA's 50^{th} anniversary activities. He crafted a celebration for the 100^{th} anniversary of Phil Morse's birth. He's organized sessions — we've done a couple together — in recognition of George Dantzig. He was involved in the creation of the Operational Research Hall of Fame and has collaborated with Arjang Assad on the history and traditions of our profession. He is, it seems to me, the one person in our profession whose contributions could be portrayed in a chart with this remarkable intensity of activities — he has been at the core of community building and community awareness. Saul is someone who has reached out and brought us all together in ways that have been singular.

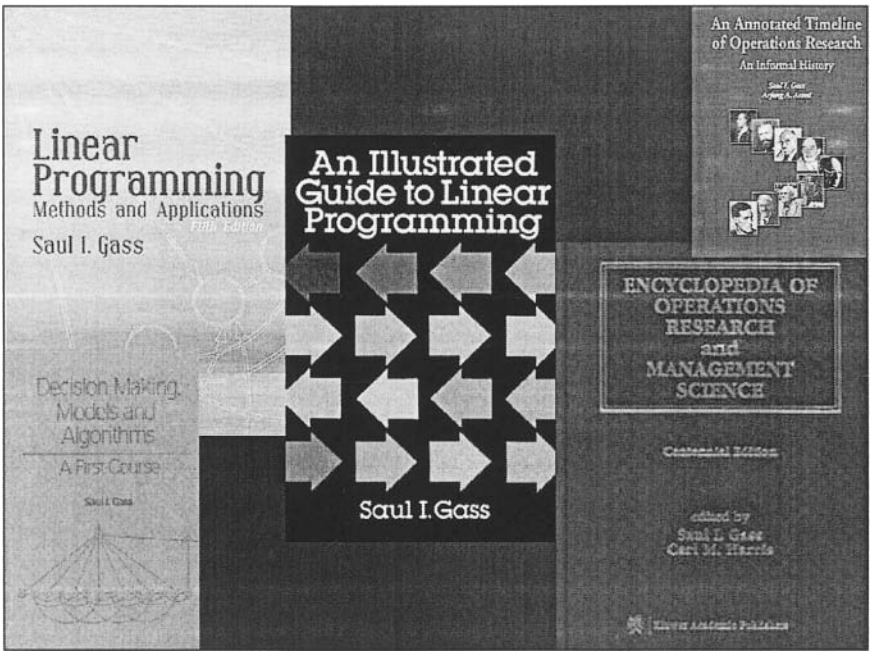


Fig. 8.

A famous book! A famous book! A famous book! A famous book! A famous book! One of the things that Saul has done remarkably well is either author or edit some very important books. Figure 8 provides us the names and a visual depiction of those books. The famous linear programming book is now in its 5th edition. I believe it was originally printed in 1958, and I think it was the first textbook on linear programming. This was the beginning for me. I've told the following story before.

I was a chemical engineering student at Syracuse University, and my undergraduate advisor at that time gave me this book. I looked at it a bit puzzled thinking, "What are all these lines and vectors?" I didn't know what any of these things were. I started reading through the book. My copy is gray (and contains lines and vectors); it doesn't have any of this fancy blue that he's got in his later editions, now that he's gone upscale. My version had the older look. I believe it was the second edition, the 1964 edition. It's still sitting in my bookcase, although I'm contemplating bringing it down to the bookstore to see if perhaps I can get a fine bottle of wine for it (as referenced before by Rudy Lamone). I bet the book is worth something at the wine store!

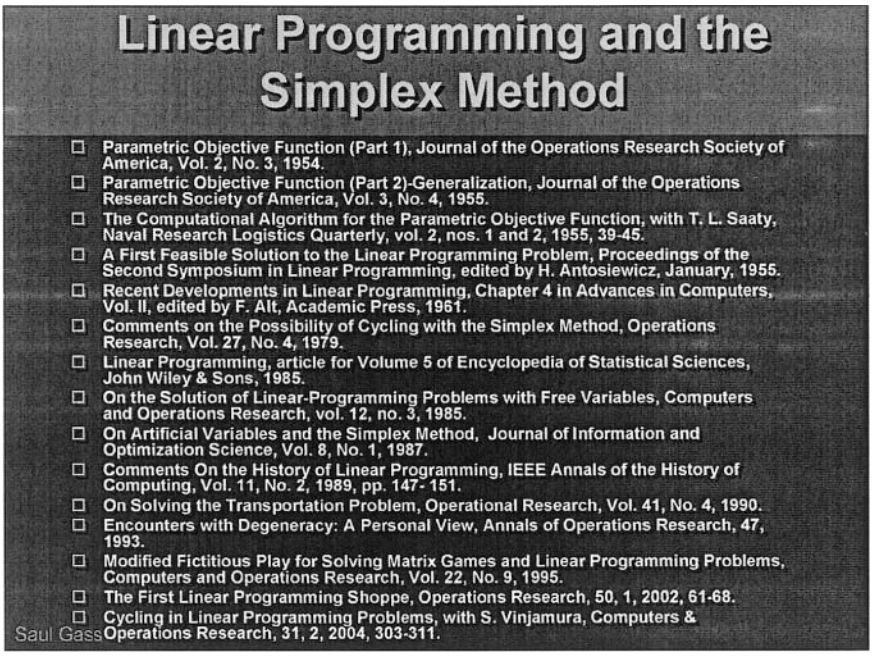


Fig. 9.

An important take away from Figure 9 is the large number of entries. These are contributions that Saul's made in the area of linear programming and the Simplex method. There's a variety of things here. I'm not going to talk about all of them, but I'll touch on a few.

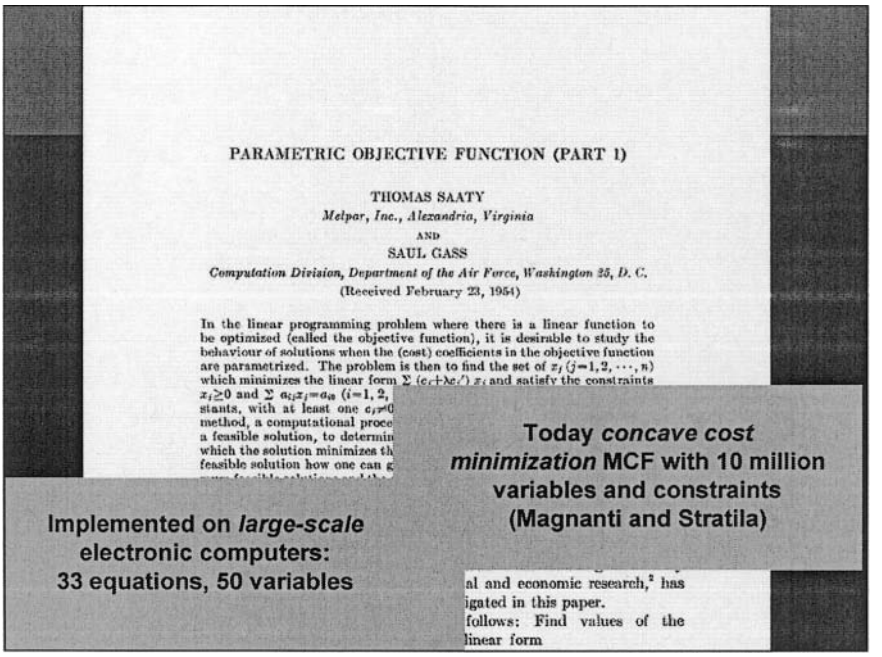
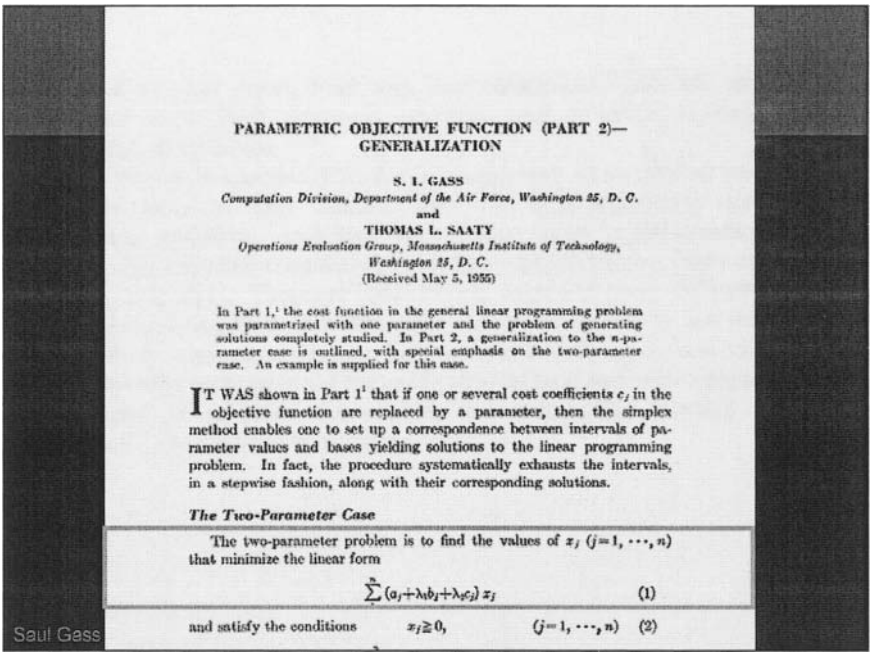


Fig. 10.

The paper shown in Figure 10 is very famous. Anyone who's ever taught linear programming has taught the parametric method. I've been spending the last few weeks reading Saul's papers, and it has been great fun. I understand this paper was done when Saul was at the U.S. Air Force and, at that time, Tom Saaty was a summer post doc student. This is the famous paper on the cost-parametric method. At the end of this paper, Saul and Tom mentioned that this method was implemented on a large electronic computer and solved a problem with 33 equations and 50 variables.

To put this in perspective, I've got this very, very smart student working with me by the name of Dan Stratila, who has been examining concave cost minimization problems, including notoriously hard, multi-commodity flow problems. He has been considering problems with 10 million flow variables and 10 million constraints. These are indeed really large and really difficult problems. Dan has been solving these to within a few percent of optimality and within only a few seconds of computation time. This tells us a little about how our field has changed since the early days and the publication of Saul and Tom's paper. With 33 equations and 50 variables, the machine they were using probably came to a grinding halt. The problem was just too big!



PARAMETRIC OBJECTIVE FUNCTION (PART 2)—
GENERALIZATION

S. I. GASS
Computation Division, Department of the Air Force, Washington 25, D. C.
and
THOMAS L. SAATY
Operations Evaluation Group, Massachusetts Institute of Technology, Washington 25, D. C.
(Received May 5, 1955)

In Part 1,[1] the cost function in the general linear programming problem was parametrized with one parameter and the problem of generating solutions completely studied. In Part 2, a generalization to the n-parameter case is outlined, with special emphasis on the two-parameter case. An example is supplied for this case.

IT WAS shown in Part 1[1] that if one or several cost coefficients c_j in the objective function are replaced by a parameter, then the simplex method enables one to set up a correspondence between intervals of parameter values and bases yielding solutions to the linear programming problem. In fact, the procedure systematically exhausts the intervals, in a stepwise fashion, along with their corresponding solutions.

The Two-Parameter Case

The two-parameter problem is to find the values of x_j $(j=1, \cdots, n)$ that minimize the linear form

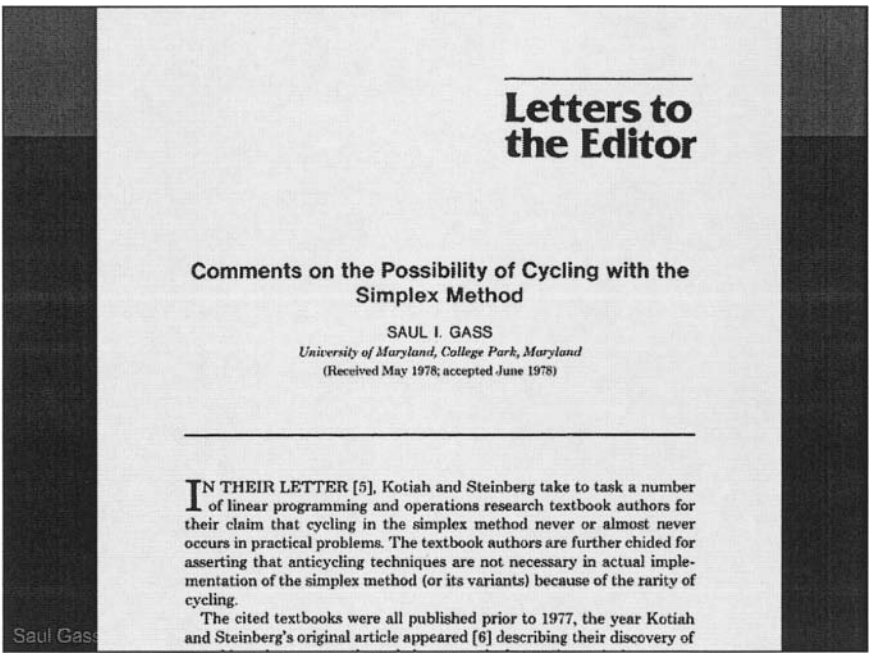
$$\sum_{j}^{n} (a_j+\lambda_1 b_j+\lambda_2 c_j)\, x_j \qquad (1)$$

and satisfy the conditions $x_j \geq 0, \quad (j=1, \cdots, n) \quad (2)$

Saul Gass

Fig. 11.

The subsequent paper shown in Figure 11 examines situations with two parameters instead of one; so now you can't just examine (basis) transitions on a line. You've got to look at the problem as a polyhedron, and it becomes more complicated. These papers represent one (important) aspect of Saul's work.



Letters to the Editor

Comments on the Possibility of Cycling with the Simplex Method

SAUL I. GASS
University of Maryland, College Park, Maryland
(Received May 1978; accepted June 1978)

IN THEIR LETTER [5], Kotiah and Steinberg take to task a number of linear programming and operations research textbook authors for their claim that cycling in the simplex method never or almost never occurs in practical problems. The textbook authors are further chided for asserting that anticycling techniques are not necessary in actual implementation of the simplex method (or its variants) because of the rarity of cycling.

The cited textbooks were all published prior to 1977, the year Kotiah and Steinberg's original article appeared [6] describing their discovery of

Saul Gass

Fig. 12.

The paper shown in Figure 12 is an interesting paper that Saul wrote on cycling in the Simplex method. Rather than focusing on new theory or algorithms, it is a discussion paper about whether or not cycling occurs. Saul makes the distinction between algorithmic cycling in terms of what we teach in textbooks and computer cycling in terms of how the computer solves the problem. He follows this up with subsequent data. Again, I think the type of things that Saul does so well is look at what's happening in the world, at how we're all thinking, and what we are all doing, and then add new insight and new understanding.

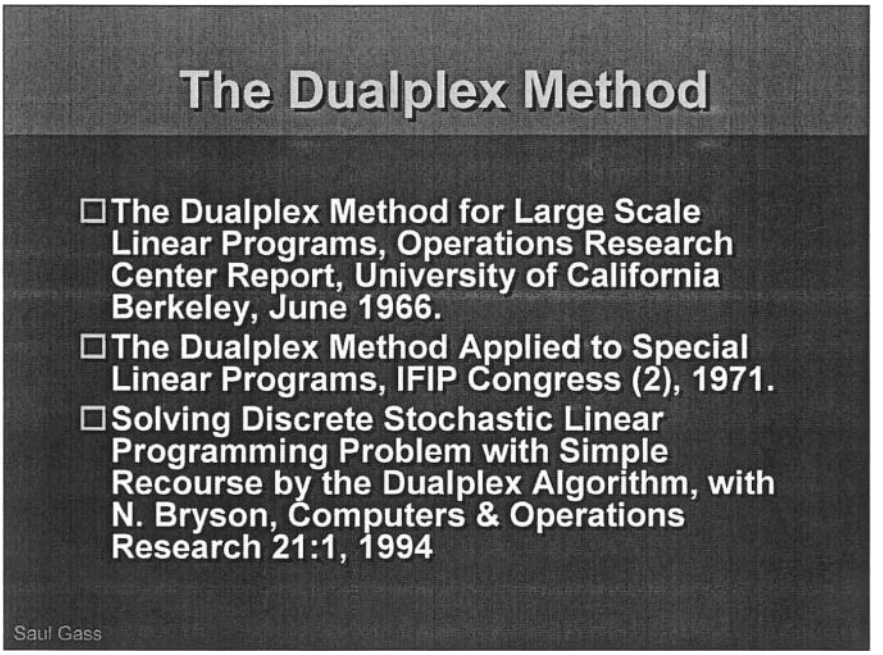


Fig. 13.

Figure 13 cites a couple of papers on the Dualplex method. To the best of my knowledge, this work has never been published except as a Berkeley report. There were subsequent papers in proceedings that described the use of the method, but not its development, per se. Papers like the third one are applications of the Dualplex method, in this case to an important, and generic, problem class.

I have a little story about this material. About four or five years ago, while I was in my current position as dean, I came across a paper that two fairly famous people in our field had written. Although it wasn't published yet, it was submitted to a journal, and although I wasn't a referee, it seemed very familiar to me. It had an application content, but also an algorithm. So I went back to the Dualplex method. The algorithm used in the submitted paper was indeed the Dualplex method. So, good things come up again and again.

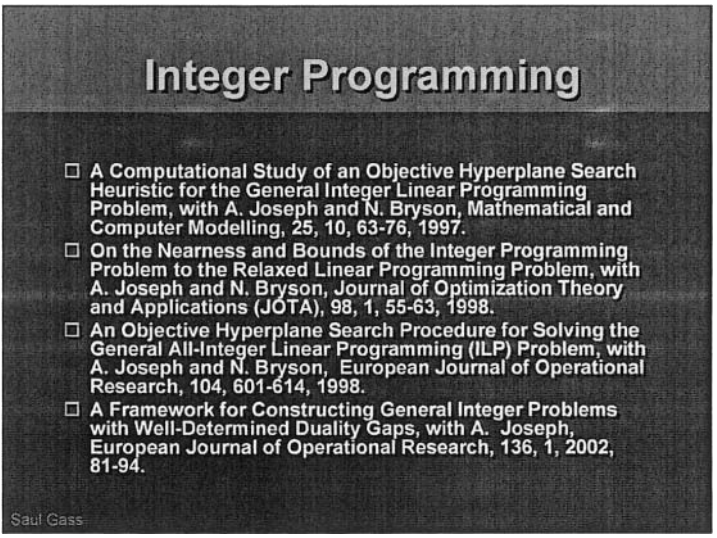


Fig. 14.

Figure 14 cites Saul's papers on integer programming. Note that these are fairly recent. These papers show both Saul's continued vitality and activity level, as well as the fact that Saul has changed his focus, moving from linear to integer programming.

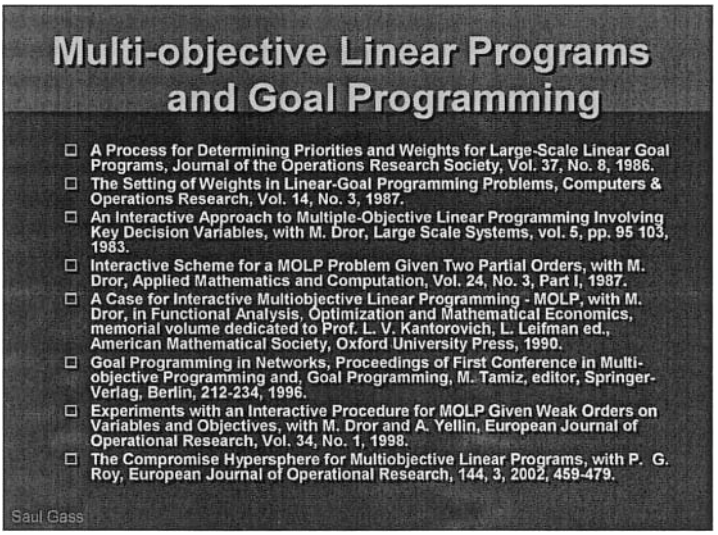
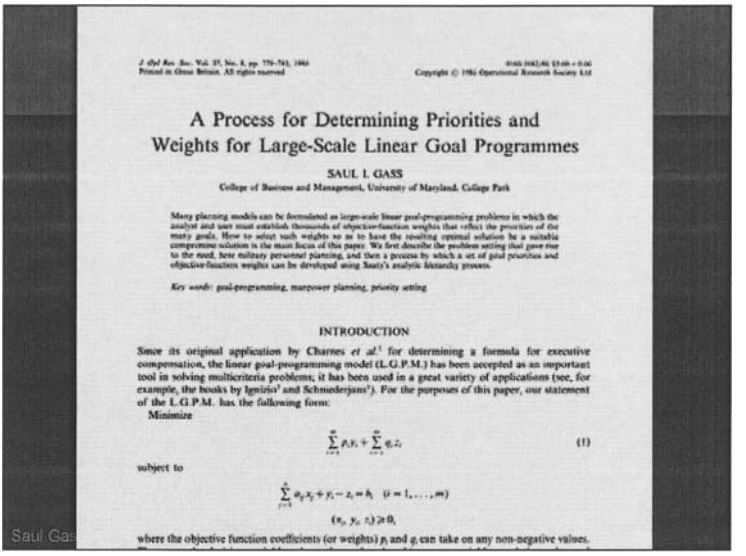


Fig. 15.

Figure 15 provides a list of Saul's papers in the area of multi-objective linear programming and goal programming. Figure 16 shows the introductory page of one paper that Saul has written in this area.



A Process for Determining Priorities and Weights for Large-Scale Linear Goal Programmes

SAUL I. GASS

College of Business and Management, University of Maryland, College Park

INTRODUCTION

Since its original application by Charnes *et al.*[1] for determining a formula for executive compensation, the linear goal-programming model (L.G.P.M.) has been accepted as an important tool in solving multicriteria problems; it has been used in a great variety of applications (see, for example, the books by Ignizio[2] and Schniederjans[3]). For the purposes of this paper, our statement of the L.G.P.M. has the following form:

Minimize

$$\sum_{i=1}^{m} p_i y_i + \sum_{i=1}^{m} q_i z_i \qquad (1)$$

subject to

$$\sum_{j=1}^{n} a_{ij} x_j + y_i - z_i = b_i \quad (i = 1, \ldots, m)$$

$$(x_j, y_i, z_i) \geq 0,$$

where the objective function coefficients (or weights) p_i and q_i can take on any non-negative values.

Saul Gas

Fig. 16.

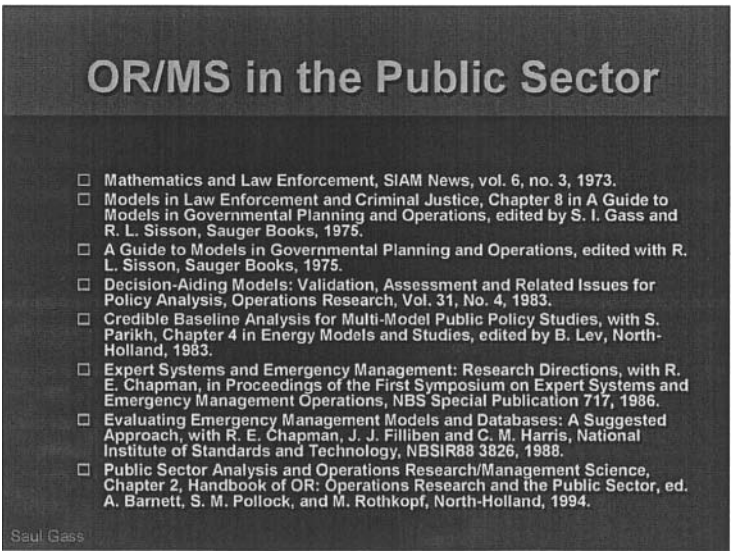


Fig. 17.

I'm not going to say much about the next area but will provide just a glimpse into an entirely different field in which Saul has worked. He's written not one or two but an entire series of papers on the public sector. In addition to what we've heard from Al Blumstein and Dick Larson today about Saul's work in this problem domain, Figure 17 lists his publications concerning the public sector. Even though his body of work has been well represented in the previous presentations today, my own modest account wouldn't be at all representative if I didn't at least mention these contributions.

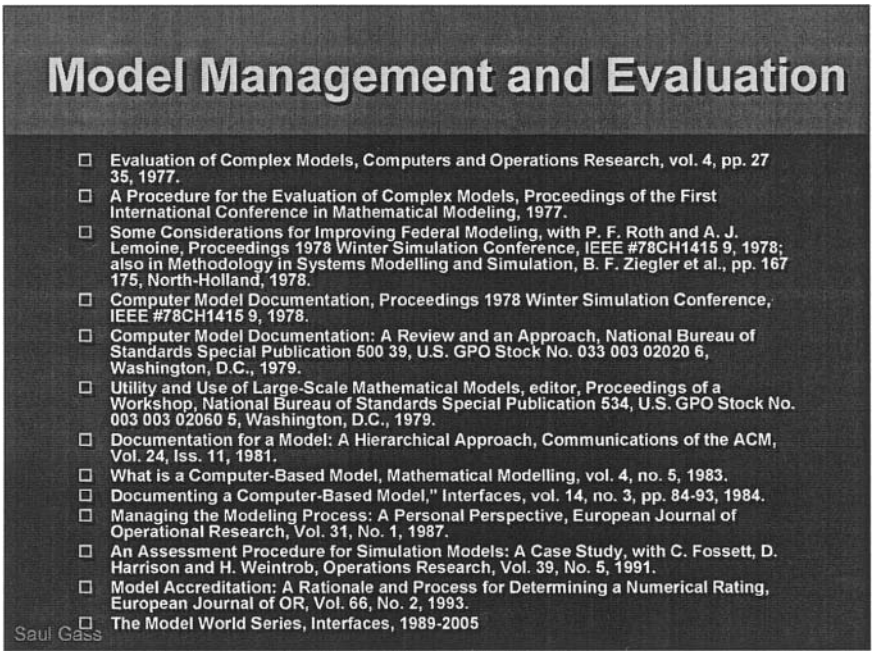
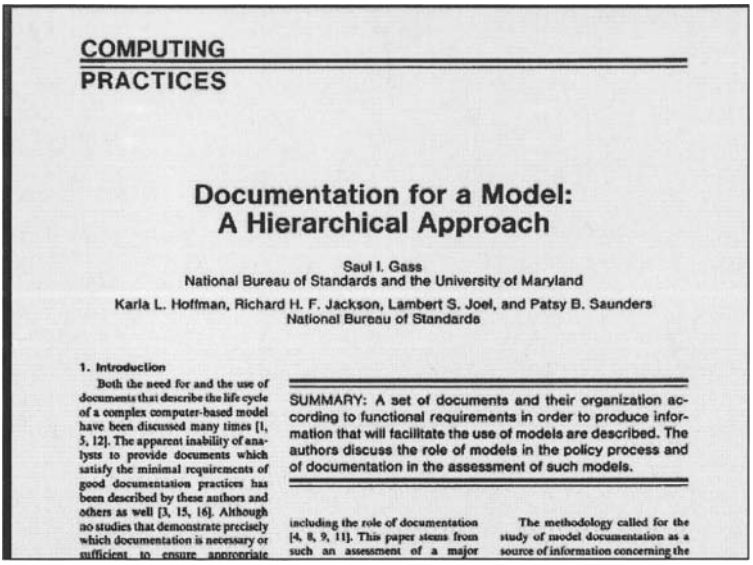


Fig. 18.

Another area in which Saul has had a significant impact is modeling. How does one think about modeling? How do you know when you have a good model? How do you know when you do not have a good model? How do you assess a model? How do you evaluate one? Some of the ideas Saul has developed on this topic are conceptual; some are approaches concerning how to do modeling. Again, I'm not going to describe any of the contributions listed in Figure 18. My intent is merely to acknowledge this body of work.



COMPUTING
PRACTICES

Documentation for a Model: A Hierarchical Approach

Saul I. Gass
National Bureau of Standards and the University of Maryland

Karla L. Hoffman, Richard H. F. Jackson, Lambert S. Joel, and Patsy B. Saunders
National Bureau of Standards

1. Introduction

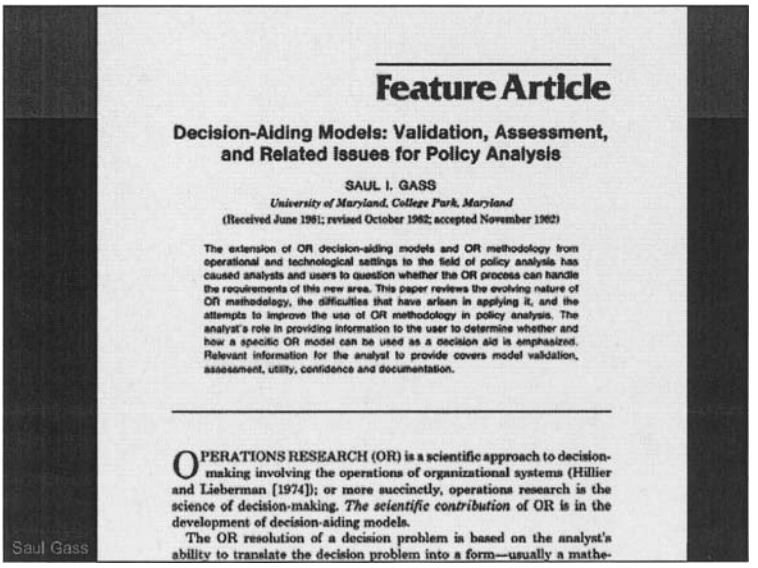
Both the need for and the use of documents that describe the life cycle of a complex computer-based model have been discussed many times [1, 5, 12]. The apparent inability of analysts to provide documents which satisfy the minimal requirements of good documentation practices has been described by these authors and others as well [3, 15, 16]. Although no studies that demonstrate precisely which documentation is necessary or sufficient to ensure appropriate

SUMMARY: A set of documents and their organization according to functional requirements in order to produce information that will facilitate the use of models are described. The authors discuss the role of models in the policy process and of documentation in the assessment of such models.

including the role of documentation [4, 8, 9, 11]. This paper stems from such an assessment of a major

The methodology called for the study of model documentation as a source of information concerning the

Fig. 19.



Feature Article

Decision-Aiding Models: Validation, Assessment, and Related Issues for Policy Analysis

SAUL I. GASS
University of Maryland, College Park, Maryland
(Received June 1981; revised October 1982; accepted November 1982)

The extension of OR decision-aiding models and OR methodology from operational and technological settings to the field of policy analysis has caused analysts and users to question whether the OR process can handle the requirements of this new area. This paper reviews the evolving nature of OR methodology, the difficulties that have arisen in applying it, and the attempts to improve the use of OR methodology in policy analysis. The analyst's role in providing information to the user to determine whether and how a specific OR model can be used as a decision aid is emphasized. Relevant information for the analyst to provide covers model validation, assessment, utility, confidence and documentation.

OPERATIONS RESEARCH (OR) is a scientific approach to decision-making involving the operations of organizational systems (Hillier and Lieberman [1974]); or more succinctly, operations research is the science of decision-making. *The scientific contribution* of OR is in the development of decision-aiding models.

The OR resolution of a decision problem is based on the analyst's ability to translate the decision problem into a form—usually a mathe-

Saul Gass

Fig. 20.

The paper with Karla Hoffman and others, as shown in Figure 19, examined a hierarchical approach to model documentation. This is a well-known paper. Saul

also wrote a featured article (Figure 20) in the journal *Operations Research*. One nice thing about this publication for me was that it appeared while I was the journal's editor. It was published during my first year as editor, and I'm very proud that this appeared as a featured article during my editorship.

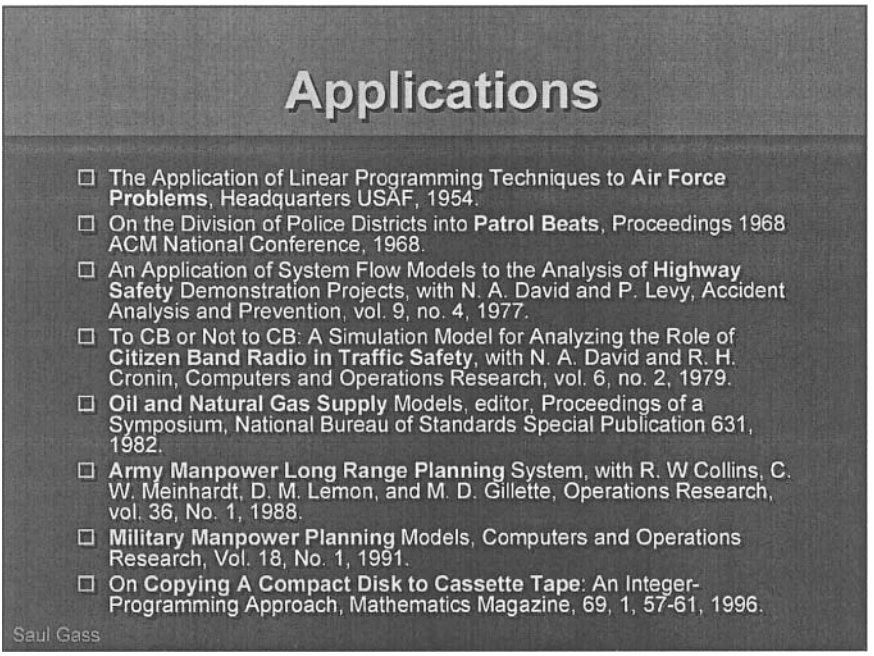


Fig. 21.

Figure 21 shows a set of applications that Saul has examined. I've tried to indicate in bold some sense of the range of these applications. In some ways, one of the most impressive things about this list is its eye-catching range — from Air Force problems, patrol beats, and highway safety, to citizen band radio (which seems related to the transportation theme I mentioned earlier), oil and natural gas, manpower planning, military planning, compact discs and cassette tapes. What a remarkable spectrum of investigations, both for Saul and for our profession!

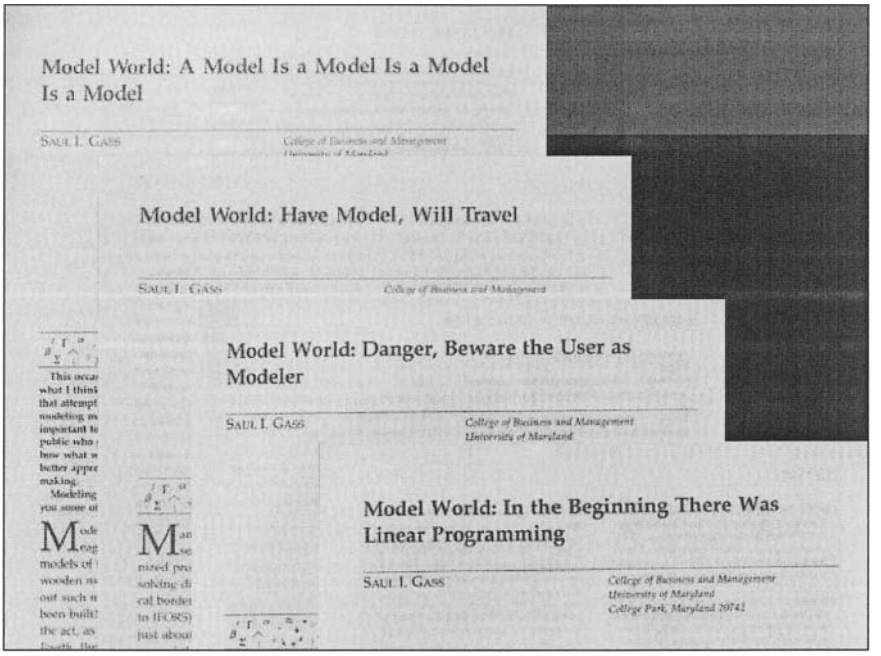


Fig. 22.

Another thing that Saul has done is serve, as I mentioned previously, as an expositor, an ambassador, and a communicator in a variety of forums. I think that perhaps one of the most impactful aspects of these contributions is captured in a series of papers he has written called *Model World* (as shown in Figure 22). I believe the first in that series was published in 1989. It asked the simple question, "What's a model?" and Saul gives us some insight into this question. That was the beginning. The second installment was, "*Have Model, Will Travel.*" This paper addressed modeling in developing countries and highlighted the fact that you could not just translate and transfer a model created in a developed country to a developing country. Saul addresses some of the cultural issues and some of the barriers that might occur in making this transition. Again, a thoughtful piece! The next installment is, "*Danger, Beware the User as Modeler.*" Is it a good thing or a bad thing that anyone can develop spreadsheets and create OR models? Saul discusses this issue in that article.

You can sit through an evening or two and read through all of these papers, and you will learn a lot, an awful lot, about our field.

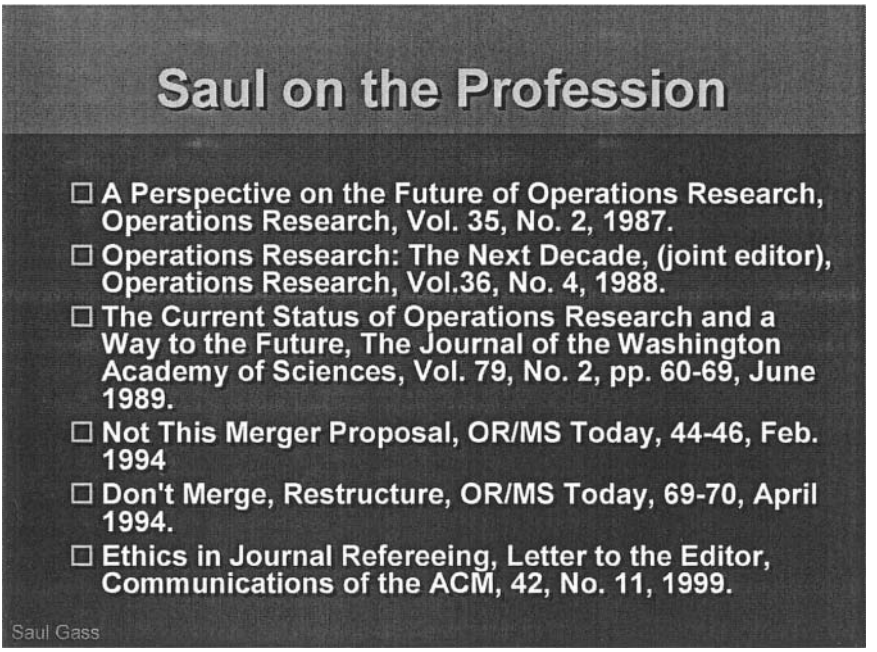


Fig. 23.

Figure 23 shows a bit about Saul and the profession. Again, he's written a set of perspective pieces, some projecting into the future. The second paper is the one in which we are co-authors. Thirty-seven of us contributed, and Saul helped lead this activity. Conceivably, this room contains half of the people who were involved with this project. The CONDOR Report examined the professional wave of the future. Items 4 and 5 show that Saul played the role of the loyal opposition when TIMS and ORSA were being merged, and his input had some fairly significant impact on the character of the merger and helped to improve it. I remember vividly those conversations because it came from his heartfelt commitment to Operations Research/ORSA and what the profession meant, and his concern that we might be losing some of its character and substance through the merger. Although he and I were on opposite sides of the merger debate, his passion and formidable articulation was simply wonderful. In addition to these many other professional contributions, he has also sparked considerable dialog concerning ethics in the profession in a variety of ways.



Fig. 24.

Figure 24 is in some ways for me where it all started, because there's the poppa and there's the guy who wrote the book that got me into this field. This might not be the beginning for the rest of you, but I think we could argue that linear programming started with those two wonderful gentlemen standing before you in this slide.

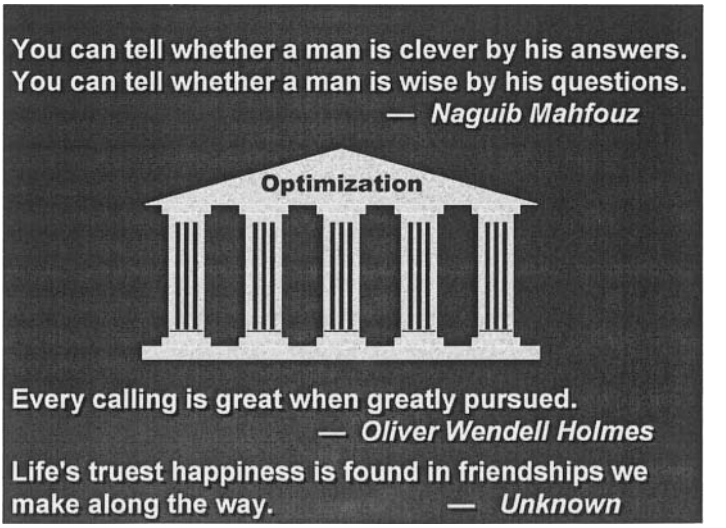


Fig. 25.

I will close with the three quotes in Figure 25 that surround the optimization house that Saul has built. They somehow say it all for me. Saul has asked penetrating questions, penetrating theoretical questions, penetrating questions about models, and penetrating questions about applications; he has also asked penetrating questions about the profession. Along the way he has contributed the answers to some of these questions and insight about many of them. But always penetrating, stimulating, and provocative questions! Saul has greatly pursued this thing we call Operations Research and in doing so has helped make the profession itself great.

Happy Birthday, my friend, and thanks for all that you have and will continue to do for all of us.

Looking Backwards, Looking Forwards: Reflections on Definitions of Operations Research by Morse and Kimball

Richard Larson

Massachusetts Institute of Technology
Building E40-231b
Cambridge, MA 02139-4307
rclarson@mit.edu

Summary. This paper is an edited transcription of the author's lecture given on February 25th, 2006, in honor of Saul Gass.

Key words: Real problems; scientific method; interdisciplinary teams; hot topics.

1 Introduction

I hesitate to talk at all about history today, because the biggest, the greatest historian about OR I know is sitting here in the front row: Saul Gass himself. Therefore, you will have to excuse me and give me a B+ or perhaps even a B- if I do this wrong.

As Al Blumstein remarked earlier today, we met in the summer of 1966 when Al was in charge of the Science and Technology Task Force of the President's Commission on Law Enforcement and Administration of Justice. Saul was there along with Ron Christensen, who happened to be a UC Berkeley physicist, and others. We weren't just focused on toy problems but real problems.

It was really great for me to have two mentors, such as Saul and Al, early on in my life, even though I was only a young child at the time. Saul was quite interesting. He was always thinking about models. Now I know this isn't politically correct, but it's excusable since it's in his book, *An Illustrated Guide to Linear Programming*. For those of you who weren't around, you can find it (Figure 1) in his book.



Fig. 1.

The young Saul Gass is shown in Figure 2, and he really hasn't changed over the years. If Trudy were here, I would ask her to comment on that handsome fellow shown there.



Fig. 2.

As shown in Figure 3, he was always using the latest computer technology because with linear programming you must do that. Figure 3 is also out of his book so I know it's politically correct to use it. I also know your lawyers will be calling me in the morning about a possible copyright violation since it's out of his book, but Saul has granted me permission and he owns the copyright.

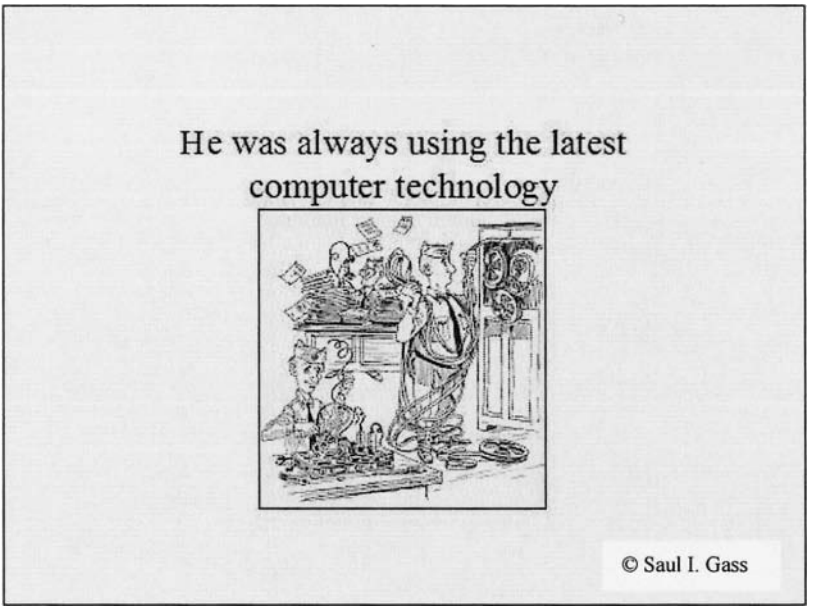


Fig. 3.

This talk is kind of like a sandwich. There is the top piece of bread, then we have the stuff that's in the middle of the sandwich, and then we have the bottom piece of bread. That was the end of phase one, which took very little time.

Basically my theme is that OR works best when it's working on real problems. Saul, being a modeler for all of his life, also knows the history very well. He knows the history of OR in Great Britain and the United States during WWII. The biggest accomplishments were made by people working on real problems, not toy problems, not virtual entities manufactured out of the imagination and transferred to the blackboard. As shown in Figure 4, that's my extreme point here: OR works on real problems. As Al Blumstein said earlier today, Operations Research is research on operations. I know that's heresy to some, but that really describes the essence of Operations Research.



Fig. 4.

The circumstance that brought me to the Science and Technology Task Force was a phone call between Philip Morse and the Institute for Defense Analysis when I was a starting graduate student at M.I.T. Phil Morse was a physicist on the U.S. side and was very important in creating OR as a named field. He was also the founding President of ORSA. In his presidential inauguration, Phil Morse said, "Operations Research is what operations researchers do." That quote was not quite so informative as what is shown in Figure 5. Since there are different people who may have said these quotes earlier, and since we don't know for how long this particular book, *Methods of Operations Research*, was classified, it is uncertain who said these quotes first. Some of the quotes follow.

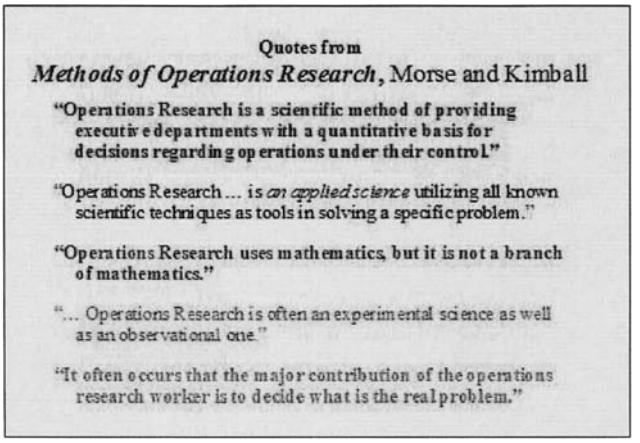


Fig. 5.

The first quote notes that OR utilizes the scientific method. The second emphasizes that OR is an applied science. The third emphasizes that OR uses but is not a branch of mathematics. Please don't be too concerned if you are from a mathematics department. The fourth quote is quite interesting because it points out that OR is often an experimental as well as an observational science. If you go back and see what they were doing in WWII, you will discover that almost all of the time they started by observing first, then collecting data and then looking at what the data is saying. I particularly like the last quote that tells us to determine the real problem. Those are the founding pillars.

The first OR book published in the US appeared over 50 years ago. It pertains to a topic that's dear to my heart, Queueing. I like that word because if you spell it the British way, Q-U-E-U-E-I-N-G, it has 5 vowels in a row. This is the only word I know that has that property. A. K. Erlang was a Danish telephone engineer who invented queueing theory. Even though he didn't know anything about Queueing, the Copenhagen Telephone Committee asked him to find a scientific way to figure out the size of the new central switching systems so that people didn't have to hang a wire between their place and everyone they wanted to call. As a result, he invented queueing theory by determining the real problem. His solution was published in 1917. Although there have been numerous textbooks, and hundreds, if not thousands of articles, written on Queueing, Erlang's equations are still the most widely used in practice today.

Another real problem is the Chinese Postman Problem. The word "Chinese" appears in the problem description not because it pertains only to the postmen in China, but due to the fact that the problem was described by a Chinese author and appeared in *Chinese Mathematics* (1:273-277, 1962). The first sentence of this classic paper states: "When the author was plotting a diagram for a mailman's route, he discovered the following problem: 'A mailman has to cover his assigned segment before returning to the post office. The problem is to find the shortest walking distance for the mailman.'" Therefore, he discovered the mathematical definition of what we now call the Chinese Postman Problem and its first solution. The Chinese Postman Problem is a real problem and something we teach most of our OR students today in a network optimization course.

Another very real problem is a facility location problem described in a paper co-authored by Hua Lo-Keng that appeared in *Chinese Mathematics* (2:77-91, 1962). The paper's title is "Application of Mathematical Methods to Wheat Harvesting," and the following quote is from the first paragraph:

"...the work of wheat harvesting in the Peking suburbs was participated in by teachers and students...The objective ...was experimental use of mathematical methods in the selection of the threshing site most economical of transportation."

In this work, Hua Lo-Keng formulated the solution for the '1-center problem' on a tree and actually did much more. By the way, this fellow never graduated from high school, never got a college diploma, and yet became the most famous theoretical mathematician in China in the second half of the twentieth century. He is viewed by many as the founder of OR in China. The Hua Lo-Keng Prize in Mathematics is the highest honor awarded in the Chinese mathematical

community. This shows that you can be rigorous and still work on applied problems, and the applied problems can yield fundamental results.

By the way, the story of how Hua Lo-Keng got to be a full professor at Tsinghua University and received honorary doctorates here and there, while never having graduated from high school, is a very fascinating story to read.

Most of you have heard of Bernard Koopman, the 6th President of ORSA, who worked with Philip Morse and others during the WWII effort. His invention of Search Theory is mathematically rigorous, and Koopman was a very talented mathematician. However, I ask you to realize that the work was a multi-disciplinary team effort of psychologists and others in the Social Sciences as well as mathematicians, and involved measurements of pilots' visual performance and much more. The results were rigorous, fundamental and classified for 15 years. The work was published in the late 1950s in *Operations Research*. Someone else eventually received the Lanchester Prize for follow-on search theory work..

Recall the last quote in Figure 5 about what OR is from Morse and Kimball: "It often occurs that the major contribution of the operations research worker is to decide what is the real problem." This quote provides the theme of the next problem I am about to describe.

You probably all know the folklore by now about the elevator delays in New York City in the mid 1950s when people started living in high rise apartments, worked in high rise offices, and stayed in high rise hotels. The elevators sometimes even had elevator operators in them before the advent of complex computer controls. People started complaining about the delays in waiting for elevators, the old fashion ones. You know the story.

Russ Ackoff sent a junior person, whose name he forgets, from Wharton to New York to look at the problem. The person studied the problem. He looked at people waiting for elevators who were looking at their watches (that's why I don't wear one, no watch here), and other things. He said, "Aha!" A narrow technocratic approach would have been to say that this is a multi server-queueing system without enough servers. Therefore, let's dynamite the building and start all over again with twice as many elevator shafts. But this fellow with lateral thinking said, "The problem is not with the delays for elevators, but rather the complaints about the delays for elevators. If we can reduce the complaints about the delays for elevators, we'll solve the problem." Then in a spirit of great insight, he said, "Let's put mirrors next to the elevators and see if people can entertain themselves. If they have hair, they can comb it; if they have a tie, they can adjust it." It worked! They put mirrors next to elevators, and the complaints dropped to near zero, even though the statistics of the delays were unchanged. Problem solved! To me this is exemplary OR. How many theorems were proved? Zero. How many equations were there? Zero. How many variables? Zero. How much is common sense and lateral thinking? A lot! If this were written up today and submitted to any of our refereed journals, I think the referees would laugh out loud and toss it away. To me this is exemplary OR! We need to go back and think about these things.

Coincidentally, the year that this happened, 1955, was the year that Disneyland opened in Anaheim, and they have become the Machiavellian Masters of the psychology of queueing. There are a lot of industries that use psychology in

queueing, and I've listed some of them in Figure 6. The psychology of queueing actually has become a separate subfield in itself.

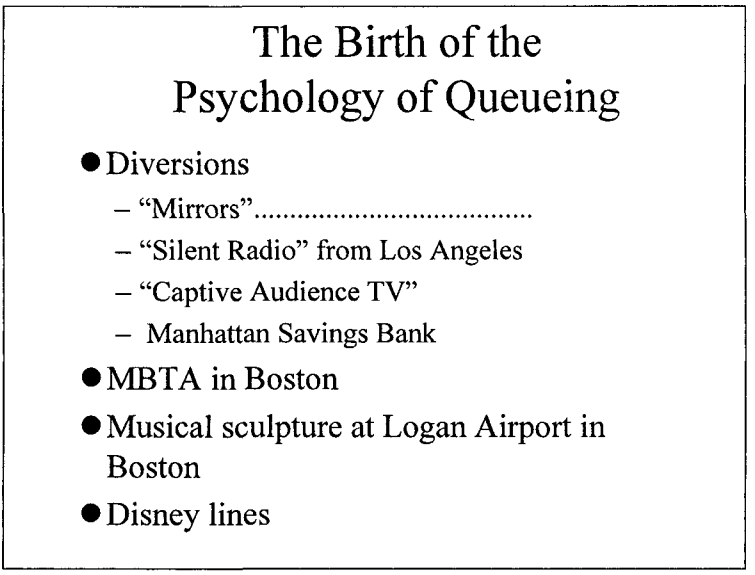


Fig. 6.

Now, Philip Morse authored or co-authored at least three books in OR and many articles. I previously gave quotes from the first book, coauthored with Kimball. The second one was *Queues Inventories and Maintenance* published in 1958. The third one that I know of is *Library Effectiveness: A Systems Approach*, which won the Lanchester Prize in 1968. How often in these books do you find the words, "Theorem" and "Proof"? So far I have been unable to find those words. Since they are not in electronic format and I couldn't do a computer search, they may be there somewhere but I could not find them. How often do you find those words in any physics text or paper? What I'm asking is: "Isn't it the job of OR people to discover the physics of complex systems involving people and technology?" Aren't we really physicists and not mathematicians or engineers? Might it be like the founders of OR in the late 1930s in the UK and the 1940s in the US who were physicists?

The dominant paradigm for an OR person should be that of a physicist who is both an experimentalist and a theoretician, rather than mathematician. Or, if he or she is a mathematician, it is their responsibiltiy to join interdisciplinary teams when looking, as Koopman did, at complex problems.

A proposition I have, as shown in Figure 7, is that "The Key Value Added Components of an OR Study are Three." One is framing the problem. Look at the elevator delays in New York City as an example. The second component is Problem Formulation. Sometimes if you're going to employ equations and

variables, that's when you reduce it to equations and constraints. The third component is Problem "Solution," often requiring optimization, which is dear to all of our hearts. What relative values would you place on each of these three components in doing a real OR study? We've heard two excellent ones earlier today, with the most recent being on spectrum bidding in auctions and the other in criminal justice systems. Each of these involved lots of work in all three of those steps. I ask you again, "What relative values would you place on each of these steps?" You each have your own personal weight that you'd put on them. Then I'd ask you to compare those weights to our emphasis when we teach OR students in our universities. How do those weights compare? Are we teaching them correctly? Are we emphasizing the right things? I leave that as a question to ponder.

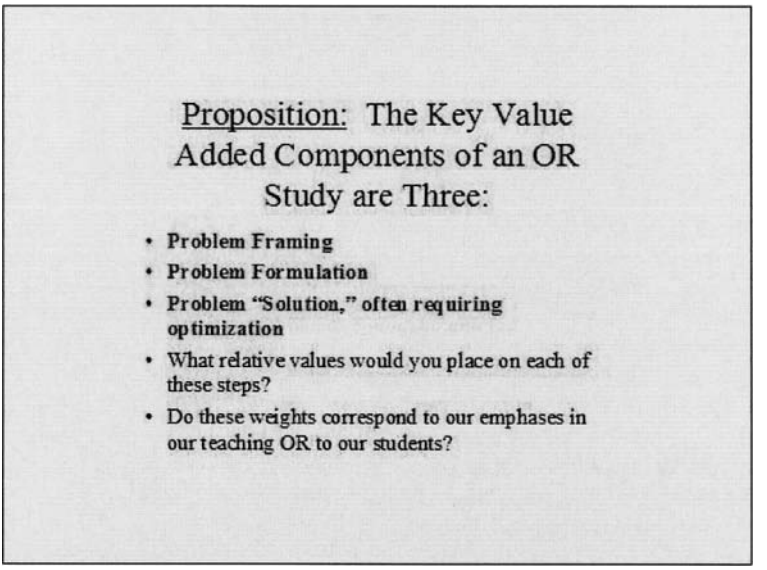


Fig. 7.

The theme of this part of the sandwich is that OR does solve real problems. My assertion is that OR's best theoretical work has been driven by real problems. Therefore, I reject the partitioning of INFORMS members and OR people into practitioners and academics. I don't think that is helpful, and I don't think it's realistic because the best OR people are both contributors to the methodology and solvers of real problems. Those other things I mention in Figure 8 also were driven by trying to solve real problems.

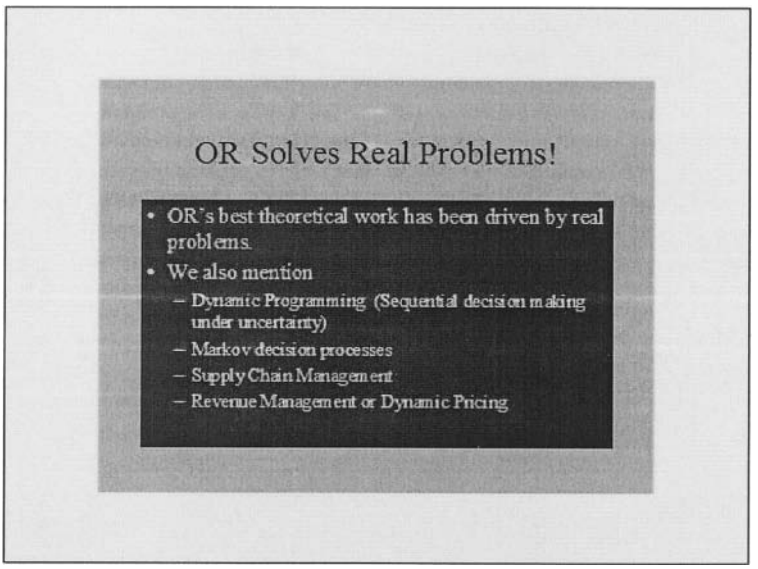


Fig. 8.

That was back in the 1950s. If we jump ahead to today, we see that we still do traditional OR work but things have really expanded in both depth and breadth. For instance, the following list shows all the refereed journals that INFORMS publishes today.

- Decision Analysis: A Journal of the Institute for Operations Research and the Management Sciences
- INFORMS Journal on Computing
- Information Systems Research
- Mathematics of Operations Research
- Transportation Science
- Organization Science
- Management Science
- INFORMS Transactions on Education
- Operations Research
- Marketing Science Journal
- Interfaces: Where Practice and Theory Converge
- Manufacturing and Service Operations Management

That is a very impressive list!

About a year and a half ago, I polled some of the editors-in-chief regarding hot topics. After all, we need to know what some of the hot topics currently are in OR. One of the editors was Larry Wein, who was still the Editor-in-Chief of *Operations*

Research. Larry gave me some of his hot topics. Without any editing, this is what he said:

- Homeland Security
- Call centers
- Internet modeling
- Revenue management
- Game-theoretic supply-chain management

Since those are all real problems, it's nice to see them. From Dr. Nimrod Megiddo, the Editor-in-Chief of *Mathematics of Operations Research*, which is the most theoretical of the publications that INFORMS publishes, we see more current hot topics:

- Internet modeling (heavy tail distributions)
- Auction theory
- Financial engineering
- "Price of anarchy"
  - Aims at analyzing the difference between performance under "selfish behavior" and under coordinated optimization. The methods here are game theoretic, involving Nash equilibrium and competitive equilibrium. Methods of both discrete optimization and continuous optimization arise in the analysis.

Even though it is a somewhat theoretical list in terms of the mathematics used, the topics are driven by real problems. *Management Science* is another flagship journal of our profession. The following list of even more current hot topics is from Dr. Wallace J. Hopp, Editor-in-Chief:

- **Social networks** (particularly as they relate to quantitative models of organizational structures).
- **Risk management** (e.g., use of finance and other tools to quantify and mitigate risk in a range of practical decisions beyond asset pricing and portfolio management)
- **Data mining** (the massive data bases becoming available via enterprise data warehouses and CRM systems ensure that the basic idea will remain of interest for quite some time).
- **Strategic planning** (much recent work has incorporated economic frameworks to consider competition and other high level market effects, but we still need research that brings these insights to the decision support level so that decision makers can use them in practical settings).
- **Service operations** (call centers and health care have been getting lots of attention in recent years, but with the growth of the service sector there are many other service industries where analysis could have a significant impact, such as primary/secondary education, retail sales, consulting, food service, etc.).

It is true though that if you open, let's say the flagship journal, *Operations Research,* to a randomly selected article, sometimes the only data you will see in that article is the date of submission and the date of acceptance. It would be nice to revisit what our founders did in the 40's and 50's and perhaps add some more data. Now Al Blumstein is a huge counter-example, as is Karla Hoffman. Again, you open the journal and sometimes the amount of data that you see, the data that is driving the analysis, is not as much as you might like.

Here is the good news! We have over a dozen world class refereed journals today. We have scores of textbooks, covering virtually every aspect of OR theory. We have approximately 7 generations of OR Ph.D.'s. I think I'm a great-great-grandfather, if you think about fathering Ph.D. students. There are others, believe it or not, who are older than me who have started this process earlier. We also have many specialty sub-societies. That's the good news! So what's the bad news?

That's also the bad news because what has happened is that a lot of this has channeled us into small canyons where we have little fiefdoms or tribes going around in these little canyons and they can't see over the edge of the canyons to the people in adjacent canyons. The idea of inter-disciplinary and cross-disciplinary teams that was so successful during the war effort and is still successful today in lots of consulting and other companies seems to be lost as the dominant paradigm of the profession. You might say, "Where are all the psychologists, social scientists, physicists, economists, empirically based research results, multi-disciplinary research teams?" We do have them here and there, but they no longer represent the dominant paradigm for OR, and I think at our risk.

I got an e-mail yesterday from an M.I.T. PhD in theoretical physics. He said he'd read some of my discussion of this topic last year in *OR/MS Today* and said, "I'm raising my hand; I'm converting to OR; tell me how to do it." So, physicists stand up and step forward. I've got to write him a serious response. Saul, any suggestions?

Problems still exist. There are myriads of problems out there that still need to be addressed in their full complexity, not some subset about which it is easy to write equations and prove theorems. You might want to check out the newly formed Council of Engineering Systems Universities (CESUN) http://www.cesun.org. I think there are about twenty-five members now, including MIT, Stanford, Georgia Tech, and RPI. I don't know if the University of Maryland is a member of CESUN or not. This movement in "Engineering Systems" embraces complex 'messy problems' in their entirety, utilizing an interdisciplinary approach by bringing in three different disciplines that are traditionally separate: OR, Industrial Engineering and other traditional engineering disciplines; Management, broadly defined; and the Social Sciences.

When I read through the literature that describes this new entity that is still less than a year old, I read that they are embracing an interdisciplinary approach, and then I look at these ancient documents that Saul sent me. One is by Philip Morse, "Of Men and Machines" and was published in 1946 at MIT (*MIT Technology Review*, 49, 1, 1946). To me, this description of OR in 1946, as shown in Figure 9, is virtually identical to what this is all about. Why is it that we have lost the

breadth and multi-disciplinary nature in studying complex problems? By the way, I am the official liaison representative between INFORMS and CESUN.

Philip M. Morse, 1946, *MIT Technology Review*)

"...the method of science is applicable to social as well as technological problems."

"The procedure which was developed for studying an operation is an obvious one...: The operation is analyzed to find out what happens rather than what should happen; the combined knowledge of physics, chemistry, psychology, and other sciences is utilized to find out why a particular operation occurs the way it does; this deeper understanding suggests possible causal relationships which might be tested by operational experiments, and finally tested understanding can then be used to modify the operation to give best results under various conditions."

"The new ingredient (from the War effort) was the cooperative effort of teams of specialists, drawn from any and all branches of science. In the solution of important problems, these teams were free to carry on a program of operations in which considerable emphasis was placed on the experimental method, and they combined their knowledge and techniques to achieve the maximum desired goal in the shortest period."

Fig. 9.

Those of us from the President's Crime Commission are Still Playing Cops and Robbers and Chasing Ambulances! Example: The hypercube queueing model is something we created at MIT in the mid 1970s and is still alive and well. We're adapting it to provide a transient solution for transient situations such as terrorist attacks in major cities. This work is sponsored by the Department of Homeland Security so it is being used and also generalized. By the way, Figure 10 shows a version of a hypercube that was implemented in the five boroughs of New York City to optimally deploy ambulances. By looking at Figure 10, you get a feel for what it looks like on a computer screen. It helps the 911 system in deploying emergency vehicles. Now, if you have dyslexia, you should go to Japan because their 911 is 119.

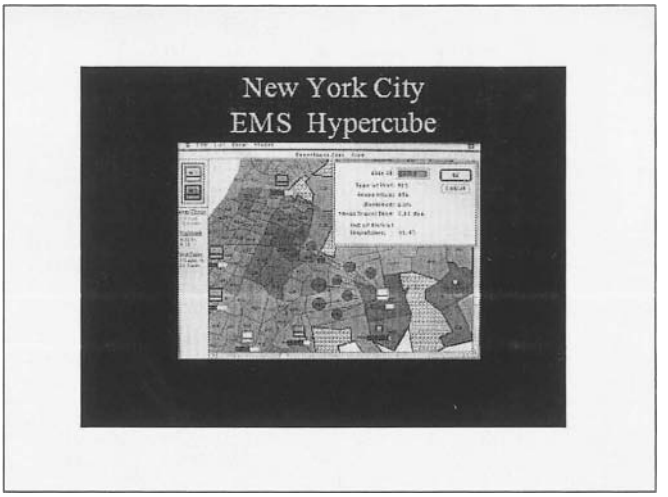


Fig. 10.

In keeping the spirit of the President's Crime Commission alive, yours truly recently visited Japan, as shown in Figure 11. By the way, I managed to give away a lot of hats in many different countries which read, "OR is the Science of Better".



Fig. 11.

We actually have used linear programming in some of our stochastic modeling. For instance, Rick Jarvis, a doctoral student of mine, used LP to generalize to end servers within the hypercube framework. Rick's work generalized the results that Carter, Chaiken, and Ignall got at the New York City Rand Institute in 1972. This

work uses Markov Decision Processes. Basically, it shows how you partition a city into zones so if you get a call for service or a call for a fire engine or for an ambulance and you know what the set of available servers are, which one do you send? It is not optimal to take the greedy heuristic of always sending the closest available server. Sometimes, it is best to send the second or even the third closest one available if your overall long term average is to minimize response time over the ensemble of all calls received over say 24 hours.

Due to Saul's excellent and best selling textbooks including his famous *Illustrated Guide to Linear Programming* and his great leadership in teaching and other accomplishments, he's been a major contributor to the correct modeling of applications that require optimization and to solving real-world problems.

We're going to celebrate, as we close here, some of LP's applications. There's supply chain management. The photo shown in Figure 12 was taken in Pakistan about 10 days ago, and that fellow is delivering milk. That is a Supply Chain Management real problem.

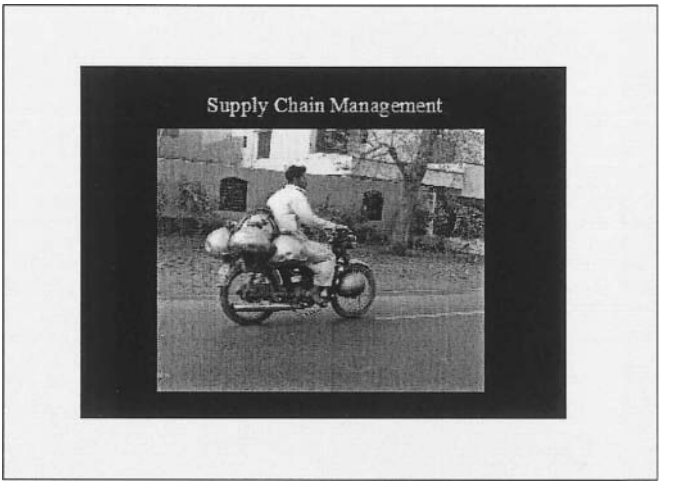


Fig. 12.

Other applications include manufacturing (such as the printing of currency and lottery tickets), inventory theory (such as the amount of grain and hops needed at Guiness Brewery), scheduling classrooms and operating rooms, water distribution systems, facilities layout, and the diet problem shown in Figure 13.

The Diet Problem is one of the earlier applications of LP and the objective is to minimize the cost of the food you buy subject to certain dietary constraints. I remember the first time they worked it out. It turned out you were just eating something like chalk all of the time. It was really awful and not very diverse. With the diet problem, they sometimes don't get the constraints right. When this happens and people follow it routinely without really understanding the black box from which it came, bad things can happen. As shown in Figure 14, becoming a

skeleton is one of the bad things that can happen if you use incorrect constraints on the diet problem!

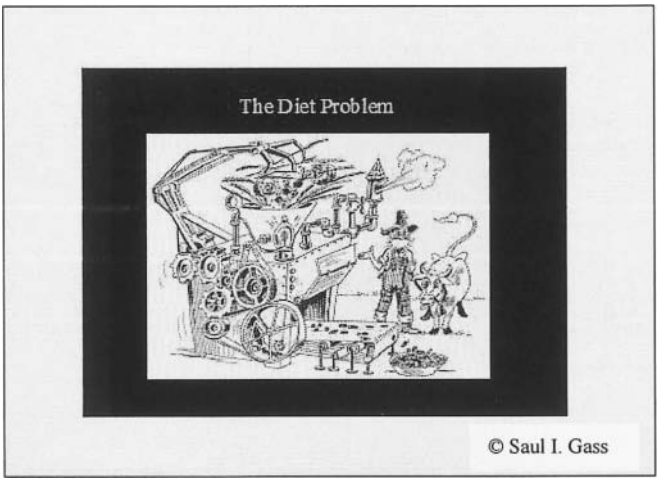


Fig. 13.

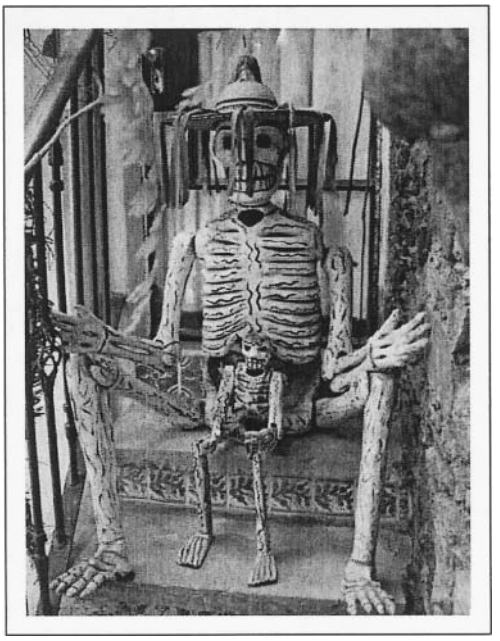


Fig. 14.

I want to thank Saul for 80 years of great OR and I am looking forward to another 80. I also want to thank Trudy, his optimal solution. Thank you very much!

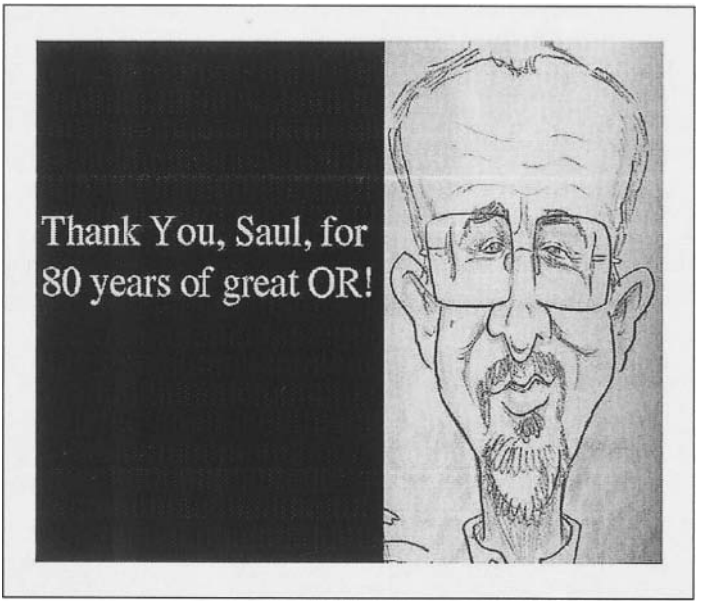


Fig. 15.

Ben Franklin: America's First Operations Researcher

Bruce L. Golden

Robert H. Smith School of Business
Department of Decision and Information Technologies
University of Maryland
College Park, MD 20742
bgolden@rhsmith.umd.edu

Summary. This paper is an edited transcription of the author's lecture given on February 25th, 2006, in honor of Saul Gass.

Key words: Ben Franklin; Franklin squares; constraint satisfaction; mathematical programming; decision making; data-driven modeling.

1 Introduction

I'm really happy to be here and to present a talk in honor of Saul Gass. I'm going to talk about Ben Franklin, and I'm going to argue, or at least I'm going to sketch an argument in support of Ben Franklin's being America's first operations researcher, defining operations research rather broadly. Let me start off by giving you some background. I graduated from the University of Pennsylvania as an undergrad. And if you're on the Penn campus, you're never very far away from a statue of Ben Franklin. And so I always thought I knew about Ben Franklin's life. I read his short autobiography when I was in college, but this past summer I read a full-length biography of Ben Franklin by Walter Isaacson [1]. I was frankly blown away by what this man accomplished in a single lifetime, and I was amazed at the range and variety of his talents. Some of you may know that we recently celebrated his 300th birthday. This is a very big deal in Philadelphia; I guess across the country it's not as well known, and it hasn't been as widely celebrated. But he was born on January 17, 1706.

So the more I thought about Ben Franklin, the more I felt that we could actually make a good case for his being America's first operations researcher. The plan for the talk is as follows: We'll take a quick look at the Ben Franklin you already know about, and then we'll take a peek at some of Ben Franklin's work in operations research. I suspect that many of you don't know about this work. And finally we'll try to answer the question, "So, what does this have to do with Saul Gass, anyway?"

2 Discussion

Okay, so what do we know about Ben Franklin? He was born in Boston. (Saul was also born in Boston.) He was the best scientist, the best inventor, the best diplomat, the best writer, and the best businessman in America in the 1700s. He was a very successful printer and publisher. Indeed, he became quite wealthy and retired at a very early age. He was a great political and practical thinker. He proved that lightning was electricity. We've all read the stories about his kite experiment. He had a number of important inventions. He founded a library, a college – namely the University of Pennsylvania in 1740, a fire department, and many other civic associations.

He's the only person to have signed all of these documents: the Declaration of Independence, the Constitution of the United States, the Treaty of Alliance with France (this guaranteed that the French would support the colonists during the Revolutionary War, which was very important because the colonists didn't have any money), and the Treaty of Paris with Great Britain that officially ended the war. So these four documents essentially assured the independence of the colonies, and he's the only one to have signed all four of them. As I mentioned, he retired early from business at age 42, and he devoted the next 42 years of his life to conducting experiments in science and providing service to Pennsylvania and then his country. But he also made, in his spare time, some significant contributions to the field of recreational mathematics.

Let me give you some background. In 1750, the population of London was about 750,000 people; there was only one city larger, and that was Beijing. What do you think the population of Philadelphia was in 1750? It was about 20,000. Imagine the difference in terms of level of sophistication between London and Philadelphia. We would now regard this as a town or maybe a small village, but Philadelphia was the largest city in the Americas in 1750. Of course, the difference between 20,000 and 750,000 is just incredible.

Now, how old was Benjamin Franklin when he ended his formal education? It turns out that he was 10 years old. He went to work for his dad, who was a candle maker, at age 10. Then, at age 14, he became an apprentice to his brother who was a printer. He remained in printing for the rest of his working life. In his spare time, he made some significant contributions to recreational mathematics. He designed these magic circles and magic squares, which were regarded as works of art by his friends, and even today some of them are regarded as masterpieces in the field of recreational mathematics.

Now, his explanation was that he was an assemblyman in Pennsylvania in the early 1730s, and he had to attend these incredibly boring meetings. So, he would start to doodle, and he would construct these remarkable magic squares.

With a magic square like the one in Figure 1, we have an 8 by 8 grid, and we want to assign the numbers 1 through 64 to the cells of the grid such that a variety of properties are satisfied.

So here you see that each row and each column will add up to a magic sum. The magic sum is the sum of numbers 1 through 64 divided by 8, because there are eight rows and eight columns. That is, every row and every column sums to 260

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 1. Franklin Magic Square

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 2. Each row and each column sum to 260.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 3. Any half-row or half-column sums to 130.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 4. Any 2 by 2 block sums to 130.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 5. Any wrap-around 2 by 2 block sums to 130.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 6. Any set of four cells forming a rectangle with an even number of cells on each side sums to 130.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 7. Any set of four cells forming a rectangle with an even number of cells on each side sums to 130.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 8. The shaded entries sum to 260.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 9. Each of the two bent rows above totals 260.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 10. Each of the two bent columns above totals 260.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 11. Each of the five bent parallel columns above totals 260.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 12. Each of the three wrap-around bent columns totals 260.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 13. Each of the five bent parallel rows above totals 260.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 14. Each of the three wrap-around bent rows totals 260.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 15. Each of the five parallel bent rows totals 260.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 16. Each of the three wrap-around bent rows totals 260.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 17. Each of the five parallel bent columns totals 260.

52	61	4	13	20	29	36	45
14	3	62	51	46	35	30	19
53	60	5	12	21	28	37	44
11	6	59	54	43	38	27	22
55	58	7	10	23	26	39	42
9	8	57	56	41	40	25	24
50	63	2	15	18	31	34	47
16	1	64	49	48	33	32	17

Fig. 18. Each of the three parallel wrap-around columns totals 260.

(Figure 2). It is also the case that any half-row or half-column sums to half of 260, or 130 (Figure 3).

Another property is that any 2 by 2 block (including wrap-around) sums to 130 (Figures 4 & 5). I'm just showing you some of these blocks, but you can see that every single one of these has to sum to 130. Next, any set of four cells forming a rectangle with an even number of cells on each side sums to 130. This property is illustrated in Figures 6 & 7. Now a corollary to the above is that if you take the interior cells and then add the corner cells, the sum is 260 (Figure 8).

Franklin also defined the notion of bent rows and bent columns. In Figure 9, you see a bent row at the top and a bent row at the bottom. They each sum to 260. And the bent columns in Figure 10 also sum to 260.

In Figure 11, we have parallel bent columns. Each of these parallel bent columns sums to 260. It's pretty amazing what he did. Not only that, but with wrap-around bent columns, you also get 260 (Figure 12).

In Figure 13, we have parallel bent rows in which each sum is 260. In Figure 14, we see bent rows with wrap-around, again summing to 260. In Figures 15 and 16, we see more of the same.

In Figure 17, we see parallel bent columns summing to 260. And in Figure 18, we see each wrap-around bent column also summing to 260. If you think about it, this is amazing that he was able to construct these magic squares (and magic circles also) with so many interesting properties.

It turns out that there are many additional properties that I could show you, but we've already made the point. This is a very impressive Franklin square, and yet it pales in comparison to Franklin's masterpiece. His masterpiece is a 16 by 16 Franklin square. This Franklin square has 4 times as many numbers as the 8 by 8 square, and so it would be difficult for me to display it on a computer screen. But it's also the case that if I tried to show you the 16 by 16 Franklin square, I would need a large number of figures to show you all of its properties, because it has many more properties than the 8 by 8 square that we have been working with. If you want more details, see [2]. One of the interesting facts that I came across is that mathematicians today are trying to figure out exactly how he was able to construct these squares; after all, he was working with pencil and paper and somehow he was able to construct some very intricate magic squares.

Now, if we think about this problem from the vantage point of today, we can view this as a constraint satisfaction problem. We have a number of these summation constraints. We can think of variables as being defined by x_i which is the number assigned to cell i. We have constraints that tell us x_i is not equal to x_j if i is not equal to j. So, we can use an arbitrary linear objective function, and we can model this as an integer program in order to take advantage of the very powerful solvers that exist today like CPLEX, and solve them. So, I would argue that Franklin's work on magic squares is a first example of his interest in Operations Research (OR), even before OR developed as a field.

The second example of Franklin's work in OR involves his interest in decision making. He devised a fairly simple tool in 1730. The tool was based upon the notion of using a chart where you essentially have two columns (Table 1). One column has the heading *Pros*, and the other has the heading *Cons*. You place the

Table 1. Franklin's Decision-Making Tool was a precursor to cost-benefit analysis and AHP.

Pros	Cons
Factor A	Factor D
Factor B	Factor E
Factor C	Factor F
	Factor G

relevant factors under each one of these headings. And that was fine, but he took it one step further. Each factor has a weight, and the weights aren't necessarily the same. So, for example, if factor A was as weighty as factors D and E combined, he would cross out those three factors from this chart. If factor B was as weighty as factors F and G, he would cross out these three factors, as well. Given that there is a factor (C) remaining in the left column, the argument in favor of the proposal is the winning argument.

Now it's a fairly simple idea, but you can think of it as a precursor to cost-benefit analysis and the Analytic Hierarchy Process (AHP) in which you have factors and sub-factors (see Saaty [3]). The factors and sub-factors have weights. Instead of representing them visually as a chart with two columns, we use a tree-like or hierarchical structure in AHP.

The third example, which is our last example, comes from a paper that Ben Franklin wrote entitled, "Observations Concerning the Increase of Mankind." He wrote this in 1751. It was based upon his analysis of existing data. This was a data-driven model. He looked at the data, and he observed that the Colonists were more likely to marry than the English, they married younger, and they averaged twice as many children. He was able to quantify these observations, and then he went through some calculations, which we would regard today as back-of-the-envelope calculations. But he was able to conclude that America's population would exceed that of England in 100 years. And his forecast turned out to be exactly right.

Again, one of the things that amazed me about this is that the two preeminent economists of the time built upon Franklin's work. Adam Smith and Thomas Malthus were college professors and college educated, and they cited Ben Franklin's paper in their work. Adam Smith, who came up with the notion of the Invisible Hand, and Thomas Malthus, who was so pessimistic about the impact of population growth, were two giants in the field of economics. They devoted their lives to the field. As we mentioned, Ben Franklin went to school only up to the age of 10.

3 Conclusions

So, in conclusion, I think we can say that Ben Franklin was one of the few American operations researchers to pre-date Saul Gass. It's also curious that the interests of Ben Franklin that I've focused on coincide very nicely with Saul's interests. Without being aware of it, Ben Franklin was interested in constraint satisfaction, mathematical programming, decision-making, and data-driven modeling, as Saul has been over the last 50 years. Both were leaders in their fields. I, therefore, think it is appropriate to congratulate both Ben Franklin on his 300th birthday, and Saul Gass on his 80th birthday.

In closing, let me say that Saul and I have been colleagues now at the University of Maryland for 30 years, and it has been my privilege and my pleasure to have him as a colleague. Happy 80th Birthday and many, many more.

Acknowledgments

The author thanks his secretary, Ruth Zuba, and his Ph.D. student, Kok-Hua Loh, for their numerous contributions to this paper.

References

1. W. Isaacson. *Benjamin Franklin: An American Life*. Simon & Schuster, New York, 2003.
2. P.C. Pasles. The Lost Squares of Dr. Franklin: Ben Franklin's Missing Squares and the Secret of the Magic Circle. *The American Mathematical Monthly*, 108 (6): 489-511, 2001.
3. T.L. Saaty. *The Analytic Hierarchy Process*. McGraw-Hill, New York, 1980.

Good Management, the Missing XYZ Variables of OR Texts

Kenneth Chelst and Gang Wang

Department of Industrial and Manufacturing Engineering
Wayne State University
Detroit, MI 48202
kchelst@wayne.edu, wanggang@wayne.edu

Summary. Introductory operations research and management science textbooks survey a standard set of modeling techniques: linear programming, queueing theory, inventory control, decision trees, etc. The discussion throughout focuses on decision making and often explores the possibility of purchasing additional resources to improve performance. All of the modern texts provide a rich array of real-world examples of successful OR projects. By referring to recognized best practices, we argue that there is another factor: management's multiple roles that are missing from the problem context. We believe this gap is a factor in why other approaches to problem solving are often more visible and more marketable than OR. We provide examples of discussions of managerial oversight, leadership, and effort that could easily be added to chapters on mathematical programming and queueing, so as to place OR models in a broader context and increase their value.

Key words: Management; operations research; linear programming; queueing; inventory control; decision theory.

1 Managerial Effort vs. Decision Making

Executives and managers spend much more time managing processes and the results of a decision than actually making a decision. We do not argue that this is how it should be. There is ample evidence that too often inadequate time and energy is spent seriously assessing a wide range of alternatives when making decisions. Often only one alternative is considered, and the question reduces to yes or no, and then for how much.

Not all bad decisions can be overcome with a heroic investment of managerial effort. The press and literature are replete with data on the infrequent success of the decision to merge with or acquire another company. The failures can be attributed to the decision itself, as well as a lack of understanding as to the managerial challenges involved in implementation. Similarly, the data on

the decision to "re-engineer" a process or corporation are equally unnerving. The success of management of change is not much better [2].

Despite these cautionary words about the importance of good decisions, managing consumes the lion's share of people's time. It is our thesis that it is important for OR modeling efforts to be better aligned with how people spend their time and energy. It is our contention that textbooks in Operations research and Management Science miss out numerous opportunities to incorporate management thinking into their examples.

1.1 What Do Managers Do Besides Make Resource Allocation Decisions?

Managerial responsibility includes, but is not limited to, the following tasks, in addition to making resource allocation decisions.

Managers define and track performance measures: First and foremost, managers define and establish appropriate performance measures for their manufacturing or service system and set targets to be achieved and continuously improved upon. These measures may be internal as well as customer focused and must be communicated to the staff of the organization and in many cases to customers as well. Along the way, the manager may help his staff prioritize work including highlighting problem areas for immediate attention.

Managers establish and monitor policies: Managers establish operational policies that, for example, might directly influence call rates and service discipline in a queueing system. If strict adherence to a set of policies and procedures is critical to performance, the manager would regularly review reports on compliance and take action against serious deviations from policy.

Managers direct problem solving efforts: Through a process of constant feedback and monitoring, a manager should create an environment of accountability and rewards for the organization to meet and even exceed established targets. If the systems fall short, he may lead cross-functional and possibly cross-organizational teams to identify the root cause of the problem, help the team develop creative solutions, and monitor progress towards solving the problem. More specifically, managers challenge and ask tough questions of their staff and expect them to get back to them with solutions. The "experienced" manager can provide guidance as to where to look for solutions. If a crisis arises, he may lead the firefighting effort and temporarily reassign resources to the troubled area.

Managers lead waste reduction and other continuous improvement efforts: One broad area of managerial interest is what to do about waste. This issue has gained greater recognition because waste reduction, Muda, is a critical dimension of Toyota management philosophy. In this context, a manager may set an objective for his organization to reduce waste, establish a culture that constantly seeks ways to eliminate waste, and possibly lead creative problem-solving sessions that search for the root cause. A corollary to this concern over

waste is that managers do not like to waste their own time solving the same or similar problems over and over again. For example, if a schedule frequently changes, the manager and his staff will be constantly running around fighting new fires created at every schedule change. These production and logistics hiccups make it hard for the manager to focus on the big picture of meeting or surpassing key performance measures.

The above description in no ways covers all managerial responsibility and action. The tasks we have described, however, are relevant to the OR examples we discuss in the paper.

In "The Trouble With Optimal," Clauss [5] argued for a managerial perspective to be included in the presentation of OR models and the interpretation of results, especially for OR courses taught in business schools. He argued that:

1) a model-generated optimal solution may be bad policy,

2) managers can do better than the optimal, and

3) OR modeling chapters, such as inventory modeling, may ignore the biggest current issues and best practices to focus just on the models.

In this paper, we build and expand on his theme, but also argue that OR survey courses, including those in math and industrial engineering departments, should place the OR models in the proper management context. We use common categories of examples from survey texts to illustrate how management's multi-faceted role should be integrated into the modeling presentation and results discussion. Our examples are drawn from the following topics within these three main OR subjects.

1) Mathematical Programming: managing constraints, changing parameters, robust stable solutions and beyond optimality.

2) Queueing Theory: waiting experience, managing arrivals, managing service time and accountability.

3) Decision Tree and Inventory Models.

2 Mathematical Programming

Every introductory OR/MS text has several chapters about mathematical programming. One of the core powerful concepts of LP is duality theory. The values of the dual variables associated with the various constraints provide an estimate of the marginal value of the resources associated with the constraint. There are two almost universal observations. Constraints that are not binding and have zero shadow prices are of limited concern. Conversely, highly priced constraints indicate a need for more resources and the shadow price helps answer the question, "What is the maximum amount that a manager should be willing to pay for an additional unit of a resource that is tightly constrained?"

2.1 Attacking Non-Constraints

Clauss [5] attacks the idea that non-binding constraints are of little concern and argues that "optimal solutions may not be good management." Underutilized capacity is a symptom of a broader problem for the company and a good manager would expend significant energy to find a strategy to generate demand to utilize this capacity. Alternatively, he may strive to reduce this capacity, so that it no longer appears on the corporate books and so as to improve the corporate ROI (return on investment).

We built into the implementation of a vehicle prototype optimization model at Ford the recognition that experienced managers could improve on the optimal by addressing concerns regarding underutilization of expensive resources [4]. A vehicle prototype can cost as much as $500,000. In the model we developed, design and test engineers specify the tests they need to perform and the characteristics needed on the test vehicle. The model considers all of the tests and the timeline required for completion to design an entire fleet of vehicles that cover all of the test requirements while meeting deadlines.

The experienced prototype fleet manager focuses his review on those vehicles with relatively low utilization. They were typically required to perform only a few tests that were highly specific or tightly time constrained. The fleet manager sits down with the test requestor who was assigned one of these vehicles and enters a dialogue as to how his test requirement needs might be met otherwise. His questions explore various implicit and explicit constraints of the model.

1. Could part or the entire test be carried out on something less than a full prototype vehicle?
2. Have you over-defined the vehicle specifications? Could you use a more generic vehicle for the tests?
3. Could the test be performed simultaneously with other tests on a similar vehicle?
4. Is the timeline really fixed or can the test be delivered at a somewhat later date?

The immediate goal of this discussion is either to move the test to another prototype with spare capacity or use something other than a full-scale prototype. If he succeeds in moving all of the tests assigned to the underutilized vehicle, the fleet manager can reduce the prototype fleet by one vehicle. It is easy to justify this dialogue dynamic as managerial time well spent when one vehicle costs $500,000.

2.2 Tackling Constraints

The OR text recommendation to buy more capacity for highly priced constraints is one of the last thoughts that will pass through a manager's mind, especially if it requires going through multiple layers of management to get approval for a capital outlay. Even a suggestion to shift capacity from one

highly priced constraint to a lower priced constraint or to a constraint that is non-binding is rarely a trivial task and can cost millions of dollars. In the automotive industry, shifting assembly capacity between distinct vehicle lines may cost hundreds of millions of dollars and would rarely be done, except during a new model changeover that already requires a major plant redesign and capacity investment. Good managers have other ways of using the insights provided by shadow prices. Goldratt's *The Goal* [7] addresses this very issue, and the Japanese have an entire manufacturing process management philosophy that focuses on addressing constraints.

Goldratt [7] developed the concept he boldly labeled the Theory of Constraints. His examples do not acknowledge linear programming nor the concepts of duality theory, but in practice he is looking closely at shadow prices within an LP type context. He recommends that managers apply significant attention to that constraint to make sure that the resource is never wasted. As he points out, few production processes are totally deterministic and any amount of variability in a sequence of operations could leave a constrained resource under-utilized. Goldratt identifies ways that constrained resources are typically wasted. The resource may sit idle while workers are on break or for some other reason such as machine breakdown. The resource may be wasted by producing defective parts or producing parts that are not in demand.

Managers in their role of oversight and team leaders might set new policies on breaks or challenge a team to reduce machine down time. In addition, by establishing appropriate performance measures, machines would not simply be kept running while producing parts that are defective or not needed.

GM came in first in the 2005 Edelman prize competition with work that significantly increased the throughput of its plants. The project leaders widely distributed Goldratt's book. The array of OR models they developed were used to identify the ever changing bottleneck constraints on throughput. This helped plant teams identify problem areas to work on, so as to reduce or eliminate these bottlenecks. Plans of action were then developed by "multidisciplinary teams of engineers, line operators, maintenance personnel, suppliers, managers and finance representatives. Next they committed to implement the solution by assigning responsibilities and obtain commitment of implementation resources." [1]

Japanese manufacturing management principles also include an emphasis on dealing with constraints. First and foremost, they push the principle of Kaizen, "the process of making incremental improvement, no matter how small, and achieving the lean goal of eliminating waste that adds cost without adding value"[11].

In the context of LP, this would be captured by working to reduce the left-hand side coefficients of highly priced constraints. In addition, the Japanese support a philosophy that sees constraints not necessarily as unbending givens, but as something that teams of creative people can ultimately control and break. For example, the long time required to change over stamping dies or even whole product lines constrained management's ability to shift produc-

tion, leaving highly priced resources underutilized when demand shifted. They broke the back of the die changeover constraint by designing equipment that took minutes to change instead of hours.

Toyota developed another constraint-breaking concept for their assembly plants. They redesigned their vehicles off of a common platform, so that now multiple models can be produced on the same line. This totally changes the left-hand side of any vehicle assembly line constraint, as it allows for more decision variables. Unfortunately for the US, it will take GM and Ford more than a decade to fully implement the same idea, as they incrementally introduce more and more new products off of a common build-to-process protocol.

Besides the above specific examples, we recommend that teachers broaden the discussion of constraints to encourage out-of-the-constraint-box thinking. Students at all ages can readily respond to the question, "What constraints do they face in their academic lives, work lives, and or personal lives?" Example responses might include: limited time to study, course schedule conflicts, and course prerequisites. The basic decision variables are generally allocation of limited time or money. However, you can explore with them creative strategies for dealing with these constraints. I am sure, like Clausen, constraints on money and time that are non-binding will be of special interest to them. They will also understand the need to not waste high value constrained resources.

2.3 Stable Solutions and Schedule Robustness

OR textbooks often use a production scheduling problem to introduce a time dimension to the math programming problem. The optimal schedule specifies the number of units to be produced during each time period in response to significant seasonal variations in demand. In one textbook example, the optimal schedule involved 34 units in March and 66 units in April [15]. No comment was made as to the managerial ramification of such wide swings in production. An OR instructor might ask his students what it would be like for them, their teachers, and classroom utilization if class schedules changed every week instead of once every four months. How long does it take them each semester to develop an efficient routine around their new schedule? The instructor could also explore with the students the concept of adding constraints that limit total production variability from one time period to the next to a maximum of X percent.

Toyota, which is often presented as the benchmark for production planning in the automobile industry, lists as one of its core management philosophies Heijunka, which means to level out the workload across time periods. Without workload balance, they argue that it is not possible to develop a pull system, to drive the manufacturing schedule, nor is it possible to standardize work so as to balance an assembly line [11]. In addition, unbalanced production schedules at final assembly can wreak havoc in the supply chain that supports this production schedule.

Another core management principle is Kaizen. "Continuous improvement (Kaizen) can occur only after a process is stable and standardized. When you make processes stable and have a process to make waste and inefficiencies publicly visible, you have an opportunity to learn continually from your improvements" [11]. A production schedule as presented in the OR text that jumps by 100% from one month to the next would make it harder if not impossible to monitor efforts at continuous improvement in productivity.

OR texts generally ignore another core managerial production management concern, that the schedules be robust. The fundamental basis for all of linear programming is that the optimal solution always lies at a corner. This is its strength but also, we argue, a potential managerial weakness. A small change in a coefficient could send the optimal off to a neighboring corner point with significantly different values for the decision variables, while adding one new decision variable and dropping out another. This solution tendency undermines a primary concern of production managers, the robustness and stability of the schedule.

The concerns over schedule robustness become even more critical in environments which require "concurrent setups and hot staging" or when "careful sequencing" is required. This criticism is not really limited to the application of LP to production as presented in OR textbooks. McKay and Pinedo [14] point out that production scheduling research has had relatively little impact on manufacturing practice due to the lack of scheduling robustness. They categorized measures of the robustness of a schedule along two dimensions, performance robustness and execution robustness. "Performance robustness is the degree to which a schedule can absorb shocks without significant degradation in schedule quality, whereas execution robustness is the degree to which a schedule can absorb shocks without the need for rescheduling."

In light of the above, we are not arguing that all production examples should be taken out of OR texts. They do illustrate how a time dimension can be incorporated into math programming models. We are, however, saying that theses examples should be accompanied with a modest introduction as to the practical issues and concerns of production planners. In addition, the optimal solution should be discussed in terms of challenges it might create for managing the production environment.

3 Queueing Theory: Waiting Experience, Managing Arrivals, Service, and Accountability

Queueing theory models represent the earliest success stories of operations research and are incorporated into all basic texts. After defining the key measures, the texts proceed to introduce the basic single server model and a multi-server model. The single server model discussion typically includes both exponential service time and general service time formula. This is used to highlight the impact on waiting time of service time variability. The text

may include an analysis of the value of using a more expensive but faster service mechanism. In response to long waiting times, the texts explore the value of adding servers, i.e., increasing service capacity. They may also demonstrate how a two-server queueing system with one waiting line outperforms two separate lines, one before each of the servers. This observation is correct only if customers cannot easily move back and forth between queues.

Before jumping into the mathematics of queueing theory, it is important to ground the concepts in the reality of waiting in line. This is especially critical since the primary measure of service, waiting time, is non-commensurate with the key decision variable, service capacity, and its associated cost. In many situations it is difficult to place a dollar value on waiting. In addition, the cost of waiting is incurred by the customer, and the cost of service is assumed by the service provider. As a result, a manager will explore every option available to improve service before spending his organization's money on added capacity.

Managers in their role of establishing performance measures for a queueing system besides just mean waiting time should keep in mind David Maister's [13] "the first law of service" concept:

$$\text{Satisfaction} = \text{Perception - Expectation.}$$

It is possible to manage the perception or expectation or both to increase customer's satisfaction. Maister [13] proposes eight factors that influence a customer's satisfaction with waiting time:

1. Unoccupied time feels longer than occupied time.
2. Preprocess waits feel longer than in-process waits.
3. Anxiety makes waits seem longer.
4. Uncertain waits are longer than known, finite waits.
5. Unexplained waits are longer than explained waits.
6. Unfair waits are longer than equitable waits.
7. The more valuable the service, the longer people will wait.
8. Solo waiting feels longer than group waiting.

Item 4 on this list is the most linked to the mathematics of queueing as it will enable a manager to forecast the estimated waiting time and provide that information to a customer at the point of joining the line. Disney theme parks were among the earliest organizations to make a priority of managing customer satisfaction with the waiting experience. One of their first initiatives was to post signs along the queue that state if you are here, this is your expected waiting.

Besides improving the waiting experience, there are other actions that managers can take that build on the mathematics of queueing but do not significantly increase the service or facility costs. Every queueing model includes the arrival rate and the service rate, both the average and distribution. It would be useful if the texts provided illustrations of managerial actions and their impact on waiting time. For example, more than 90% of the burglar alarm calls that come into 911 are false alarms. Managers in their role of policy makers have imposed fines for the second or third false alarm which in

turn have dramatically reduced the rate of false alarms and, thereby, reduced queueing delays.

Perhaps the most significant change in routine problem solving in the last ten years is a whole array of strategies companies have employed to answer Frequently Asked Questions (FAQ), so as to reduce the need for extended direct person-to-person communication that often produced long waits on hold. The Internet has increased the importance of managing FAQs and reducing calls to operators.

Managers can also reduce the variability in the arrival rate. For example, museums that arrange for a special exhibit often sell tickets that contain a specific time slot for the visitor to arrive.

Managerial oversight and continuous improvement efforts can also decrease the average service rate. This might entail creating better standards for processing customers. A special service agent might be available to handle difficult cases that are likely to take a long time so as not to delay the vast majority of routine customer transactions. A number of medical emergency rooms have established and successfully achieved a goal of seeing patients within thirty minutes. They have often used the principles of Lean and Six Sigma to root out sources of delays without adding capacity. (Lean refers to a systematic process of removing non-value added activities and applies to both manufacturing and service delivery. Six Sigma is a structured methodology for reducing process variability.)

In the context of continuous improvement, the recommendation of forming a single queue may, in fact, be counterproductive. Imagine if all customers arriving at a bank were divided approximately equally between three regular tellers based on their last name into three separate lines. The standard queueing analysis presented in introductory texts indicates this is inefficient. However, if through this dichotomy, a teller gets to know regular customers, he may be able, over time, to reduce their average processing time sufficiently so as to reduce the customer's overall time spent in the bank. More importantly, the friendly knowledgeable teller may be able to deliver a higher level of customer satisfaction and sell more products and services through increased familiarity with the customer and his personal or business needs. This last point highlights the importance of managers establishing the right performance measures.

The merging of lines for efficiency might also go counter to Toyota management's principle of using visual controls so no problems are hidden [11]. In their environment individual accountability and initiative is critical to continuous improvement. A separate line before each teller provides a more visible sense of performance ownership and accountability than a single long line shared by all. Consider the following example. Assume that server utilization is 95%. With two separate lines, the average number of customers (in steady state, assuming exponential interarrival and service times for simplicity) will be two times nineteen for a total of thirty eight customers in the system. A single queue reduces that number dramatically to 19.5. Consider,

however, that if greater visibility and accountability associated with separate lines can reduce the service time slightly so that average utilization drops to 90%, the combined queue is only 18. A further decrease to 85% utilization reduces the average to 11.4. This example highlights the importance of the managerial issues associated with operations that should be considered before using queueing theory to justify something that at first glance seems to be a no brainer improvement in performance.

4 Decision Tree and Inventory Models

We have provided detailed examples of the importance of bringing a managerial element into the application and interpretation of OR models in mathematical programming and queueing theory. Analogous arguments can be presented in the context of decision tree and inventory models. Chelst [3] explored in "Can't See the Forest Because of the Decision Trees" the need to change the way introductory OR texts present decision trees. One point that was made is that managers do not consider the "states of nature" as immutable. Often the first example is an oil drilling problem with uncertain states of nature referring to oil in the ground. The real state of nature is "recoverable oil," and the state changes depending upon how much effort a company is willing to put in as result of increases in the price of oil.

With regard to inventory modeling, the single biggest deficiency in the introductory remarks is the lack of acknowledgment and discussion of the fundamental changes in inventory management that have occurred as a result of the information explosion as well as other managerial efforts. WalMart grew to be the world's largest retailer, in part because of its leadership in the deployment of IT to manage its inventory. Dell controls its inventory by better managing its supply chain, from customer order to delivery. The Japanese leadership in Just-in-Time inventory was not driven by the inventory models in standard OR texts. Vender Managed Inventory (VMI) improves performance as a result of the tacit mutual information sharing between retailers and manufacturer, which can eliminate the bullwhip effect, and improve the flexibility of production and delivery planning.

We recognize that OR texts are designed to teach inventory models and not the theory of inventory management. However, we argue that it is important for the student to understand the context in which these models are used and the relative value of OR models as compared to other managerial initiatives. Without that perspective, they may push complex solutions of modest value when compared to the other alternatives available to management for reducing inventory costs.

5 What's a Teacher To Do?

We have explored a number of practical managerial concepts and best practices and related them to different aspects of OR model application and interpretation. The OR and Management Science texts in vogue are unlikely to change quickly just because of these observations. Thus the teacher of OR has to take initiative to complement his standard text description. For OR professors with limited practical managerial experience, we urge they stay at least superficially abreast of the latest management buzzwords such as Lean Manufacturing, Six Sigma, and Process Reengineering. When reading a book such as the *Goal* or *The Toyota Way*, they should continually ask themselves what are the ramifications of these concepts for OR modeling and interpretation. They should pay special attention to any case studies provided, so that they can enhance their classroom discussions with real organizational experiences. It is critical for the teachers of OR to understand that management's r! ole in the OR context is not just to sponsor an OR project, but also to help validate a model and facilitate its implementation. Little argued that the role of OR models is to "update intuition." Every aspect of the model formulation, model results, and related communication contributes to this intuition building and subsequent decisions and managerial actions.

References

1. J.M. Alden, L. D. Burns, T. Costy, R. D. Hutton, C. A. Jackson, D. S. Kim, K. A. Kohls, J. H. Owen, M. A. Turnquist and D. J. Vander Veen. General Motors increases its production throughput. *Interfaces*, 36 (1):6–25, 2006.
2. M. Beer and N. Nohia. Cracking the code of change. *Harvard Business Review*, 78:133–141, 2000.
3. K. Chelst. Can't see the forest because of the decision trees: a critique of decision analysis in survey texts. *Interfaces*, 28 (2):80–98, 1998.
4. K. Chelst, J. Sidelko, A. Przebienda, J. Lockledge, and D. Mihailidis. Rightsizing and management of prototype vehicle testing at Ford Motor Company. *Interfaces*, 31(1):91–107, 2001.
5. F. J. Clauss. The trouble with optimal. *ORMS Today*, 24(1), 1997.
6. D. A. Garvin and M. A. Roberto. What you don't know about making decisions. *Harvard Business Review*, 79(8):108–116, 2001.
7. E. M. Goldratt and J. Cox. *The Goal: A Process of Ongoing Improvement.* North River Press, Inc, Great Barrington, MA, 1985.
8. F. S. Hillier, M. S. Hillier, and G. J. Lieberman. *Introduction to Management Science.* McGraw-Hill, New York, NY, 1999.
9. F. S. Hillier and G. J. Lieberman. *Introduction to Operations Research.* McGraw-Hill, New York, NY, 2004.
10. K. L. Katz, B. M. Larson, and R. C. Larson. Prescription for the waiting-in-line blues: entertain, enlighten, and engage. *Sloan Management Review*, 3(2):44–53, 1991.
11. J. K. Liker. *The Toyota Way.* McGraw-Hill, New York, NY, 2004.

12. J. D. C. Little. Models and managers: the concept of a decision calculus. *Management Science*, 50(12):1841–1853, 2004.
13. D.H. Maister. The psychology of waiting in lines. *Harvard Business School Note 9-684-064*, Boston, 1984.
14. K. McKay, M. Pinedo, and S. Webster. Practice-focused research issues for scheduling systems. *Production and Operations Management*, 11(2):249–258, 2002.
15. J. H. Moore and L. R. Weatherford. *Decision Modeling with Microsoft Excel.* Prentice Hall, Upper Saddle River, NJ, 2001.
16. H. A. Taha. *Operations Research: An Introduction.* Prentice Hall, Upper Saddle River, NJ, 1997.
17. B. W. Taylor III. *Introduction to Management Science.* Prentice Hall, Upper Saddle River, NJ, 2004.
18. W. L. Winston. *Operations Research: Applications and Algorithms.* Duxbury Press, Stamford, CT, 2003.

The Operations Research Profession: Westward, Look, the Land is Bright

Randall S. Robinson

12 Sparks Station Road
Sparks, MD 21152
randyrobi@cs.com

Summary. As the worthy profession of operations research enters the twenty-first century, the outlook is bright in many respects. Yet various impediments to progress deserve attention. In this essay I offer a personal survey of the situation, focused on the delivery of operations-research services in practice. I begin by reviewing the nature and scope of modern operations research. Then I examine the current status of the profession – what's going well and what isn't. Finally, I propose actions that I believe promise to improve the opportunities for operations-research professionals and their clients to undertake beneficial projects.

Key words: Operations research profession; marketing operations research; operations research practice.

1 Introduction

To work in the field of operations research (OR) can be exhilarating. In 1995, the marketing director of the Institute for Operations Research and the Management Sciences (INFORMS) made phone calls to randomly selected INFORMS members to see how they were feeling about things in general. Some members indicated such a great affection for the profession that they became teary talking about it.

In many respects this outstanding profession is in excellent shape. Yet few OR professionals I know would say we are without serious challenges.

Winston Churchill, in a wartime speech on April 27, 1941, encouraged his country's citizens with a poem by A. H. Clough [1] that ends with the now-famous line: "But westward, look, the land is bright!" This might apply today to the OR community.

The idea in what follows is to examine the professional landscape with emphasis on practice, note the good and the not so good, and suggest how the profession might want to address specific difficulties.

The thoughts offered are my own, and therefore not in any sense officially endorsed.

2 The Nature and Scope of Modern Operations Research

2.1 Overview

What generic activities are most central to the practice of OR? I think it's clear the answer is analysis and related consulting. The client usually, but not always, is management. And the goal ordinarily is to help management improve, in some way management wishes, their organizational effectiveness – not always the same as "efficiency."

Seemingly similar basic services of management-application analysis and consulting have been provided, and will in the future be provided, by the thousands of staff analysts and management consultants who are not in the OR community, and not qualified to be.

What sets OR apart is that its practitioners ideally are able to draw upon pertinent analytical methods and knowledge throughout technology – mathematics, engineering, and science.

The situation is analogous to the difference between a medical doctor and a less-well-trained medical professional (nurse, paramedic, medical technician). Ideally, the medical doctor is able to determine, among all available treatments, sometimes with assistance from specialists, the treatment that best fits the patient's need. The non-physician lacks the requisite comprehensive knowledge and skill to permit offering the highest quality medical care.

An OR practitioner is not limited to employing advanced methods, just as the physician is not limited to performing brain surgery. Simple methods may suffice in a particular case. But the ideal OR professional is uniquely able to skillfully prescribe and apply advanced methods where appropriate, while the non-OR counterpart is not.

2.2 The Complete Practice Agenda

The foregoing omits various details. One worth noting is that the ideal OR practitioner is skilled in all steps from beginning to end of a project, not just skilled in analysis. Those steps include diagnosing a client's needs or opportunities, working with a project team that contains client-organization members, and assisting with or guiding implementation.

Often the OR task is to develop a computer-based system, called a "decision-support system" in some instances. Following the custom of application-software developers within the information-technology (IT) world, the inputs and outputs of such a system should be specified by the client or "end user." This standard IT approach worked well enough during the early days of replacing manual systems with functionally similar computer systems. But now the core of a managerial system designed to provide cutting-edge capabilities frequently cannot be specified in detail by clients, or by other domain experts who know the application in its current state.

Enter OR. It is the field that specializes in helping executives turn a general, even vague managerial-improvement idea into concrete plans and system

specifications. That is, OR specializes in determining what raw data is worth collecting, how to convert data into useful information, how to interpret and present the information for management consumption, and how to implement all of this in a system.

To borrow a few clichés, when management wants leading-edge decision support, or a best-of-breed smart system (sometimes better than best-of-breed), "OR rules." Moreover, OR professionals are the highest-tech management analysts and highest-tech management consultants on the planet.

In contemporary OR practice, application activities have expanded beyond former boundaries. Notably, while the clients most often are management, sometimes the clients may be non-management, such as front-line engineers or consumers. And while the OR end-result most often is analysis and consulting to help improve a managerial function, sometimes the result is a new product or new service. And, of course, system development frequently is central.

2.3 The Complete Profession

The purpose of OR is to deliver beneficial applications in practice. This is carried out by full-time practitioners and also by university faculty members who pursue practice part-time. Moreover, some professionals working in other disciplines, but who are qualified to do OR, work on OR projects occasionally.

The profession importantly contains those university faculty members who, whether or not they participate in practice projects, create new methods and knowledge for use in OR practice. Faculty members in addition perform the vital role of teaching the next generation of OR professionals and the next generation of prospective customers for OR services.

All of these contributors together make up the OR profession.

2.4 Is OR Science?

Is operations research a science? Yes and no.

"No," because it more closely resembles a branch of engineering than a branch of science (pure science). Morse and Kimball, in their famous early textbook ([10], p.1), say that "engineering" suggests "the construction or production of equipment," so they decided instead to call OR an "applied science." These days, we are more comfortable with the concept of an engineering product being information rather than hardware.

We may also say "yes," OR is science, or at least is scientific, because OR and all branches of engineering are "applied science" in that they draw upon formal knowledge and upon analytical methods created by members of the scientific community – scientists, mathematicians, and academic engineers who do research. Medicine is scientific in the same sense, except the things they draw upon come from different sectors of the scientific community.

2.5 Is OR Multidisciplinary?

Our heritage from the early days contained mixed signals about whether OR should or should not be multidisciplinary.

For example, the definition of OR in Morse and Kimball ([10], p.1) stipulates that "Operations research is a scientific method of providing executive departments with a quantitative basis for decisions regarding the operations under their control."[1] The implication of "quantitative" is that if it's not mathematical, it's not OR.

Yet, other voices from our heritage advocate drawing upon all helpful scientific methods, both mathematical and non-mathematical. Indeed Morse himself later wrote ([9], p.152) that an OR group should contain, in addition to the customary mathematically inclined members, "economists... and, in some cases, psychologists or other types of social scientists."

A modern view, in my opinion, is that the ideal OR practitioner will diagnose the client's need and then select appropriate methods and knowledge from the full array of analytical methods – both quantitative and qualitative – and from the full array of associated formal knowledge. In other words, OR should be multidisciplinary. Still, just as quack medicine is not acceptable in the medical community, methods and knowledge applied in OR (from the disciplines included in "multidisciplinary") should be scientific, in that they emerged from the scientific system of peer review.[2]

3 Status of the OR Field

3.1 The Good News

Applications of operations research are far ranging and diverse. You will find actual or prospective applications in all types of organizations – business, government, military, health care, non-profit; at all levels – top, middle, front line; and in most or all organizational functions. While only the tip of the application iceberg has been documented, by now the documented record of success stories is substantial. To cite just one example, the Franz Edelman Award competition finalist papers were archived since 1982 in the INFORMS journal *Interfaces*, and the associated competition presentations were made available on video or DVD.

Meanwhile, the environment in which organizations operate has become more likely to stimulate demand for assistance from OR. Indeed, there are strong pressures to improve performance beyond what has been achieved without OR.

---

[1] Gass and Assad [6] note that this famous definition originated with Charles Kittel, whose proposed version was modified by Charles Goodeve into the form put forward by Morse and Kimball.

[2] The view of OR portrayed in all of the above is close to that offered earlier (Robinson [12], [13], [14]). This is a broad view that tries to recognize current activities and trends in practice.

Those well-known pressures come from such sources as budget squeezes, the rapid introduction and phasing out of new technology, and global competition.

An important phenomenon that has strengthened the prospects for OR's future is the development of more and better pertinent analytical methods, along with the development of greatly improved delivery systems – computers, software, communications. Also, data that might serve as input to OR applications are more readily available and more easily computerized.

Accompanying this has been the increase, slowly but surely, of highly valued and widely used OR applications – e.g., revenue management ("dynamic pricing"), simulation of manufacturing operations, quantitative analysis of supply chains, data mining, valuation of investments, production scheduling and other kinds of scheduling, forecasting, and performance measurement.[3]

Reflecting on the points just noted above, one would judge that OR is positioned for fast growth in the immediate future. This would be true, I think, if it weren't for the obstacles I'll next review. I will say more about the obstacles than I did about the good news, because I want to focus later on overcoming obstacles.

3.2 The Problem of Weak Perceived Value

A fundamental obstacle confronting OR practitioners – possibly the single greatest obstacle – is that many prospective clients haven't accepted the proposition that expertise in advanced analytical methods and knowledge is sufficiently valuable to warrant seeking OR services. This contrasts sharply with medicine, where few patients question the value of professional assistance from physicians.

In other words, prospective customers for OR services may either not know of OR or else know of it but not see much value.

3.3 The Problem of Competition

A related roadblock has been that OR is sometimes viewed as direct competition by the non-OR providers of managerial analysis and consulting, who are well established, abundant in numbers, and sometimes pleased when OR is absent.

Fortunately, this attitude does not characterize every case or most cases; often enough the resident analysts are happy to team up with the OR professionals. It happens with sufficient frequency, however, to be significant. And when it happens, the consequence may well be to reinforce the impression that OR will not "add value."

3.4 The Problem of Poor Branding

The fundamental difficulty – lack of visibility and lack of appreciation even if visible – is referred to in the marketing world as poor "branding."

For operations research, various things besides competition contribute to poor branding. One is the extreme diversity of applications, making commonality

[3] A similar list of reasons to be "bullish on OR" were enumerated by Cook [2].

unclear. Another is that practitioners too often do not consider themselves to be part of the umbrella profession or that, even if they do, they fail to mention the profession to their clients. An analogy would be a medical profession with its diverse specialties but no overall field of medicine, no general title "doctor" or "physician," and no white coats.

In order to brand, the place you begin, I think you will hear from most or all marketing specialists, is with a brand name. Here is a colossal weakness in OR.

OR is in fact undertaken under a blizzard of different names. They include synonyms and near-synonyms (e.g., management science, decision technology, operations analysis, analytics), names of OR specialties in particular applications (e.g., financial engineering, marketing engineering, operations management), and the main professional fields of individuals who occasionally undertake OR projects (various branches of engineering, applied mathematics, and science). While the medical profession guarantees that all their practitioners regardless of specialty are recognized to be doctors, within the umbrella medical profession, OR lacks assured identification of practitioners with its umbrella profession.

This embarrassment of riches in names makes it difficult to take a census of practitioners. Because the job title sought in a census usually is operations research (or "operational research" in many countries) and possibly one or two synonyms, a count typically understates the true total.

3.5 The Problem of Adopting a Brand Name

With so many different names in use, and the possibility of coining new ones, adopting a single brand name for the OR profession is not a simple matter. More about this later.

3.6 The Problem of Ambiguous Scope

Branding of OR is rendered still more difficult than it already is because of different views concerning the boundaries of the field. The discussion in Section 2 gives a suggested rough guide, but it leaves ambiguities.

While documented OR applications address various managerial functions, the function most often mentioned when characterizing the OR field is decision-making. A particularly tricky related question is this: Does OR offer help with managerial decisions only, or does it offer help with all decisions?

The traditional scope of OR was to help executive management with their managerial decisions. A modern view is that OR also may assist the non-managerial client, perhaps an engineer. An extension of the traditional scope, then, would be that OR assists the engineer with matters ordinarily decided by management. Example: Applying advanced analytical methods to help decide about production layout or plant expansion (OR), about design of machinery or

robots (not OR). For those who wish to consider OR's scope to encompass all decision-making, however, everything in the foregoing example is OR.[4]

Such a lack of clear boundaries creates a problem for branding when it interferes with OR receiving visibility and credit for its professional work. If only acknowledged OR professionals did OR projects, we might propose that whatever they work on is OR. But the blizzard-of-names situation prevents that approach from solving the problem. If we are to brand the field not only with clients but also with professionals who presently use non-OR names when they practice OR, this issue further clouds the branding message.

I think this particular difficulty, while noteworthy, ranks low on the hierarchy of obstacles to progress.

3.7 The Problem of Overlap Between OR and Other Fields

Here's yet another source that gives rise to trouble with the umbrella branding of OR. Some hold the view that a particular project comes from either OR or from another field, but it cannot come from both. Example: An industrial engineer does an OR project, labels it industrial engineering, and feels that it can't be OR, too.

This view overlooks the numerous overlaps in technological disciplines. For instance, many branches of engineering draw heavily on physics. Medicine draws heavily on biology. Thus, it clearly is routine for work to be simultaneously from different disciplines (e.g., physics and electrical engineering). In the example, then, the work is simultaneously OR and industrial engineering. It certainly is not just industrial engineering.

I judge that much OR is being carried out by professionals who consider themselves to belong primarily to some other professional field, and therefore use the name of that field for their OR projects.

3.8 The Problem of Fragmentation

The problems of the name blizzard, ambiguous scope, and not recognizing discipline overlaps are aggravated by a fourth impediment. Too many professionals who in fact undertake OR work do not recognize that they are part of the umbrella OR profession, no matter what it's called or how you define the scope.

Such fragmentation of the OR field obviously is damaging. Where it occurs, successes in one sector do not help other sectors, clients who are served by only one sector miss opportunities in other sectors, and professionals in one sector do not learn from professionals in other sectors. Picture the dilemma for patients and physicians if the medical profession consisted of separate specialties that did not interact, refer from one to another, share knowledge, or otherwise combine efforts.

---

[4] Gass and Assad ([5], Preface) reflect this perspective when they say: "OR is the science of decision-making, the science of choice."

3.9 The Problem of Misunderstandings

When OR is visible, it too often is misunderstood. I'll note a few of the stereotypical misunderstandings that restrain increases in the demand for OR.

Misunderstanding 1: OR is pursued by impractical thinkers who are not likely to produce a truly useful result. The facts: A strong practical bent and sociability are survival skills in most OR-practice scenes. Experienced practitioners are well aware of this.

Misunderstanding 2: OR is defined to be work done with a specified few "OR methods" – linear programming (LP) is almost always mentioned. For instance, "if you're applying LP, it must be OR." Or, "if there's no math model, it can't be OR." The facts: Similar to a medical doctor, the OR practitioner starts by diagnosing a client's needs or opportunities, and then, ideally, selects the most appropriate method. Furthermore, the list of potentially helpful methods is quite long, is not restricted to mathematical modeling, and I suspect none of the methods are applied exclusively in OR.

Misunderstanding 3: OR works on a limited list of applications, associated with the name "operations," and generally at the lowest organizational level. The facts: OR applications, past and prospective, are found in just about all managerial functions, in all types of organizations, at all levels. Further, the application ideal is, and has been, the high-impact, strategic application (save the company, win the war) – just the opposite of low-level (optimize light-bulb replacement).

3.10 The Problem of Believing that Mathematical Methods are Too Crude to Be Helpful

Does OR's heavy use of mathematical modeling reflect an ivory-tower mentality that favors overly simple methods not capable of yielding helpful results in real-world applications?

Mathematical modeling is not something concentrated in OR. Indeed, I believe it's true that the fundamental analytical method shared in all of engineering and quantitative science is the mathematical model. In other words, mathematical modeling is ubiquitous, not rare and extreme. Moreover, these days high degrees of complexity are addressed successfully with mathematical models. One example whose outputs are familiar to the general public are the enormously complex weather-forecasting models.

The notion that managerial applications are so complicated that models are too simple to assist is, frankly, out of touch with the true situation today.

3.11 The Problem of Not Being Multidisciplinary

In practice over the years, I observed that practitioners were happy to apply whatever respectable method seemed to help. Most of them, I believe, were not limiting themselves to mathematical modeling out of an ideological belief that modeling offered the only worthwhile source of results.

Still, many or most practitioners were in fact limited in their repertoire of methods and knowledge of things quantitative. I think we can attribute this to the fact that academic degree programs in OR are really surveys of different mathematical methods, and so too are OR textbooks. The idea of being multidisciplinary – to include going beyond mathematical methods and quantitative knowledge to also cover qualitative methods and knowledge – hasn't substantially influenced courses and textbooks.

This situation continues. Meanwhile, the quantitative methods have grown considerably in their power to give assistance to management.

3.12 The Problem of Weak Interpersonal and Sales Skills

A continuing difficulty is that successful OR practice calls for interpersonal skills – the counterpart of a physician's bedside manner – not ordinarily cultivated in OR training. Whereas doctors are explicitly trained along these lines in their required residencies and internships, OR graduates typically enter practice with little or no such training.

A closely related difficulty is that also in short supply in OR practice are superior skills in conceiving worthy new projects and selling them to clients. Again, OR professionals rarely receive training in those key skills before they enter practice.

Clients who hire OR professionals sometimes fail to screen for those skills. It's a mistake that's easy to make, because clients naturally seek a strong technical background.

3.13 The Problem of Cutbacks in OR Departments and Jobs

I'm not aware of good data on the subject of cutbacks in OR. From my personal experience and conversation with many others, I think it's safe to say that cutbacks have been major. Corporate-level OR departments were common in large business firms; no longer. Required OR courses (or quantitative-methods courses) were common in business schools; no longer. And so on. Notable exceptions exist, of course. But the overall picture is one of retrenchment.

Naturally, cutbacks constitute a symptom, not a root cause.

4 Summarizing the Problems

The various difficulties just reviewed include the highlights, I believe. Others could be mentioned. For example, we don't know to what extent project failures have led to diminishing the reputation of OR in practice.

Summing up, while certainly not the whole story, the core problems appear to revolve primarily around under-marketing the profession and under-marketing individual OR projects. Significant difficulties also emanate from cases where some OR practitioners could have a stronger professional identity, greater technical breadth, and improved social skills.

5 Overcoming Major Obstacles to Further Progress

5.1 Goals

In my opinion, the ultimate goal of an initiative to improve the external standing of operations research should be to increase the global sum of benefits realized by organizations from successful OR practice projects. This would be a natural consequence of increasing the demand for OR services, i.e., "unlocking the latent demand for OR solutions" (Cook [3]). For OR professionals, both practitioners and academics, success in pursuing those goals – more OR-service demand leading to more organizational benefits – almost surely would result in enhanced jobs, pay, research funding, courses in the university curriculum, and reputations.

5.2 Topics

It's convenient to organize my suggestions for improvement under four headings: marketing the profession to customers; marketing an umbrella identity to OR professionals; promoting a modern, multidisciplinary view of the field; and increasing emphasis on marketing and social skills for OR practitioners.

I believe these are high-priority topics. They are not meant to cover all that might be said about removing obstacles.

5.3 Marketing the Profession to Customers

A fundamental difficulty is that a large number of our prospective customers or clients are either unaware of OR or have impressions of OR such that they do not call upon it for assistance. The apparent remedy is somehow to increase OR's visibility, explain the field better, and show enough value that clients are inclined to sign up for OR services.

The status quo is that individual members of the OR profession have been explaining OR and its value on their own, in their personal jobs. That the difficulties continue, with no indication they would be resolved by natural forces, leads us to see the need for collective action.

In short, we urgently need a campaign to market the OR profession. Fortunately, this need has been recognized and action has started.

A marketing-the-profession program was launched early in this decade by the Institute for Operations Research and the Management Sciences (INFORMS) of the United States. A close partner in this endeavor was the Operational Research Society (ORS) of the United Kingdom. The EURO unit of the International Federation of Operational Research Societies (IFORS) subsequently joined the effort.

In the campaign of INFORMS, a crucial move was reframing. Originally attention focused on the essential first step of adopting a brand name. In the U.S., that proved contentious. But when the program was explicitly reframed as marketing the profession, which was the goal all along, people began to feel more comfortable. Another point that increased comfort was that a brand name is for

marketing to the outside, so there was no push to rename things inside the profession or to rename INFORMS.[5]

The naming matter was investigated thoroughly by INFORMS with assistance from an outside marketing firm. The brand name "operations research" (or "operational research") was then selected, thus agreeing with a choice made independently by the ORS. The basic reasons were that it is by far the best established of the general names worldwide and that the estimated cost of promoting an alternative sufficiently to be adopted by the profession and its customers was much too high to consider. Further, the fact that "operations research" is not self-explanatory is acceptable, because a brand name does not require that property to be successful (e.g., Exxon).[6]

The approach taken by INFORMS was to travel two paths to prospective customers and those who influence them. First, a direct effort. Second, an effort of enlisting assistance from members of the profession, who reach out within their own work scenes.

Early outputs from the INFORMS program included a poster to rally members of the profession [7], a booklet to explain the field to prospective customers and others outside the profession [8], a website with resources for members of the profession (www.orchampions.org), and a website for prospective customers, the media, and others in our target external audiences (www.scienceofbetter.org).

The ORS has a website similar to the science-of-better site of INFORMS (www.theorsociety.com, click on "OPERATIONAL RESEARCH: SCIENCE OF BETTER"). The EURO unit of IFORS has a link to the ORS site (www.EURO-online.org, click on "OR Science of Better") and also a section to promote branding the profession (www.EURO-online.org, click on "Branding OR").

Continuing the INFORMS experience, an advertising campaign followed next. This didn't seem to produce immediate dramatic results, so then attention shifted elsewhere. The subsequent initiative was to let the internally prestigious Franz Edelman competition for achievement in operations research practice be more externally visible, in the spirit of the Baldrige Award for quality, and to combine this with reaching out to the media. At the time of this writing (spring 2006), that is the latest development.[7]

I think the multi-country marketing-the-profession program deserves the enthusiastic support of everyone in the OR profession. I hope the program will continue and strengthen.

---

[5] The marketing-the-profession program and the requirement to adopt a brand name for that purpose were described clearly by Cook [3].

[6] The INFORMS and ORS choices were explained well by Oliva [11] from the United States and Sharp [16] from the United Kingdom. Both are marketing professionals who have close-hand knowledge of the OR situation.

[7] The expanded Franz Edelman event includes a gala dinner and ceremony, the inauguration of an Edelman Academy for present and past finalist organizations, a formal program book with articles such as those by Cook [4] and Robinson [15], a keynote talk by a distinguished speaker, and the showing of professionally prepared videos.

5.4 Marketing an Umbrella Identity to OR Professionals

We turn now to the problem of fragmentation of the profession. To try to fix it, we would like to facilitate two things. First, professionals who perform OR work in practice should come to recognize that they are part of an umbrella profession and should take part in the profession's combined activities. Second, OR professionals who engage in practice should use the brand name operations research (or operational research), where they may do so without creating a problem for themselves.

Let me focus on the professional society INFORMS, based in the United States. The OR champions website (www.orchampions.org) offers motivation and resources for OR professionals who choose to support the marketing-the-profession program. But it doesn't achieve buy-in from those who don't care to take part in the program. And it misses especially those who practice OR, yet don't identify their work with the OR umbrella profession.

Of course the marketing-the-profession effort may in fact reach and influence a number of professionals. Still, the program is not really directed toward speaking to isolated segments of the OR profession. Therefore we should launch an additional effort. The method of reaching out would be, presumably, to identify each OR specialty, giving high priority to those where interaction with the umbrella profession seems weak, and to direct marketing communications there. This effort might complement the marketing of memberships in INFORMS, and could be part of the membership program.

5.5 Promoting a Modern, Multidisciplinary View of the OR Field

I've talked about bringing OR professionals together in their umbrella profession, and about marketing that profession to prospective clients and those who influence prospective clients. The next logical subject is to consider the quality of the work performed in OR practice projects.

Here I think we encounter an obvious important need. To increase the probability of a successful outcome, and, more than that, to offer the client the truly best service, an OR professional should have a broad view of OR, a view that includes OR being multidisciplinary.

The essence of a broad view is that OR is much more than just building a math model. OR is not a "tool," as you sometimes hear. OR is a process that extends from the earliest stage of a vague managerial-improvement idea or a first-conference with a client to the final stage of being up and running with a responsive result and having that result beneficially influence the intended managerial action – policy, plan, major decision, or other managerial output.

Furthermore, OR should be drawing upon, ideally, the most appropriate (the best) of pertinent methods and knowledge in all of technology – mathematics, engineering, and science. The single code word for this is "multidisciplinary."

One thing that would help with these matters would be a few university faculty members becoming excited about filling this need. We could use more team teaching. Our OR textbooks, which presently are mostly collections of

mathematical methods, could be expanded to cover the broader scope. The OR course programs could be restructured, especially for those students who intend to go into full-time practice.

What would be added to the traditional curriculum and texts? The first thing is consideration of the entire OR practice process. Second is surveying potentially useful methods and knowledge from all the disciplines. This may seem to be a tall order, yet even modest progress would be worthwhile.

Asking students to take overview courses in, for example, economics and social sciences (the disciplines mentioned by P. M. Morse, as I noted earlier) certainly is doable. A genuine further advance would be the addition of material that clarifies how methods and knowledge from such non-OR disciplines have been and could be brought to bear in OR projects.

5.6 Increasing Emphasis on Marketing and Social Skills for OR Practitioners

Practitioners who have a broad view of OR's scope and are multidisciplinary still could be ineffective. The prospects for success are greatly increased when practitioners also possess social skills, organizational savvy, and the ability to propose and sell applications. Not every member of an OR team requires the full set of such skills. But the skills should be present in the team.

Right now, marketing and social skills are addressed primarily by considering them (possibly) when screening for hiring, and then by cultivating them on the job. I suggest we could do more, for instance by mentioning them in modernized OR curricula and texts, and by monitoring them in cooperative study-work programs.

The point is these skills are crucial to OR success in practice. Whatever we could do to cultivate them would help the profession.

6 Conclusion

OR should be the answer to many managerial prayers for a promising new way to boost organizational performance. OR professionals are the folks who can bring you smart systems, cutting-edge decision support, and state-of-the-art managerial processes. OR is a true twenty-first-century technology. When it comes to management analysis and management consulting, OR professionals are the highest-tech experts on the planet.

But OR's ability to assist organizations is restrained by some troublesome impediments:

- Too many prospective executive clients have never heard of OR, or else they misunderstand it. In either case, they are not inclined to ask for OR services.

- Many qualified professionals who practice OR contribute to poor client perceptions by failing to operate in the way they should within an umbrella profession. They don't refer work outside their specialty, learn from others

outside their specialty, use the brand name of the umbrella profession, or take part in any collective activities on behalf of the whole profession.

- Some professionals see OR from a narrow, outmoded perspective, giving rise to such things as OR is a "tool," and the entire OR job is to build a math model. This spills over into the perceptions of clients. Modern OR, as most successful practitioners know well, is a process that takes you from the beginning of an idea or first client interview all the way through to implementation and, most important, to the organization's realizing benefits. Furthermore, modern OR would be more powerful if it revisited its historical roots by becoming once again avowedly multidisciplinary.

- To succeed, OR professionals in practice need more than technical skills. They need interpersonal skills, a sense of organizational "politics," and marketing skills – especially skills in marketing new projects. Not every OR professional has to be good in all things. But an OR team needs these skills at least among some members.

I have offered a few suggestions regarding how to surmount those obstacles, so that OR can flourish, which would be good news for organizations worldwide, where many opportunities to improve managerial functions are missed every day. We should assertively market the OR profession, and have started to do so. We should reach out to OR professionals who aren't participating in the umbrella profession. We should broaden and otherwise modernize curricula and textbooks. And we should pay attention as best we can to cultivating social and marketing skills, particularly among students headed for full-time practice.

Acknowledgments

I thank the many friends who contributed over the years to concepts set forth above, in particular to Saul Gass, in whose honor the book containing this paper has been prepared. For helping me to better appreciate the great potential and broad scope of the OR profession, the key role of a professional society, the profession's untapped beneficial impacts and associated high opportunity cost, its specific difficulties, the resulting call to action, and strategies for action, I am indebted especially to these additional colleagues (in alphabetical order): Al Blumstein, Tom Cook, Art Geoffrion, Carl Harris, Gordon Kaufman, Dick Larson, Gary Lilien, John Little, Irv Lustig, Ralph Oliva, Lew Pringle, and Graham Sharp. And I am most grateful to numerous other colleagues not named here. I also thank Maurice Kirby, with whom I exchanged the rough notes from which this paper grew.

References

1. A.H. Clough. Say Not the Struggle Naught Availeth. *The Oxford Book of English Verse,* A.T. Quiller-Couch, ed., Clarendon , Oxford, UK, 1919.
2. T.M. Cook. Bullish on OR. *OR/MS Today*, 30(1): 6, 2003.
3. T.M. Cook. The Branding Dilemma. *OR/MS Today*, 30(5): 6, 2003.
4. T.M. Cook. Running a 21st-Century Organization Requires OR. *The Franz Edelman Award Program Book*, 5-6, The Institute for Operations Research and the Management Sciences, Hanover, MD, 2006.
5. S.I. Gass and A.A. Assad. *An Annotated Timeline of Operations Research: An Informal History.* Kluwer Academic Publishers, New York, NY, 2005.
6. S.I. Gass and A.A. Assad. Model World: Tales from the Timeline – The Definition of OR and the Origins of Monte Carlo Simulation. *Interfaces*, 35(5): 429-435, 2005.
7. INFORMS. *Promote Operations Research: The Science of Better* (poster). The Institute for Operations Research and the Management Sciences, Hanover, MD, 2004.
8. INFORMS. *Seat-of-the-Pants-Less: Executive Guide to Operations Research.* The Institute for Operations Research and the Management Sciences, Hanover, MD, 2004.
9. Operations Research Center, MIT. *Notes on Operations Research.* The Technology Press of MIT, Cambridge, MA, 1959.
10. P.M. Morse and G.E. Kimball. *Methods of Operations Research.* The Technology Press of MIT, Cambridge, MA, 1951. Republished with Introduction by S. I. Gass, Dover Publications, Mineola, NY, 2003.
11. R.A. Oliva. Time to Move Forward. *OR/MS Today*, 31(2): 24, 2004.
12. R.S. Robinson. Welcome to OR Territory. *OR/MS Today*, 26(4): 40-43, 1999.
13. R.S. Robinson. A Business Executive's Guide to Modern OR. *OR/MS Today*, 27(3): 22-27, 2000.
14. R.S. Robinson. More Profit, Productivity, and Cost Reduction – From Operations Research. White paper available at www.scienceofbetter.org, The Institute for Operations Research and the Management Sciences, Hanover, MD, 2004.
15. R. S. Robinson. OR: Innovation as a Profession. *The Franz Edelman Award Program Book*, 8-9, The Institute for Operations Research and the Management Sciences, Hanover, MD, 2006.
16. G. Sharp. What's in a Name? *OR/MS Today*, 31 (2): 25, 2004.

Part II

Optimization & Heuristic Search

Choosing a Combinatorial Auction Design: An Illustrated Example

Karla Hoffman

Systems Engineering and Operations Research Department
George Mason University
Fairfax, VA 22030
khoffman@gmu.edu

Summary. This paper summarizes a talk given in honor of Saul Gass' 80th Birthday celebration. The paper is modeled after Saul's well-known book, *An Illustrated Guide to Linear Programming,* and presents some of the illustrations provided during that talk. In this paper, we explain why specific rules might be chosen within a general combinatorial auction framework. The purpose of such rules is to assure that the market mechanism is fair to both buyers and sellers, and so that the auction will end in an efficient outcome, i.e., the goods are won by those that value them the most. The paper describes some of the issues, both computational and economic, that one faces when designing such auctions.

Key words: Combinatorial auctions; winner determination problem; pricing; efficiency; threshold problem; exposure problem.



1 Introduction

We begin this paper by asking the simple question: Why do people sell or buy goods via an auction mechanism? There are a number of answers to this question. Often the price of the good or goods has not been determined and the bidders wish to determine the minimum price that they must pay given that they must compete with others for the ownership of this good. From the seller's perspective, submitting goods to an auction may increase the number of buyers, thereby increasing the potential for competitive bidding and higher selling prices. Thus, an auction is a relatively simple mechanism to determine the market-based price, since the bidders who are competing against each other set the price. This mechanism is dynamic and reacts to changes in market conditions. The determination of the selling price by auctions is therefore perceived as both less haphazard and fairer than if the price were set by bilateral negotiations. In the auction case, many players are allowed to participate and all are playing by the same set of rules. Most importantly, if the rules are well designed, the result will have the goods given to the entity that values them the most.

In this paper we will discuss only one-sided auctions and restrict our attention to the case where there is a single seller and multiple buyers. Since the multiple seller/single buyer and the multiple buyer/single seller are symmetric, all results

follow for either case. We will also assume that there are multiple items being sold and that, for at least some of the buyers, the value of buying a collection of the items is greater than the sum of the value of the items individually. This is specifically the case where *combinatorial* or *package-bidding* auctions can provide greater overall efficiency as well as greater revenue to the seller. This auction design is sufficiently general to allow bidders who are willing to buy more quantity but as the quantities increase, the value of the total quantity decreases (often called substitutes condition), and the case where there is neither substitutes nor compliments.

We begin by first classifying auctions into a number of major types: One of the simplest mechanisms is the **first-price (sealed bid) auction.** In this design, all bidders submit their bids by a specified date. The bids are examined simultaneously with the bidder with the highest bid being awarded the object and paying the amount that was bid. The problem with this approach is that the winning bidder may suffer from the "winner's curse", i.e., the bidder may pay more than was necessary to win since the second highest bidder may have submitted a bid price on the item that was far less than the winning bid amount. For this reason, first price auctions encourage bidders to shave some amount off of the highest amount that they are willing to pay in order to not pay more than is necessary. An alternative is the **second-price (sealed bid) auction** whereby the bidder that has submitted the highest bid is awarded the object, but he pays only slightly more (or the same amount) as that bid by the second-highest bidder. In second price auctions with statistically independent private valuations, each bidder has a dominant strategy to bid exactly his valuation. The second price auction also is often called a *Vickrey auction.*

However, often the value of the good is either not completely known or not totally a private valuation that is independent of the value that other bidders palce on the bid. Instead, there is a *common component* to the bid value – that is, the value of the item is non statistically independent of the others bidders but rather there is a common underlying value and bidders have to guess the value that rival bidders may place on the bid. In a common value auction, the item has some unknown value and each agent has partial information about the value. Many high-stakes auctions, such as antique, art and horse auctions fall into this class. In this case, ascending bid auctions, often known as **English (open bid) auctions** are used. Here, the starting price is low and bidders place bids. A standing bid wins the item unless another higher bid is submitted. All bidders observe the previous bids and the new bid must increase the going price by some increment set by the auctioneer. The auction ends when no one is willing to top the standing high bid. With statistically independent private valuations, an English auction is equivalent in terms of payoffs to a second price sealed bid auction.

An alternative auction design is the **Dutch auction** whereby the price is set very high and the price is gradually lowered by a clock until it is stopped by some bidder that is willing to buy some or all of the items up for sale at the current price. Dutch auctions are strategically equivalent to first price sealed bid auctions. The name derives from the fact that tulips were sold via this mechanism and this auction mechanism is still employed for many agricultural products.

Finally, for a large group of identical objects, a variation of the ascending auction is one in which the auction price ascends until there are no new bids. At the end of the auction, all buyers pay the lowest amount bid by any of the winning bidders. Such auctions are called **all-pay one-price auctions.**

Another important auction format issue is whether the auctioneer should auction each of the items sequentially or auction all items simultaneously. The problem with the sequential approach is that bidders must guess the resulting prices of future auctions when determining how to respond to prices in the current auction. A simultaneous ascending auction is where all items are auctioned *simultaneously,* and the auction ends when there are no new bids on any of the items. This auction design allows a bidder to switch back and forth between items as the prices change.

In contrast to the sequential approach of auctioning items, a simultaneous ascending multi-round design specifies that: (a) All items are up for sale at the same time; (b) Bidding takes place in a series of discrete rounds; (c) The auctioneer raises by some increment the price on each item that has multiple bids in the current round; (c) Each bidder is forced to participate throughout the auction; (d) Rounds continue and all items remain available for bidding until there is a round when there are no new bids on any item; and (e) The auction closes at the end of such a round and the items are awarded to the high-standing bidders at the bid price at which they became the high-standing bidder.

We now specify a few additional rules to complete the description. Firstly, the auction has discrete rounds whereby each bidder is required to participate throughout the auction through imposed activity rules. When bidders are forced to bid throughout the auction, all bidders have better price information and are encouraged to bid sincerely. Unlike eBay, the stopping rule relates to the activity of the auction and not a fixed stopping time. Fixed stopping times allow bidders to stay silent until very late in the auction and then present last minute bids in the hopes that other bidders do not have time to respond. Conversely, activity rules force bidders to participate throughout the auction, thereby helping price discovery.

The ascending simultaneous multi-round auction design allows bidders to see the price of all items during the auction and to adjust their strategies and business plans as the auction progresses. Thus, there is a higher likelihood of increased efficiency. This is the general framework that this paper will discuss. However, there are still many details that affect auction outcomes and we will discuss these details in the following section.

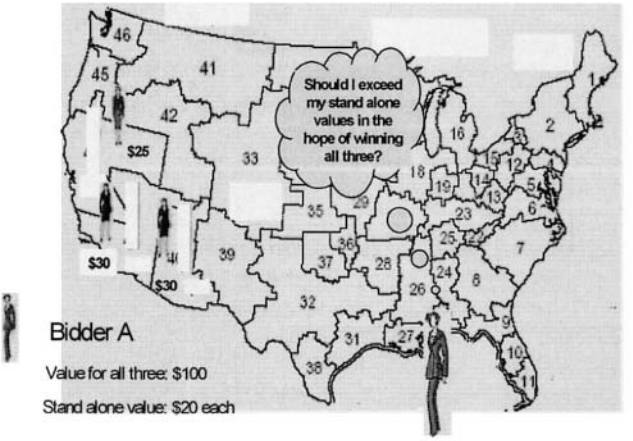
We must first highlight a major shortcoming of the above design: the design as stated above does not consider the issue of synergies among the multiple items being auctioned. Without the ability for a bidder to state that his value for a collection of items is greater than the value of the sum of the individual items, an *exposure* problem can exist. When the winner of each item is determined independently, bidding for a synergistic combination is risky. The bidder may fail to acquire key pieces of the desired combination, but pay prices based on the synergistic gain. Alternatively, the bidder may be forced to bid beyond his valuation in order to secure the synergies and reduce its loss from being stuck with

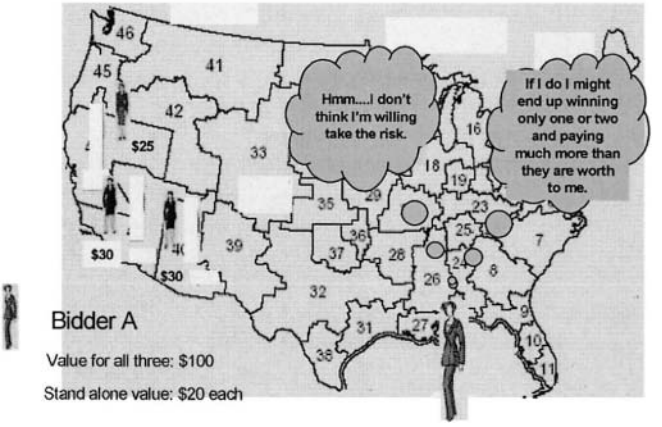
less valuable pieces. Bidding on individual items exposes bidders seeking synergistic combinations to aggregation risk.

Consider the following example whereby a bidder wishes to buy licenses to provide services in a geographic area of the country. Consider that the license is for a technology, such as cellular phone business. Such service providers need enough bandwidth in contiguous regions to support the technology being offered. In the cellular case, too small a geographic presence may make the service less useful to the customer and may make the cost of building the network too expensive to be viable for the provider. Alternatively, larger region coverage also benefits the provider because it incurs less roaming costs of using a competitor's network.

We now show how these concerns play into the auction pricing issues. Consider a bidder that wishes to buy such licenses. To Bidder A, the value of owning the three licenses – covering parts of the Southwest US – is $100, whereas owning the three licenses individually is only $20 each. After many rounds of bidding, Bidder B has bid up one of the licenses to $25, and Bidder C has bid up the other two licenses to $30.each. Bidder A has a dilemma: Should she place bids on these licenses and risk not winning all three, thereby paying too much for what she does win. Or, alternatively, should she leave the auction? And, if this bidder leaves the auction without further bidding, what are the consequences to the seller? If she stops bidding then the auction revenue may total only $85, since the other bidders are satisfied with the current situation. Notice in this case, not only is the revenue low but the auction did not result in an efficient outcome assuming Bidder A had the highest total value on these three items[1]. Alternatively, if the bidder continues to bid, she risks only winning one or two of the items and paying too much.

[1] An efficient outcome is one in which the auction allocates all items to the bidders that value them the most.


Should I exceed my stand alone values in the hope of winning all three?
$25
$30
$30
Bidder A
Value for all three: $100
Stand alone value: $20 each


Hmm....I don't think I'm willing take the risk.
If I do I might end up winning only one or two and paying much more than they are worth to me.
$25
$30
$30
Bidder A
Value for all three: $100
Stand alone value: $20 each

To overcome the exposure problem, many have proposed auction designs that allow bidders to construct packages of subsets of the items. Such auctions are called ***combinatorial auctions.*** We next discuss the issues in designing an auction mechanism that allows bidders to provide package bids. For a more complete description on non-package simultaneous ascending auctions see [5].



2 The Basic Components of a Combinatorial Auction Design

We now begin to describe the details of an ascending multi-round combinatorial auction design with *package* bids. The design must consider how to assure that both the seller and buyers perceive the auction as fair. Under such conditions, the auction is likely to have maximum participation with the end result being a market mechanism that allocates the goods in an optimal manner. We begin the discussion by listing a collection of attributes that we would like this auction to have. We will then discuss possible rules that attempt to obtain – to the degree possible – each of these attributes.

Our goal is to design an auction whose outcome has the following attributes:

(1) Bidders can create and place bids on any and all collections of the objects being auctioned.

(2) The auction results in maximum revenue to the seller.

(3) The auction results in an efficient outcome, i.e., all items are collectively allocated to the bidders that value these items the most.

(4) The auction is fair to all bidders.

(5) The auction ends in a reasonable amount of time.

(6) The auction has limited transaction costs, i.e., the rules are not so difficult or the bidding so complicated that a straightforward bidder finds it difficult to participate.

(7) The auction cannot be gamed, i.e., truthful bidding is an optimal strategy for all bidders.

(8) The auction allows price discovery; to accomplish this end, we choose a multi-round auction framework.

(9) The auction is computationally feasible and scalable, i.e., the auctioneer can calculate the results (both the determination of the provisional winners and be able to set next round bid prices) in a reasonable amount of time. Also, the bidders are capable of determining their next collection of most-profitable packages.

It is not possible to have all such attributes simultaneously. However, the better we are able to design an auction toward achieving these goals, the more likely it is that the outcome will be efficient. For more on what really matters in auction design, see [12]. For an alternative ascending package-bidding design, see [20]. We begin by discussing each of these attributes in turn.



2.1 Package Creation

A package bid expresses a collection of objects that the bidder wishes to win simultaneously and the price that he is willing to spend to obtain that package. Ideally, the auction design should allow bidders to submit as many packages as they might desire. However, as the number of items in the auction increases, the number of possible packages increases exponentially. We want the auction to remain computationally tractable; and we do not want allow a bidder the ability to subvert the auction by overloading the system with a near-infinite number of bids.

We, therefore, want a bidding language that is concise and does not force the bidder to *enumerate* every package that they might at some point consider profitable. Thus, we want a bidding language that provides the bidders with the flexibility to legitimately describe what they want. There are a number of languages that are both concise and allow complete expressivity. We refer the reader to the research of [3, 11, and 18] for complete descriptions of the alternative languages.

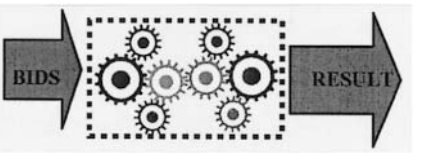
To begin, let us consider the simplest language: bidders submit package bids and may win any combination of these bids, as long as each item is awarded only once (i.e., bids are mutually exclusive if two bids contain the same item). The problem with this language, known in the computer science literature as an "OR" language, is that a bidder can have a new type of *exposure* problem: a bidder wants to let the system know that he is willing to win either of two packages but cannot afford to win both simultaneously. Fujishima et al. [11] propose a generalization of this language called OR* that overcomes this problem. Namely, each bidder is supplied "dummy" items (these items have no intrinsic value to any of the participants). When a bidder places the same dummy item into multiple packages,

it is telling the auctioneer that he wishes to win at most one of these collections of packages. This language is fully expressive as long as bidders are supplied sufficient dummy items. We note that this language is relatively simple for bidders to understand and use, as was shown in a Sears Corporation supply-chain transportation auction. In that auction, all bids were treated as "OR" bids by the system. Some bidders cleverly chose a relatively cheap item to place in multiple bids thereby making these bids mutually exclusive, see [14]).

By using this OR* language, the auctioneer can treat all bids as "OR" bids, and the optimization problem that determines the optimal set of packages will automatically treat bids with the same dummy item as mutually exclusive. For computational feasibility, the auctioneer may still limit the maximum number of bids that a bidder can submit throughout the course of the auction.

Bidders are not required to identify or create their packages before the start of the auction, but may create their packages as the auction progresses. A bidder may modify or delete a package it has created up until the point where it has bid on the package and the round has closed. If the bidder submits a bid on a package and subsequently removes the bid during the same round, the bidder has the option of also deleting or modifying the package. However, once a bidder bids on a package and the round closes, the package may not be modified or deleted and counts as one of the bidder's allowable packages.

Another rule imposed on bids is that all bids that have been placed when the round ends will be considered active throughout the auction – that is, these bids will remain in the optimization problem and may become winning at any subsequent round in the auction. This rule on extending the life of a bid throughout the auction forces bidders to provide sincere bids. In addition, it also helps to overcome a problem inherent in package bidding, that of "package fit". We will say more about this in the section on efficient outcomes.



2.2 The Auction Maximizes Revenue

The auctioneer receives a collection of bids and wishes to compute the set of bids that maximizes revenue and that assures that each bid is awarded precisely once. We will define *winning bids* in a package bidding auction as the set of bids on individual licenses and packages that maximizes revenue when the auction closes, assigning each license only once, and we will define *provisionally winning bids* as the set of bids that maximize revenue in a particular round (these bids would win if the auction were to close in that round).

When determining winning and provisionally winning bids, all bids made in every round throughout the course of the auction (except for bids that are placed and subsequently removed during the same round) will be considered. In addition, each license is treated as having a bid placed by the auctioneer at some small amount less than the minimum opening bid or reserve price, whichever is larger.

This procedure will ensure that a bid on a license or package is winning only if it exceeds this amount.

Let $x_b = 1$ if bid b is awarded, and 0 otherwise. The matrix A, represents the packages being bid, where each row of the matrix corresponds to an item in the action and the entry a_{ib} of the matrix will equal 1 if bid b includes item i, and 0 otherwise. The *winner determination problem* can now be stated as:

$$\underset{x}{Max} \sum_{b=1}^{\#Bids} BidAmt_b x_b$$

$$subject\ to:$$

$$Ax = 1 \text{ (each item is awarded once)}$$

$$x \in \{0,1\}$$

This set-partitioning optimization problem determines the maximum revenue possible. However, since there can be more than one set of consistent bids that produces the maximum revenue, we wish to have a procedure that does not bias in favor of certain collections winning. We therefore want a procedure that randomly selects among tied sets. The tie-breaking procedure works as follows: We first assign a selection number to each bid. A bid's *selection number* is the sum of n random numbers, where n is the number of items comprising the bid's package. We then solve a second optimization problem which determines the set that maximizes the sum of the selection numbers subject to the constraint that each item must be awarded exactly once and that the revenue of a winning set must equal the maximum revenue (as calculated in the winner determination problem, above).

$$\underset{x}{Max} \sum_{b=1}^{\#Bids} SelectionNumber_b x_b$$

$$subject\ to:$$

$$Ax = 1 \text{ (each item is awarded once)}$$

$$\sum_{b=1}^{\#Bids} BidAmt_b x_b = \text{ Max Revenue}$$

$$x \in \{0,1\}$$

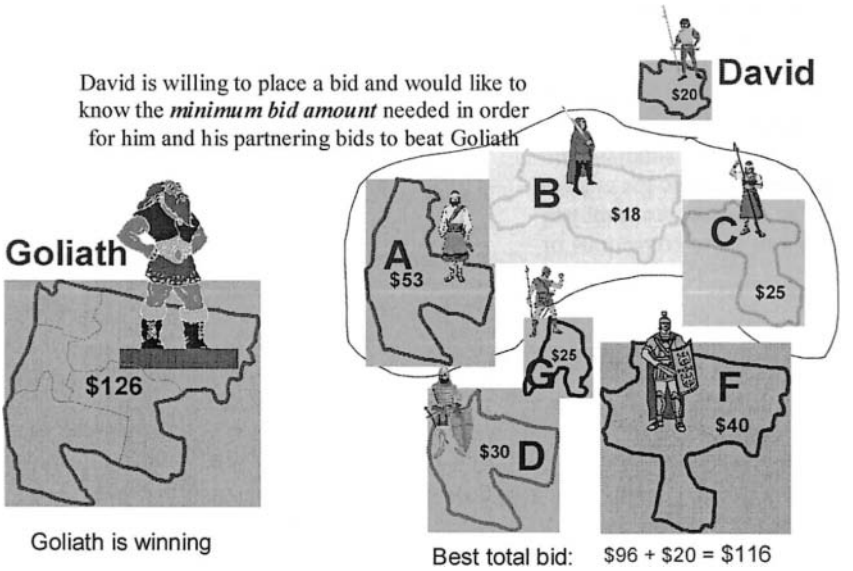
Thus, the set of provisionally winning bids is the set of consistent bids that maximizes revenue *and*, among all such sets, maximizes the sum of the selection numbers. Each bid is assigned a new selection number in each round. Note that a tied bid that was provisionally winning in some round may not be winning in a subsequent round, even if no new bids are placed on any of the items in that package. Changes in winning bid designations among tied bidders from round to round encourage sincere and continuous bidding. When there are no new bids by

any bidder, the auction closes. The winner determination problem is not re-solved, and the previous round's *provisional winners* become the *winning* bidders

The determination of a set of provisionally winning bids that assure that ties are broken randomly provides not only maximum revenue to the seller but also assures that the optimization solver does not bias in favor of some collection of bidders or bids over some other equally deserving set. Thus, the fairness criterion is enhanced through this second optimization problem.

2.3 The Auction Works Toward Efficiency Outcomes

Combinatorial auction designs are likely to increase efficiency by removing the exposure problem. The package feature does, however, create a new problem for smaller bidders, known as the *threshold problem,* whereby smaller bidders may have difficulty overcoming the tentative high bid of a large bidder. An extreme case occurs when a bidder creates a package of *all* items in the auction and places a large bid amount on this package. Smaller bidders must be able to each incrementally increase their bids so that *collectively* they can overcome the bid price of this large package bid. We explain this problem through an example: Consider Goliath who has placed a bid of $126 on the four properties that make up the Northwest region of the country. David wants to purchase only one of these areas, and he has a bid of $20 on that property and can afford to bid up to $25. David needs to coordinate his bid with the bids of the other three bidders, who collectively could overcome the bid of Goliath. However, we do not want the bidders to collude. Instead, we want to tell each of the smaller bidders the amount that they need to contribute to the total difference in the bid price of their collective bids and Goliath's bid.



Thus, bidders need *price information* that helps them determine how much of the total bid difference they should contribute. With this information, David will know the minimum amount that he needs to spend in order for him and his partnering bids to beat Goliath, assuming all are willing to increase their bids by the apportioned amount. We don't know how Goliath (or other bidders) will bid in the next round, but we would like to give David an idea of what it would take for him and his partnering bids to beat Goliath's current winning bid. After calculating the provisionally winning bids for the round, a *current price estimate* is determined for each license (the "going-rate" of the license).

The current price estimate, in our auction design, is based on dual information obtained from the linear relaxation to the winner determination problem. If the winner determination problem were a linear (rather than integer optimization), the values of the corresponding dual variables provide the proper pricing information for each item and could be automatically provided with no additional computation. However, since the winner-determination problem is an integer problem, these linear dual prices are not correct and will *overestimate* the true prices (since the linear programming value may be larger than the integer programming value). We will start with the dual concept and obtain estimates of the prices that are consistent with the total revenue obtained from the integer solution. We first present the dual to the winner-determination problem.

[*WDP LP Dual*]:

$$\min \sum_{i \in L} \pi_i$$

Subject to:

$$\sum_{i \in I^j} \pi_i \geq b_j, \qquad \forall \quad j \in B$$

$$\pi_i \geq r_i, \qquad \forall \quad i \in I$$

where, I^j is the set of items contained in bid j.

If the linear optimization relaxation to the winner-determination problem was integer optimal, then the resulting prices, π_i, would satisfy the inequalities, $\sum_{i \in I^j} \pi_i \geq b, \ \forall j \in B$ and $\sum_{i \in I^j} \pi_i = b_j, \ \forall j \in W$, where W is the set of winning bids. Since such prices might not exist, one approach is to use *pseudo-dual prices*[2] rather than the dual prices. These pseudo-dual prices are obtained by forcing the sum of the dual prices of the items comprising a provisionally winning bid to equal its respective bid amount, but allowing the prices of non-winning bids to be less than the maximum price bid for that package.

Thus, the pseudo-dual prices for each item i, denoted by π_i are required to satisfy the following constraints:

$$\sum_{i \in I^j} \pi_i + \delta_j \geq b_j, \forall \ j \in B \setminus W \qquad (1)$$

$$\sum_{i \in I^j} \pi_i = b_j, \qquad \forall \ j \in W \qquad (2)$$

$$\delta_j \geq 0, \qquad \forall \ j \in B \setminus W \qquad (3)$$

$$\pi_i \geq r_i \qquad \forall \ i \in I \qquad (4)$$

where $W \subset B$ is the provisionally winning bid set and δ_j is a slack variable that represents the difference between the bid amount of non-winning bid j and the sum of pseudo-dual prices of the items contained in non-winning bid j. Constraint set (2) assures that the bid price of winning-bid packages exactly equal the winning bid amount. Constraint set (4) assures that the price of any item cannot go below the reserve or starting price and thereby assures that the price of any package is,

[2] In our research we found this term first applied to auction pricing in the paper by Rassenti, Smith and Bulfin [19].

therefore, greater than or equal to the sum of the reserve prices of the items that make up that package.

By keeping constraints (1)-(4), we have considerable flexibility in choosing an objective function that will help in selecting among multiple solutions, while still ensuring that the sum of the pseudo-dual prices yields the maximum revenue of the round. There are likely to be many solutions (i.e., many sets of dual prices) that satisfy this constraint set. Since we want the solution to be as close to integer optimality as possible, we want to minimize the total infeasibility, i.e., we wish to minimize $\sum_{j \in B \setminus W} \delta_j$. Pseudo-dual prices for each item i, denoted π_i, can be obtained by solving the following linear program:

$[CP]$:

$$z_\delta^* = \min \sum_{j \in B \setminus W} \delta_j$$

Subject to:

$$\sum_{i \in I^j} \pi_i + \delta_j \geq b_j, \qquad \forall\, j \in B \setminus W \qquad (1)$$

$$\sum_{i \in I^j} \pi_i = b_j, \qquad \forall\, j \in W \qquad (2)$$

$$\delta_j \geq 0, \qquad \forall\, j \in B \setminus W \qquad (3)$$

$$\pi_i \geq r_i, \qquad \forall\, i \in I \qquad (4)$$

Constraint set (2) assures that the bid price of a winning bid is equal to the winning bid amount. Constraint set (1) together with the objective function tries to enforce dual feasibility as much as possible. Constraint set (4) forces the prices on all items to be at least the minimum price set by the auctioneer.

The solution to this problem is not necessarily unique. In fact, testing has shown that using [CP] in an iterative auction can result in significant changes in the pseudo-dual price of an item from round to round. Although the prices of items should be allowed to reflect real change (both increases and decreases) in the way bidders value the items over time, large oscillations in minimum acceptable bid amounts for the same bid that are due to factors unrelated to bidder activity, such as multiple optimal primal or dual solutions, can be confusing to bidders and may encourage unwanted gaming of the system. We therefore solve a second optimization problem that chooses a solution in a way that reduces the magnitude of price fluctuations between rounds. This method is known as *smoothed anchoring,* since the method anchors on exponentially smoothed prices from the previous round when determining prices for the current round.

First, [CP] is solved to obtain the minimum sum of slack. Second, a linear quadratic program is solved with an objective function that applies the concepts of exponential smoothing to choose among alternative pseudo-dual prices with the additional constraint on the problem that the sum of the slack variables equals z^* (the optimal value of [CP]). This objective function minimizes the sum of the

squared distances of the resulting pseudo-dual prices in round t from their respective smoothed prices in round *t-1*.

Let π_i^t be the pseudo-dual price of item i in round t. The smoothed price for item i in round t is calculated using the following exponential smoothing formula:

$$p_i^t = \alpha\pi_i^t + (1-\alpha)p_i^{t-1}$$

where p_i^{t-1} is the smoothed price in round t-1,

$0 \le \alpha \le 1$, and

p_i^0 = the minimum opening bid amount for item i.

The following quadratic program [QP] will find the pseudo-dual price, π_i^t for each item i in round t that minimizes the sum of the squared distances from the respective smoothed prices in round *t-1*, while assuring that the pseudo-dual prices sum up to the provisionally winning bid amounts.

$$[QP]:$$

$$\min \sum_{i \in L} (\pi_i^t - p_i^{t-1})^2$$

Subject to:

$$\sum_{i \in I^j} \pi_i^t + \delta_j \ge b_j, \qquad \forall\, j \in B^t \setminus W^t \qquad (1)$$

$$\sum_{i \in I^j} \pi_i^t = b_j, \qquad \forall\, j \in W^t \qquad (2)$$

$$\sum_{j \in B^t \setminus W^t} \delta_j = z_\delta^* \qquad (3)$$

$$\delta_j \ge 0, \qquad \forall\, j \in B^t \setminus W^t \qquad (4)$$

$$\pi_i^t \ge r_i, \qquad \forall\, i \in I \qquad (5)$$

Note that problem [QP] has the same constraints as [CP], but has added the additional restriction (3) that the sum of the δ_j's is fixed to the value z_δ^*, the optimal value from [CP].

Among alternative prices that satisfy all constraints, the objective function of this optimization problem chooses one that forces the pseudo-dual prices to be as close as possible to the previous round's smoothed price. Thus, this method is called the *smoothed anchoring method,* since it "anchors" on the smoothed prices

when solving for the pseudo-dual prices. The current price estimate for item i in round t is therefore the pseudo-dual price, π_i^t, obtained by solving [QP].[3]

There are a number of alternative pricing algorithms that have been suggested in the literature. For an overview of this research, see [9] and the references therein. In this paper, we highlight the fact that bidders need good price information and that such information can be obtained by having the auctioneer perform addition optimizations to determine how to allocate the price of a big package to smaller packages.

Although linear pricing cannot accommodate all aspects of the pricing associated with the non-linear, non-convex, winner determination problem, there are still good reasons for considering its use for determining prices and therefore future bid requirements. In an ascending bid auction, bidders need pricing information that is easy to use and that is perceived to be "fair". We say that linear pricing is "easy to use" in that bidders can quickly compute the price of any package – whether or not it had been previously bid. With linear prices, bidders can create a new package "on the fly" and know its price by merely adding the prices of the items that compose the package. When choosing such prices, we want these prices to be "competitive", i.e., using these prices assures the bidder that they have some possibility of winning in the next round. Bidders also perceive these prices to be "fair," since all bidders must act on the same information. Finally, linear prices are likely to move the auction first determining the "competitive" price and then adding an increment to this price, so that the new bids are likely to out-bid the standing high bids.

These "current price estimates" use information from all current and past bids on all packages containing that license, to approximate a "price" for each item in every round of the auction. Until a bid is placed on an item or on a package containing that item, by any bidder in any round, the current price estimate is the minimum opening bid price. If a bid on a single item is part of the provisionally winning set of bids in a round, then the current price estimate for that item in that round will be equal to the provisionally winning bid. Generally, however, if an item is part of a package bid in the provisionally winning set, the current price estimate for an individual item is calculated using all of the bids placed that involve that item, since those bids together yield information on how bidders value that license. The set of current price estimates, one for each item, is then used as the basis for determining minimum acceptable bids for the next round, under the minimum acceptable bid price for the following round. Specifically, the next round price for a package is the sum of the current price estimates for the items comprising that package, plus some percentage of the sum.

---

[3] An alternative objective function to the quadratic objective function of [QP] is to minimize the maximum distance from the smoothed prices, subject to the same set of constraints. This alternative formulation would involve solving a sequence of minimization problems as was done in [11].

Another rule that can improve the efficiency is to allow "last and best" bids on any package. This rule allows bidders to bid on a package at a price that is in between the last bid price and the current round's bid price. When a bidder exercises this option, they have placed a "last and best" bid on that package and cannot again bid on that package. Thus, the last-and-best rule helps to overcome any problems with the increment setting.



2.4 The Auction is Fair

We have already provided some of the rules that promote fairness among the bidders. Choosing among winning bidders randomly eliminates any possible bias among winning bid sets. Our price estimates provide all bidders with the same price information and help small bidders overcome the threshold problem by providing useful price information.

We also promote the idea that the increment that we use to increase the current price estimate should be based on the amount of demand for that item. Thus, items with less competition increase slowly, while items with high demand will increase more rapidly. Thus, as the price approaches its final price, the increment is sufficiently small to be able to obtain the efficient outcome at approximately second price. When there is much competition, having a large increment helps move the auction along.

Most important to a fair auction is that we impose rules of eligibility and activity that force all players to bid each round and thereby provide good price discovery to all bidders. These rules dictate how bidders act in each round. Typically, high-stakes auctions require an "up front" deposit that assures that bidders have the ability to pay for what they may eventually purchase. This refundable deposit also defines the bidder's initial eligibility – the maximum quantity of items that the bidder can bid for. A bidder interested in winning a large quantity of items would have to submit a large deposit. The deposit provides some assurance that the bids are serious and that the bidders will not default at the end of the auction, leaving goods unsold. Some auctions require bidders to increase the deposit as bid prices increase.

In addition to an upfront deposit, the auction needs activity rules for improving price discovery, by requiring a bidder to bid in a consistent way throughout the auction. It forces a bidder to maintain a minimum level of bidding activity to preserve his current eligibility. Thus, a bidder desiring a large quantity at the end of the auction (when prices are high) must bid for a large quantity early in the auction (when prices are low). As the auction progresses, the activity requirement increases, reducing a bidder's flexibility. The lower activity requirement early in the auction gives the bidder greater flexibility in shifting among packages early on when there is the most uncertainty about what will be obtainable. As the auction

progresses, if the bidder cannot afford to bid on sufficient items to maintain his current eligibility, then his eligibility will be reduced so that it is consistent with his current bidding units. Once eligibility is decreased, it can never be increased. Thus, activity rules keep the auction moving along and do not allow the sniping[4] that occurs in other auction designs.

When eligibility and activity rules are enforced, bidders are constrained to sincerely bid throughout the auction, thereby providing all bidders with adequate price information.

Since the auction requires that bidders continually be active and since there may be times when a bidder cannot participate, each bidder is provided with a few activity waivers that may be used at the bidder's discretion during the course of the auction, as set forth below. Use of an activity rule waiver preserves the bidder's current bidding activity, despite the bidder's lack of activity in the current round. An activity rule waiver applies to an entire round of bidding and not to a particular license or package. Activity rule waivers are principally a mechanism for auction participants to avoid the loss of auction activity in the event that exigent circumstances prevent them from placing a bid in a particular round. If a bidder has no waivers remaining and does not satisfy the activity requirement, his activity will be permanently adjusted, possibly eliminating the bidder from the auction.

Precisely how the activity and eligibility rules are set matters and must be depend upon the type of auction – the value of the items being auctioned, the projected length of the auction, the number of participants, etc. In many high-stakes auctions (spectrum, electricity, etc) these activity rules have proven highly successful, see [12, 16 and 17].



2.5 The Auction is Fast

In experimental studies, Cybernomics [7] showed that combinatorial auctions can be more efficient, but can also take longer to complete – in terms of the number of rounds – than a non-package simultaneous multi-round counterpart. Both the increment and the closing rule can impact speed, but can also impact efficiency. A closing rule that requires no new bids can therefore force a very long auction. Ausubel et al. [1] argue for a hybrid auction in which the first phase is a simultaneous multi-round package bidding auction that allows price discovery to occur, and then, as the revenue begins to trail off, one stops this ascending phase of the auction and implements an ascending proxy auction. Thus, one can consider using a proxy auction as a final round sealed-bid auction, where the bidders express

[4] Bid sniping occurs at the last minute of a fixed-time ending. The purpose of sniping is to give other bidders no chance of responding to an offer. In this way, a bidder can acquire price information from other bidders but does not reciprocate, since throughout most of the auction, the bidder is silent.

their final values but pay only what is necessary to win. Ausubel and Milgrom [2] show that the proxy auction results in a Nash equilibrium where an optimal strategy for each bidder is to bid sincerely. Much work is yet to be done to specify how the rules from the ascending auction phase mesh with the rules to the final round to ensure that the bidders are forced to bid sincerely rather than "parking"[5] in the ascending phase, so that the necessary price discovery is achieved. However, the ability to stop the auction early and still obtain efficient outcomes is promising.



2.6 The Auction is Difficult to Game

In early non-package simultaneous round auctions, bidders were quite creative in using bid amounts to signal both their identities and the items that they were most interested in obtaining. Through this signaling, they attempted to determine how to "cut up the pie" outside of the auction and thereby receive the items at much lower rates. To overcome this gaming, one can (1) limit the information provided back to bidders, and (2) only allow bids in specified amounts. (The FCC refers to this as "click-box" bidding, since a bidder can "click" on only one of ten pre-specified bid amounts.) Experimental results have shown that there are significant incentives for collusion, and that limiting the identities of the bidders improves the outcomes while still providing good price information. Thus, we recommend that auctioneers carefully consider limiting identifying information and not allowing bidders to supply such information through their bid amounts.

The rule that all bids remain active throughout the auction and cannot be withdrawn is deterrence to "parking," since a bid early on may become winning late in the auction. Also, by limiting the total number of packages that a bidder can submit, one forces bidders to bid sincerely.



2.7 The Auction Should Help Bidder Handle the Complexities

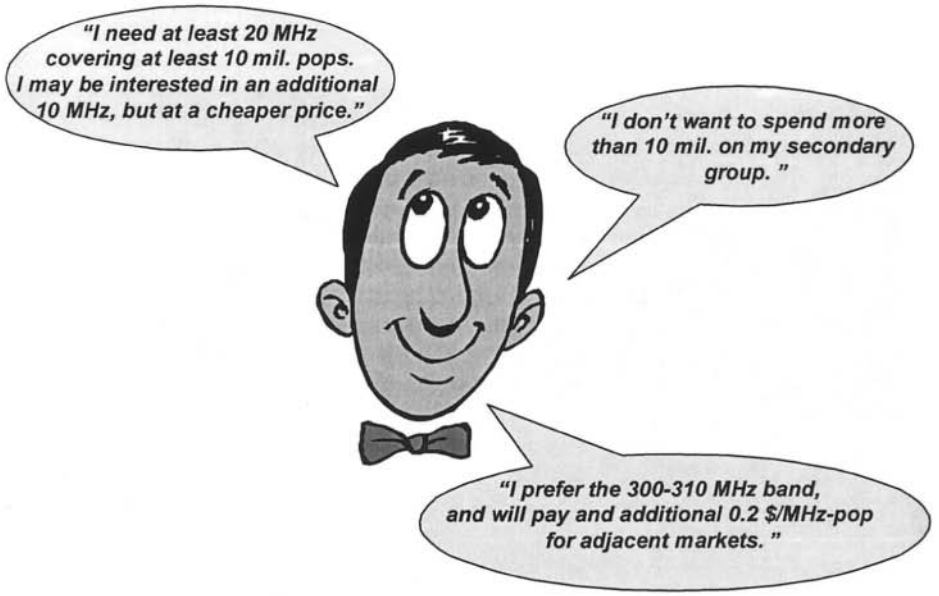
Since multi-item auctions are complex and require bidders to consider multiple alternative bid options, we believe that it important that the computer software used for communication between the bidder and the auctioneer be easy to use and understand. Good graphical user interfaces help bidders to feel comfortable that they understand the current state of the auction (they have been able to find the

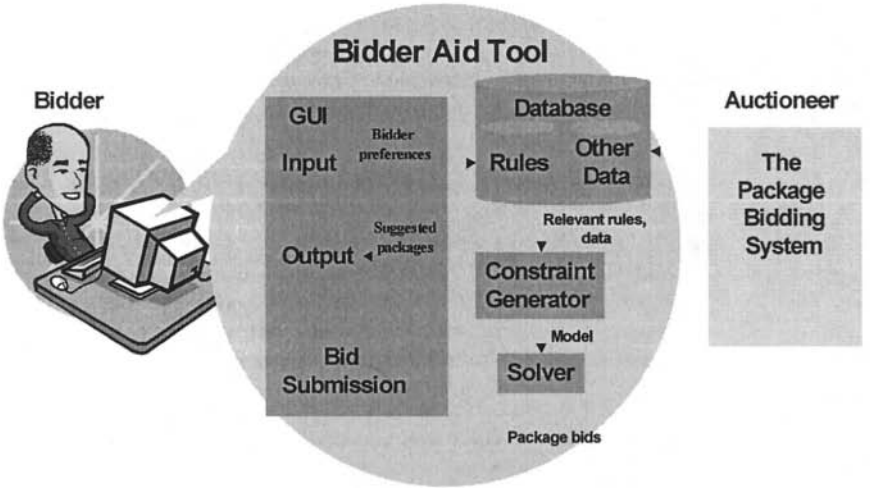
[5] Parking is a gaming strategy whereby bidders, early in the auction, bid on lower-priced items in order to keep the prices low on the items they value the most and wait until late in the auction to begin bidding on the items that they value the most

current price information, the items they are winning, the amount of bidding necessary to remain eligible, their dollar exposure based on what they have bid, etc.). The system must also provide easy ways for them to input their next moves and confirm that they have provided the system with the correct information. As the use of auctions is spreading, computer interfaces for such processes continue to improve and to provide better ways of displaying information to the users through charts, graphs and pictures. We foresee continued improvement in this area.

These tools do not, however, help the bidder determine the optimal combination of items to bundle as a package and the optimal number of packages to supply to the system. Although all of us face similar decisions regularly – we are bombarded with a near limitless variety of purchasing opportunities for even the simplest of tasks, e.g., what to eat for dinner – when the stakes are high and the number of substitutes is limited, the problem of providing decision support tools becomes more pressing. Thus, there is a need for the auction community to create tools that will help narrow the scope of the search and suggest good packages based on the bidder-specific business plans and preferences. We refer the reader to a few papers that have begun work in this very important area. For some discussion on this topic, see [3, 4, and 10], Also, [8, and 19] provide alternative ways of bidders expressing preferences that preclude the bidder specifying specific packages to the auctioneer.

We present as an example a tool that is being built by FCC contractors to assist bidders in FCC spectrum auctions. In this tool, bidders supply information in terms of their business plans: What geographic areas are their primary, secondary and tertiary markets? What is the minimum and maximum bandwidth needs? Do they prefer certain bands? How much population do they need to have a viable business? How much are they willing to spend per MHz population unit? With this information, the tool translates their needs into constraints for an optimization model that maximizes the bidder's profit, given the current purchase prices for the licenses of interest. Packages are then provided back to the user for evaluation. The process continues with the bidder having the ability to change his business plan or specifications within a plan. The tool re-optimizes and creates new packages. Once the bidder is satisfied with the outputs, he indicates the packages that are best, and these are supplied to the FCC.


"I need at least 20 MHz covering at least 10 mil. pops. I may be interested in an additional 10 MHz, but at a cheaper price."
"I don't want to spend more than 10 mil. on my secondary group."
"I prefer the 300-310 MHz band, and will pay and additional 0.2 $/MHz-pop for adjacent markets."


Bidder Aid Tool
Bidder
GUI
Input
Bidder preferences
Output
Suggested packages
Bid Submission
Database
Rules
Other Data
Relevant rules, data
Constraint Generator
Model
Solver
Package bids
Auctioneer
The Package Bidding System

Tools similar in nature are being used by the major supply-chain auction providers. They must always be designed for the specific application and include all of the constraints relevant to that business concern. However, with such tools, more entities are willing to participate in the auction thereby improving the overall competitiveness of the auction.

2.8 The Auction Should Be Computable and Verifiable

As discussed previously, the auction requires that most of the computational burden be placed on the auctioneer. This seems appropriate and the direction that is being taken by industry. There are a few corporations that supply auction technology to various companies to use in their supply-chain auctions. These companies have the incentives and the capabilities to solve the difficult optimization problems that arise in combinatorial auctions. Specifically, the combinatorial optimization problem that determines the provisional winners in each round determines how ties are broken and computes prices for the subsequent rounds.

Off the shelf software, such as CPLEX or XPRESS have proven to be able to solve such problems in reasonable times (less than 30 minutes). Thus, although there is much in the literature that argues against combinatorial auctions because of the computational burden, the optimization software has proven up to be capable of handling the problems that are currently being considered applicable for this type of auction. For more on the computational issues in computing winner determination problems, see [15].

Although work continues on fine-tuning the auction software, the computability demands on the auctioneer have not been found to be a limiting factor. We believe that tools that help a bidder construct packages will also be computationally tractable, since the decisions being made by an individual bidder compose a smaller set of reasonable alternatives.

3 Combinatorial Auction Design Components

In this paper, we have presented a general framework for a simultaneous multi-round ascending package bidding auction. This design has a number of characteristics:

1. The auction, although complex for the auctioneer, is relatively easy for the bidders. Indeed, if the bidder chooses to only submit item bids and treat the auction as a non-package ascending auction, the bidding language allows such bids.

2. The complexity of the auction for the auctioneer, although substantial, is computationally feasible with easily available off-the-shelf state-of-the-art combinatorial optimization software.

3. The design allows bidders complete flexibility in package construction and provides a language that allows bidder to declare that certain packages should be treated as mutually exclusive, while keeping the language concise and completely flexible. In addition, the language maintains the set-partitioning structure to the winner determination problem, thereby not complicating the optimization problem.

4. The auction advocates that the auctioneer limit the total number of bids that a bidder can submit, but allows for that limit to be relatively high, e.g., 5000 bids per bidder.

5. The auction requires upfront payments to ensure that the bidders are capable of paying for the items won.

6. The auction specifies eligibility and activity rules in order to protect against frivolous or insincere bidding and to encourage price discovery.

7. The auction provides price and demand information after each round.

8. The auction does not disclose the identities of individual bids, since such information provides an opportunity for bidders to collude and to engage in destructive or punitive behavior.

9. The auction allows a stopping rule that encourages bidding throughout the auction.

10. The auction allows for a final round to speed up the auction and improve efficiency.

11. The auction allows bidders to provide "last and best" bids when the price increase specified by the auctioneer is too high, and yet the bidder is willing to pay more than the prior round's bid price.

12. Bids are considered "active" throughout the auction; bid withdrawals are not allowed. The forcing of all bids to be active throughout the auction is another rule to force bidders to bid sincerely.

When one considers an auction design, one must consider how all of the rules fit together, so as to assure that the overall goals of the auction are met. Much has been learned about how to set such rules from the many non-package multi-round ascending auctions that have been held. Adding the additional characteristic of allowing package bids created some new issues, such as how to price packages and how to overcome the threshold problem. One also needed to able to answer the question of whether it was possible to compute the winners in each round in a timely fashion.

In order to study these issues, we created a simulation tool that allowed us to run auctions of varying size and complexity quickly. We could "stress test" the optimization procedures and be able to study much larger auctions than would be possible with human subjects. Although this testing allowed us to better understand how auction rules altered computability and length of auction, our simulations did not allow us to test how well the auction would do when human subjects creatively try to game the system. The FCC is therefore building a new simulation tool that allows computer agents (robotic bidders) to compete in a auction with human subjects. In this way, we can begin to see how these rules impact gaming strategies, while still providing a test environment that allows tests where there are more than a few dozen items and for auctions that are likely to be more complex and go over several days or weeks rather than hours. We look forward to this new avenue of research.

If this short description of issues associated with combinatorial auctions has sparked your interest, we recommend the book, *Combinatorial Auctions* [6], which provides detailed discussions on all aspects of this important new research area.

Acknowledgments

This position paper is an outcome of research that was partially funded by the Federal Communications Commission under a contract to CompuTech Inc. and partially funded by the NSF under grant IIS-0325074. All views presented in this paper are mine and do not necessarily reflect the views of the Federal Communications Commission or its staff. I take complete responsible for all errors or omissions within this paper but want to credit my co-researchers, Melissa Dunford, David Johnson, Dinesh Menon, and Rudy Sultana of Decisive Analytics, Inc. for their extraordinary contributions to the FCC auction project. A special thanks to Melissa Dunford, who created most of the graphics for this paper, and to Evan Kwerel of the FCC for the many discussions that helped me understand many of the practical issues in auction design.

References

1. Lawrence M. Ausubel, Peter Cramton, and Paul Milgrom. The Clock Proxy Auction, in *Combinatorial Auctions,* Cramton, Peter, Yoav Shoham and Richard Steinberg, eds., MIT Press, 113-136, 2005.
2. Lawrence M. Ausubel and Paul Milgrom. Ascending Auctions with Package Bidding. *Frontiers of Theoretical Economics* 1: 1-42, 2002.
3. Craig Boutilier and Hoger H. Hoos. Bidding Languages for Combinatorial Auctions. *Seventh International Joint Conference on Artificial Intelligence (IICAI-01):* 1211-1217, 2001.
4. Craig Boutilier, Tuomas Sandholm, and Rob Shields. Eliciting Bid Taker Non-Price Preferences in Combinatorial Auctions. *Proceedings of the National Conference on Artificial Intelligence,* V. Khu-Smith and C.J. Mitchell, eds., San Jose, CA, 204-211, 2004.
5. Peter Cramton. Simultaneous Ascending Auctions, in *Combinatorial Auctions*, Cramton, Peter, Yoav Shoham and Richard Steinberg, eds., MIT Press, 99-114, 2005.
6. Peter Cramton, Yoav Shoham and Richard Steinberg. *Combinatorial Auctions.* MIT Press, 2005.
7. Cybernomics. An experimental comparison of the simultaneous multiple round auction and the CRA combinatorial auction. Discussion paper. Report to the Federal Communications Commission. 2000. Available at: *http://wireless.fcc.gov/auctions/conferences/combin2000/releases/98540191.pdf*
8. Robert Day and S. Raghavan. Assignment Preferences and Combinatorial Auctions. working paper, Operations and Information Management School of Business, University of Connecticut., 2005. Available at: *http://users.business.uconn.edu/bday/index.html*
9. Melissa Dunford, Karla Hoffman, Dinesh Menon, Rudy Sultana, and Thomas Wilson. Price Estimates in Ascending Combinatorial Auctions. Technical Report, George Mason University, Systems Engineering and Operations Research Department, Fairfax, VA, 2003.
10. Wedad Elmaghraby and Pinar Keskinocak. Combinatorial Auctions in Procurement, *The Practice of Supply Chain Management*, C. Billington, T. Harrison, H. Lee, J. Neale (editors), Kluwer Academic Publishers. 2003
11. Yuz Fujishima, Kevin Leyton-Brown, and Yoav Shoham. Taming the Computational Complexity of Combinatorial Auctions: Optimal and Approximate Approaches. *Proceedings of IJCAI 1999* 548-553, 1999.
12. Paul Klemperer, What Really Matters in Auction Design. *Journal of Economic Perspectives,* 16: 169-189, 2002.
13. Anthony M. Kwasnica, John O. Ledyard, David Porter, and Christine DeMartini. A New and Improved Design for Multi-object Iterative Auctions. *Management Science*, 51: 419-4234, 2005.
14. John O. Ledyard, Mark Olson, David Porter, Joseph A. Swanson and David P. Torma. The First Use of a Combined Value Auction for Transportation Services. ***Interfaces***. 32: 4-12, 2002.

15. Kevin Leyton-Brown, Eugene Nudelman, and Yoav Shoham. Empirical Hardness Models, in *Combinatorial Auctions*, Cramton, Peter, Yoav Shoham and Richard Steinberg, eds., MIT Press, 479-503, 2005.
16. John McMillan. *Reinventing the Bazaar: A Natural History of Markets.* Norton Press, 2002.
17. Paul Milgrom. *Putting Auction Theory to Work.* Cambridge Press, 2004.
18. Noam Nisan. Bidding and Allocation in Combinatorial Auctions. *Proceedings of the 2nd ACM Conference on Electronic Commerce,* 1-12, 2000.
19. David C. Parkes. Auction Design with Costly Preference Elicitation. *Annals of Mathematics and AI,* 44: 269-302, 2005.
20. David C. Parkes and Lyle H. Ungar. Iterative Combinatorial Auctions: Theory and Practice. *Proceedings of the 17th National Conference on Artificial Intelligence* (AAAI-00), 74-81, 2000.
21. Steve J. Rassenti, Vernon I. Smith and Robert I. Bulfin. A Combinatorial Mechanism for Airport Time Slot Allocation. *Bell Journal of Economics,* 13: 402-417, 1982.

Label-Correcting Shortest Path Algorithms Revisited

María G. Bardossy and Douglas R. Shier

Department of Mathematical Sciences
Clemson University
Clemson, SC 29634-0975
mgbardossy@gmail.com, shierd@clemson.edu

Summary. In this paper we study label-correcting algorithms, which are routinely used to solve single-source shortest path problems. Several variants of label-correcting algorithms have been proposed in the literature, differing primarily in the strategy implemented to handle the candidate list of nodes. In particular, we study both one-list and two-list versions of the basic label-correcting algorithm; these variants implement either one or two lists to manage the set of candidate nodes. We examine the theoretical complexity and empirical behavior of these algorithms. In contrast to previous studies of shortest path algorithms, our focus is on explaining observed empirical performance in terms of certain intrinisic properties of the algorithms (namely, "sharpness" and "maturity"). In addition, a new variant of the label-correcting algorithm is proposed (PRED), which insures a type of "local sharpness" relative to the candidate list. Computational evidence suggests that this new algorithm, in both one-list and two-list versions, performs quite well in practice and deserves further study.

Key words: Shortest paths; label-correcting algorithms; sharpness; operation counts.

1 Introduction

The efficient management of communication, transportation and distribution systems often requires finding shortest paths from one point to another in the underlying network. Other situations (equipment replacement, project scheduling, cash flow management, DNA sequencing, artificial intelligence, data compression), that at first glance do not seem related to networks, can be modeled and solved as shortest path problems. Consequently, a variety of algorithms have been proposed to solve shortest path problems efficiently [1], and several benchmark computational studies [6, 7, 10, 11, 12, 14] have delineated circumstances under which certain algorithms might be preferred in practice.

Based on the number of source nodes and the number of sink nodes, shortest path problems are classified into those of finding: (1) a shortest path from node s to node t, (2) shortest paths from node s to all other nodes, (3) shortest paths to node t from all other nodes, and (4) shortest paths from every node to every other node. In the literature, problem type (2), which forms the basis of most shortest path algorithms, is referred to as the *single-source shortest path problem.*

Labeling algorithms arise as an iterative approach for solving single-source shortest path problems. These algorithms assign tentative distance labels to nodes at each step; the distance labels are estimates of the shortest path distances (namely, upper bounds). The different labeling approaches vary in how they update the distance labels from step to step and how they converge to the shortest path distances. Gilsinn and Witzgall [7] classified labeling algorithms for computing shortest paths into two general classes: *label-correcting* and *label-setting* algorithms. The typical label-correcting method starts with any tree T rooted at the source node s and updates T until no further improvement is possible. Label-setting methods begin with T consisting of the source node s and successively augment T by one node so that at each step T forms a shortest path tree (rooted at s) relative to the subgraph induced by T. Label-setting methods terminate once all nodes accessible from s are in T. This algorithm is applicable only to shortest path problems defined on acyclic networks with arbitrary arc lengths, or to shortest path problems with nonnegative arc lengths. On the other hand, label-correcting algorithms are more general and apply to all classes of problems, including those with negative arc lengths (but no negative length cycles). However shortest path distances are not revealed successively (as in label-setting methods); rather shortest path distances are known only upon termination of the algorithm.

In the following we concentrate on label-correcting algorithms applied to single-source shortest path problems. Our objective is to understand better the empirical behavior of label-correcting algorithms, taking into account their theoretical worst-case complexity as well as other intrinsic characteristics, such as *sharpness* [8, 10, 17] and *node maturity* (defined in Section 3.1). This emphasis on exploring certain features that contribute to algorithm performance distinguishes our work from previous computational studies.

1.1 Label-Correcting Algorithms

Consider a directed network $G = (N, A)$ with node set N and arc set A. Here $n = |N|$ and $m = |A|$, respectively. Let c_{ij} denote the arc length (cost, traversal time) of $(i, j) \in A$. The length of a path is the sum of the arc lengths along the path. The set of arcs emanating from node i is the *adjacency list* $A(i)$ of node i; namely, $A(i) = \{(i, j) \in A : j \in N\}$. Assume the network contains no directed cycles with negative length. For $s \in N$, a *shortest path tree* from source node s is a directed tree such that the unique path in the tree

from node s to any other node (accessible from s) is a shortest path between these nodes in the network G.

The generic *label-correcting algorithm* maintains a set of distance labels $D(\cdot)$ at every step. The label $D(j)$ is either ∞, indicating that we have yet to discover a directed path from the source to node j, or it is the length of some directed path from the source to node j. For each node j we also maintain a predecessor $pred(j)$, which records the node prior to node j in the current directed path to j. At termination, the predecessor array allows us to reconstruct a shortest path tree. The generic label-correcting algorithm involves successively updating the distance labels until they satisfy the shortest path optimality conditions [3]:

$$\begin{aligned} &D(j) \leq D(i) + c_{ij}, \quad \forall (i,j) \in A, \\ &D(s) = 0. \end{aligned} \tag{1}$$

One nice feature of the generic label-correcting algorithm (see Figure 1 below) is its flexibility: we can select arcs violating the optimality conditions (1) in any order and still assure finite convergence of the algorithm. However, without any further restriction on the choice of arcs in the generic label-correcting algorithm, the algorithm does not necessarily run in polynomial time. To obtain polynomially bounded label-correcting algorithms, the computations need to be organized carefully. A useful way is to do this is by maintaining a *candidate list* of nodes [7], initialized with $LIST = \{s\}$. At each iteration, the algorithm will add and/or delete nodes from $LIST$; it terminates when $LIST$ is empty. The typical iteration (assuming that $LIST$ is nonempty) first removes a node i to be *scanned* from $LIST$. The scanning of node i entails checking, for each arc in the adjacency list $A(i)$, the shortest path optimality conditions (1). If arc $(i,j) \in A(i)$ violates the optimality conditions, we set $D(j) := D(i) + c_{ij}$, $pred(j) := i$, and add j to $LIST$ if it does not already belong to $LIST$. In this modified label-correcting algorithm, shown in Figure 2, the procedure insert(j, $LIST$) adds node j to $LIST$.

```
algorithm generic label-correcting
    begin
        D(s) := 0; pred(s) := 0;
        D(j) := ∞, ∀j ∈ N − {s};
        while some arc (i, j) satisfies D(j) > D(i) + c_ij do
            D(j) := D(i) + c_ij; pred(j) := i;
        end while
    end
```

Fig. 1. Generic label-correcting algorithm.

```
algorithm modifed label-correcting
   begin
      D(s) := 0; pred(s) := 0; LIST := {s};
      D(j) := ∞, ∀j ∈ N − {s};
      while LIST ≠ ∅ do
         remove element i from LIST;
         for each (i, j) ∈ A(i) do
            if D(j) > D(i) + c_ij then
               D(j) := D(i) + c_ij; pred(j) := i;
               if j ∉ LIST then insert(j, LIST);
            end if
         end for
      end while
   end
```

Fig. 2. Modified label-correcting algorithm.

1.2 List Processing Disciplines

Different methods for adding and deleting nodes from *LIST* yield different shortest path algorithms. Label-correcting methods typically use some type of linear data structure (queue, stack, two-way list, partitioned queue) to maintain the candidate list *LIST*. In general, algorithms select the node to scan from the head of *LIST*. However they mainly differ in the manner of entering nodes into *LIST*. We now briefly describe some standard label-correcting algorithms.

- Bellman-Ford-Moore Algorithm [3]: *LIST* is implemented as a queue so that insertion occurs at the tail of *LIST* while removal occurs at the head of *LIST*, thus realizing a FIFO (First-In-First-Out) policy in the selection order.

- Gilsinn and Witzgall Algorithm [7]: This algorithm implements a LIFO (Last-In-First-Out) strategy; i.e., nodes are always added to the head of *LIST*.

- D'Esopo-Pape Algorithm [15]: In this algorithm, *LIST* has two insertion points. A node that enters *LIST* for the first time is added at the tail; a node that re-enters *LIST* is added at the head. The insertion routine for this algorithm (PAPE) is specified in Figure 3.

- Predecessor Algorithm: In this proposed algorithm, *LIST* also has two insertion points. However, the criterion to determine whether a node is

inserted at the tail or at the head of the list is based on the status of its children in the current tree, defined by *pred*. If the entering node is the immediate predecessor (parent) of a node currently in $LIST$, the new node is added to the head; otherwise, it is added to the tail. A justification for this strategy is that a node j is only added to $LIST$ after its label has been improved (decreased); therefore, its children's labels can be improved by at least the same amount. Consequently, it is desirable to add node j to $LIST$ ahead of its children, so that it will be scanned before its children. Otherwise, the (improvable) labels of the children will be used, likely to be corrected later on by an update from node j. The insertion routine for this algorithm (PRED) is shown in Figure 4.

```
algorithm insert(j, LIST)
    begin
        if j has been on LIST then
            add j to the head of LIST;
        else
            add j to the tail of LIST;
        end if
    end
```

Fig. 3. Insertion routine for Pape's label-correcting algorithm.

```
algorithm insert(j, LIST)
    begin
        if j has children on LIST then
            add j to the head of LIST;
        else
            add j to the tail of LIST;
        end if
    end
```

Fig. 4. Insertion routine for the predecessor label-correcting algorithm.

- Partitioning Shortest path (PSP) Algorithm [9]: This algorithm partitions the candidate list $LIST$ into two parts, NOW and $NEXT$, initialized with $NOW = \{s\}$ and $NEXT = \emptyset$. Nodes entering $LIST$ are inserted into $NEXT$. Nodes selected for scanning are removed from the head of

NOW if NOW is nonempty; otherwise, one pass of the algorithm has been completed and all nodes in $NEXT$ are transferred to NOW. The shortest path problem is solved when both lists are empty. The generic version of this algorithm (Figure 5) does not specify rules for entering nodes into $NEXT$. Thus any of the previously described variants (FIFO, LIFO, PAPE or PRED) can be used for inserting nodes into $NEXT$. These variants will be studied in subsequent sections of this paper.

```
algorithm psp
    begin
        D(s) := 0; pred(s) := 0; NOW := {s}; NEXT := ∅;
        D(j) := ∞, ∀j ∈ N − {s};
        while NOW ≠ ∅ do
            remove element i from NOW;
            for each (i, j) ∈ A(i) do
                if D(j) > D(i) + c_ij then
                    D(j) := D(i) + c_ij; pred(j) := i;
                    if j ∉ NOW ∪ NEXT then
                        insert(j, NEXT);
                    end if
                end if
            end for
            if NOW = ∅ then
                NOW := NEXT; NEXT := ∅;
            end if
        end while
    end
```

Fig. 5. Partitioning shortest path algorithm.

2 Algorithm Characteristics

The computational complexity of algorithms has been traditionally evaluated in the literature using two distinct approaches: worst-case analysis and empirical complexity. The present section will review the worst-case complexity of the various algorithms, expressed as a function of the problem size, generally measured by n, m and $\log C$, where $C = \max\{|c_{ij}| : (i,j) \in A\}$. We will also evaluate the algorithms in terms of their *sharpness*, a theoretical construct that appears to be relevant to understanding the empirical behavior of shortest path algorithms. Section 3 will be dedicated entirely to computational testing performed on the various algorithms.

2.1 Worst-Case Complexity

We summarize complexity results for the algorithms studied here, presenting first those algorithms with (pseudo) polynomial time complexity followed by those with exponential time complexity. The only new complexity result involves algorithm PRED, so it is discussed at greater length below.

Generic Label-Correcting Algorithm
The generic label-correcting algorithm has pseudopolynomial time complexity $O(n^2mC)$. To see this [1], notice that under the assumptions that all arcs have integral length and that there are no negative length cycles, the temporary labels $D(j)$ are finite and satisfy the bounds $-nC \leq D(j) \leq nC$. At each step, $D(j)$ decreases by an integral amount (at least one) for some j; so at most $n(2nC) = 2n^2C$ updates are required before termination occurs. Since it may take $O(m)$ work to find an arc violating the optimality conditions (1), this algorithm requires $O(n^2mC)$ operations in the worst case.

Bellman-Ford-Moore (FIFO) Algorithm
This algorithm [3] runs in $O(mn)$ time. The whole sequence of scanning operations performed during the algorithm can be grouped into a sequence of *passes*, where the first pass consists of scanning the source node s, while pass k consists of scanning the nodes that have been added to $LIST$ during pass $k-1$. Thus, a node can be scanned at most once during each pass, and at most $n-1$ passes are necessary (no shortest path requires more than $n-1$ arcs). Since a single pass requires $O(m)$ time, the overall complexity of the algorithm is $O(mn)$.

Partitioning Shortest path (PSP) Algorithm
The partitioning shortest path algorithm [9] also has computational complexity $O(mn)$, assuming an $O(1)$ complexity for insert(j, $NEXT$). Since the maximum number of arcs required in a shortest path is $n-1$, the method will require at most $n-1$ passes during which the candidate nodes are repartitioned from $NEXT$ to NOW. The total number of arcs examined during each pass is at most m. Therefore, the overall complexity of the algorithm is $O(mn)$.

Gilsinn and Witzgall (LIFO) Algorithm
This algorithm always adds nodes to the head of $LIST$, resulting in a depth-first exploration of the network. Shier and Witzgall [17] constructed a family of networks on which LIFO displays exponential behavior, requiring $O(2^n)$ node scans.

D'Esopo-Pape (PAPE) Algorithm
Just as in the $LIFO$ algorithm, the eventual stack nature of $LIST$ in the D'Esopo-Pape algorithm results in exponential worst-case time complexity. See [13, 17].

Predecessor (PRED) Algorithm
We now show that PRED has exponential complexity by considering the class of networks G_k, defined in the following way. Begin with a complete acyclic network on nodes $1, 2, \ldots, k$, where arc (i, j) exists only for $i > j$ and has length -2^{i-2}. Add a source node s connected to node $i = 1, 2, \ldots, k$ by an arc of length -2^{i-2}. Finally, add nodes $k+1, k+2, \ldots, 2k$ and zero length arcs of the form $(i, k+i)$ for $i = 1, 2, \ldots, k$. Figure 6 shows the constructed network for $k = 3$. Notice that G_{k+1} is formed recursively from G_k by adding two new nodes and relabeling appropriately. Let f_k be the number of node scans needed by PRED on G_k, using the convention that arcs $(s, j) \in A(s)$ are processed by increasing j and arcs $(i, j) \in A(i)$ are processed by decreasing j. Then we have the following recursion.

$$f_k = 2f_{k-1} - (k-2), \quad k \geq 1. \tag{2}$$

To see this result, notice that the execution of PRED on G_k entails the following steps:

1. scan s, placing $1, 2, \ldots, k$ on $LIST$;
2. process nodes $1, 2, \ldots, k-1$ as in G_{k-1};
3. scan k, thus placing $1, 2, \ldots, k-1$ back on $LIST$;
4. process nodes $1, 2, \ldots, k-1$ as in G_{k-1};
5. scan nodes $k+1, k+2, \ldots, 2k$.

Steps 1, 3 and 5 require a total of $2+k$ node scans, while Steps 2 and 4 each require $f_{k-1} - k$ node scans (since nodes $s, k+1, k+2, \ldots, 2k-1$ are not scanned during these steps). Altogether this gives $f_k = 2 + k + 2(f_{k-1} - k) = 2f_{k-1} - (k-2)$. Since $f_1 = 3$, solution of this recurrence gives $f_k = 2^k + k$ for $k \geq 1$, demonstrating the exponential behavior of PRED on networks G_k.

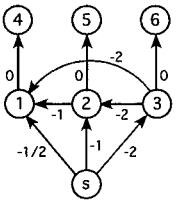


Fig. 6. Network G_3.

2.2 Sharp Labelings

We now discuss the concept of sharpness, introduced in [17]. A label-correcting algorithm maintains a predecessor graph (subtree) at every step, defined by

the array *pred*. We say that a distance label $D(i)$ is *sharp* if it equals the length of the unique path from node s to node i in the predecessor graph. An algorithm is *sharp* if every node scanned by the algorithm has a sharp distance label. (A sharp algorithm might produce nodes with nonsharp labels, but the algorithm never ends up scanning nodes with nonsharp labels.) Sharpness is a desirable characteristic in any label-correcting algorithm because it guarantees that when a node i is scanned there are no paths to node i in the predecessor graph that would currently improve its label. Consequently node i has a "good" label that can be used to update the labels of other nodes. Examples of sharp algorithms are PAPE and LIFO. On the other hand, FIFO and PSP are nonsharp; see [10, 17]. Algorithm PRED is also nonsharp since it is possible to scan a node j while it still has a (great) grandparent i on $LIST$. However, we term PRED as *locally sharp* since children will never be scanned before their parents using this list discipline.

3 Empirical Testing

The goal of empirical testing is to estimate how algorithms will behave on various classes of problem instances. Accordingly, our study involved the following steps: (1) constructing a random network generator to generate random problem instances for selected combinations of input size parameters (n, m giving the density $\delta = \frac{m}{n}$), (2) executing the computer programs written for each algorithm, and (3) collecting relevant performance data. Since our emphasis is on explaining the behavior of the various label-correcting algorithms in terms of *representative operation counts*, advocated by Ahuja and Orlin [2], we coded all algorithms in MATLAB. This provided a unified (and reproducible) framework for encoding the algorithms in a standard fashion, and it enabled us to eliminate variation due to coding abilities, compiler, language and timing routines.

The overall objectives for our computational study can be summarized as follows:

- Gain insight into the behavior of the different algorithms based on reproducible measures of computational effort.
- Identify possible improvements to the basic label-correcting algorithm.
- Understand the role played by sharpness, local sharpness and other intrinsic characteristics in the performance of shortest path algorithms.
- Compare the one-list with the two-list label-correcting algorithms.

In order to make valid comparisons, we insist on generating random directed networks in which every node is accessible from the source node s. Otherwise, a varying number of nodes will maintain infinite labels and will never be placed on $LIST$, making comparisons less reliable for generated networks of the same size. Connectedness becomes even more important with larger networks, since as the parameter n grows and δ stays fixed, it becomes

less and less likely that a randomly generated network will be connected. Consequently, we followed the random walk methodology proposed by Broder [5] by first generating a random spanning tree rooted at the source node s. Additional arcs were then randomly added to create a (simple) network on n nodes and m arcs, having no loops or repeated arcs.

In our computational study, we generated such s-connected random networks with $n = 100, 150, 200, 250, 500, 750, 1000$ nodes having various network densities $\delta = 3, 4, 5, 10, 15, 20, 25, 30$. In all the test networks generated, arc lengths were uniformly generated over the range $[1, 1000]$. All algorithms described earlier were executed to calculate a shortest path tree rooted at the source node $s = 1$ for each generated network. For each combination of size n and density δ, thirty random networks were generated; in all subsequent results, we report averages over these thirty replications for various measures of computational effort.

The performance of each algorithm was evaluated by computing the number of times the algorithm executed certain fundamental operations that we identified as indicative of algorithm performance. In particular, we collected data on the following measures of computational effort:

- **Node Scans:** the number of times a node i is removed from $LIST$. Namely, it is the total "length" of the candidate list.
- **Comparisons:** the number of times an arc (i, j) is examined for violation of the optimality conditions (1). This number is expected to be approximately the number of node scans times the average number of arcs per node (density).
- **Updates:** the number of times a node label $D(j)$ is decreased. This represents the number of updates performed until no further improvement is possible.

3.1 Results for One-List Label-Correcting Algorithms

We first present results obtained for random networks with 200 nodes in Tables 1–3. Similar data for networks of all sizes $100 \leq n \leq 1000$ can be found at the website [16]. As expected the number of comparisons (Table 2) is approximately the number of node scans (Table 1) times the network density. This relationship was observed in all of our test data instances. Therefore we omit this computational measurement in subsequent discussions.

Clearly for every network density, LIFO is by far the worst algorithm in terms of the measures shown in Tables 1 and 3. Moreover, LIFO continues to be the most inefficient algorithm for all combinations of numbers of nodes and densities. Contrary to the behavior observed for FIFO, PAPE and PRED, in which the number of node scans (and updates) increases with network density, the number of node scans (and updates) required by LIFO is only slightly affected by the density, particularly for $\delta \geq 5$.

Table 1. Number of node scans ($n = 200$).

Density	LIFO	FIFO	PAPE	PRED
3	4683.3	286.7	248.9	253.1
4	5915.5	305.4	275.7	276.5
5	6730.1	318.2	284.0	284.3
10	7237.2	385.3	389.8	350.5
15	7156.3	407.9	429.5	377.7
20	6631.1	437.6	491.7	403.0
25	7094.0	444.6	512.8	416.8
30	6482.9	445.8	526.5	416.3

Table 2. Number of comparisons ($n = 200$).

Density	LIFO	FIFO	PAPE	PRED
3	13967.0	858.9	744.8	757.9
4	23662.0	1213.5	1098.2	1104.6
5	33720.0	1595.3	1422.7	1429.6
10	72361.0	3844.2	3892.2	3506.8
15	107480.0	6107.0	6449.0	5680.6
20	132480.0	8741.7	9812.0	8054.5
25	177390.0	11132.0	12856.0	10454.0
30	194850.0	13371.0	15767.0	12502.0

Table 3. Number of updates ($n = 200$).

Density	LIFO	FIFO	PAPE	PRED
3	4689.7	340.5	325.9	325.3
4	5927.7	397.7	397.5	394.3
5	6757.0	439.3	437.3	435.9
10	7348.4	660.3	720.7	685.8
15	7378.9	772.2	859.6	823.4
20	6964.1	869.0	991.7	922.9
25	7587.0	950.0	1080.1	1014.5
30	7110.0	1001.7	1161.7	1086.0

Consistent with previous computational studies [4, 6, 7, 14], we observe from Table 1 that PAPE is more efficient than FIFO (in terms of the number of node scans) for sparse networks whereas FIFO is superior to PAPE for networks with higher densities. These observations are substantiated by carrying out paired-sample t tests on the data supporting Table 1. Namely, PAPE is significantly better than FIFO at densities $\delta \leq 5$, while FIFO is significantly better than PAPE at densities $\delta \geq 15$, in both cases well beyond the 1% significance level. There is no statistically significant difference between FIFO and PAPE at density $\delta = 10$. These results, illustrated for $n = 200$, hold

for the entire range of test problems ($100 \leq n \leq 1000$), with the breakpoint consistently occurring around $\delta = 10$.

PAPE's superior behavior (compared to FIFO) in sparse networks can be explained as follows. The first time a node j is given a label, the chance of that label being correct is very low; as a result, node j is added to the tail of *LIST* where its label can continue to be updated before it is used. The second time a node gets added to *LIST* (not necessarily the second time its label is updated), the chance of having a "good" label is considerably higher, and consequently the node is added to the head of *LIST* under PAPE so it can be scanned immediately. This strategy of PAPE is advantageous for sparse networks (say $\delta \leq 5$) where a node is scanned relatively few times ($1.2 - 1.7$) on average for $100 \leq n \leq 1000$, and thus nodes are unlikely to be added to the head of *LIST* a second time. On the other hand, the number of times PAPE adds a node to *LIST* increases for denser networks; in such networks the stack characteristic of *LIST* negatively affects its performance since inaccurate labels are being used right away to update other node labels. For example, when $500 \leq n \leq 1000$ our computational results show that PAPE will scan a node on average three times at higher densities. Thus for dense networks, PAPE will frequently add a node to the head of LIST a second time.

Also, Table 1 shows that PAPE entails fewer node scans than PRED for the sparsest networks, yet as the network density increases, PRED shows better and better performance. Formal statistical paired-sample t tests validate these conclusions. Specifically, at $n = 200$ nodes, PAPE is better than PRED for $\delta = 3$, while PRED is better than PAPE for $\delta \geq 10$. There is no statistically significant difference between PAPE and PRED at densities $\delta = 4$ and $\delta = 5$. Similar conclusions hold for the entire range of sizes $100 \leq n \leq 1000$. In addition, PRED is consistently (and statistically) better than FIFO at all network densities.

Several statistical models were considered to relate the number of node scans S to the network size; rather than using n, δ as independent variables, we used instead the number of nodes n and number of arcs m in our analyses. The best model found for representing the observed relationships was a power law model: $S = \alpha n^{\beta} m^{\gamma}$. Table 4 shows the estimated coefficients obtained when these models were fit using linear regression (after a logarithmic transformation of the power law model). In fact, by using dummy variables for the categories of algorithm type (FIFO, PAPE, PRED), only a single regression model needed to be run; it produced an excellent fit, judged by an examination of the residuals, with $R^2 = 0.998$. The $\hat{\gamma}$ coefficients in Table 4 indicate that as the density increases for n fixed, FIFO (with a smaller $\hat{\gamma}$) performs better than PAPE (with a much larger $\hat{\gamma}$). The regression coefficients $\hat{\alpha}, \hat{\beta}, \hat{\gamma}$ for PRED are intermediate between those for FIFO and PAPE, indicating that PRED represents a useful compromise between FIFO and PAPE.

We have also recorded (in Table 5) the number of times a locally nonsharp node is scanned by FIFO. A node j is said to be *locally nonsharp* when

Table 4. Regression models for one-list algorithms.

Algorithm	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\gamma}$
FIFO	0.7084	0.8887	0.2052
PAPE	0.4916	0.7508	0.3504
PRED	0.6366	0.8577	0.2285

its distance label $D(j)$ does not equal $D(i) + c_{ij}$, where $i = pred(j)$ is the predecessor of node j in the current tree. Notice that the values in Table 5 are only slightly affected by the network density. This surprising constancy in the number of locally nonsharp scans could also explain the difference in performance between PRED (which is locally sharp) and FIFO, and could also explain why FIFO is better than PAPE for dense networks, since the proportion of time FIFO scans a locally nonsharp node decreases with density.

Table 5. Number of times FIFO scans a locally nonsharp node ($n = 200$).

Density	3	4	5	10	15	20	25	30
# scans	30.33	28.47	31.13	39.10	41.67	42.87	44.43	43.67

In Table 6, we display the number of updates per node scanned for networks with 200 nodes. This ratio gives an idea of the *maturity* of a node when it becomes scanned. Namely, it measures the average number of times a node is updated while resident on *LIST*. In LIFO the labels rarely get updated before they are used, while in PRED a node label gets updated more often before the node is scanned, especially for high densities. A high maturity value increases the chance that a node label is in fact the shortest distance from s by the time it is used; hence, the higher the number of updates per scan, the better the label. Comparing Tables 1 and 6, we see that at each density there is a high negative correlation between the number of node scans (measurement of algorithm efficiency) and the number of updates per scan. In other words, the (average) maturity of a node appears to be an excellent predictor of algorithm efficiency.

Table 6 shows that maturity increases with network density, an observation that can be easily explained. Namely, as the density increases the number of arcs entering a given node k on *LIST* increases, as does the chance that $D(k)$ will be updated while on *LIST*. On the other hand, node maturity is surprisingly stable as n increases. For example, this ratio (updates per node scan) for FIFO ranged from 1.2 ($\delta = 3$) to 2.3 ($\delta = 30$) at $n = 100$; the corresponding range was $[1.2, 2.2]$ at $n = 1000$. Similarly, the ranges for PAPE and PRED were remarkably stable at $[1.3, 2.2]$ and $[1.3, 2.7]$ over the range $100 \leq n \leq 1000$. LIFO showed the lowest ratios $[1.0, 1.1]$, consistent with the algorithm's feature of scanning a node right after it is updated. On the other

Table 6. Number of updates per scan ($n = 200$).

Density	LIFO	FIFO	PAPE	PRED
3	1.00	1.19	1.31	1.28
4	1.00	1.30	1.44	1.43
5	1.00	1.38	1.54	1.53
10	1.02	1.71	1.85	1.96
15	1.03	1.89	2.00	2.18
20	1.05	1.99	2.02	2.29
25	1.07	2.14	2.11	2.43
30	1.10	2.25	2.21	2.61

hand, the highest values ($2.3 - 2.7$) were achieved by PRED, for the densest networks, independent of network size.

In summary, our computational study has verified that PAPE outperforms FIFO for sparse networks, as noted in previous studies [4, 6, 7, 14]. Here we provide several explanations for this phenomenon. The computational overhead of scanning locally nonsharp nodes puts FIFO at a disadvantage compared to the sharp algorithm PAPE. However, as the network density increases this essentially fixed overhead becomes negligible and the relative performance of FIFO improves. Moreover, for higher densities the number of node scans per node is larger for PAPE compared to FIFO; since nodes re-enter $LIST$ at the head for PAPE, these subsequent node scans serve to propagate out-of-date information and slow down PAPE. We also observe that PRED is generally better than PAPE, especially as the network density increases. We argue that this may be a result of the larger maturity values achieved by PRED.

3.2 Results for Two-List Label-Correcting Algorithms

In the Partitioning Shortest path (PSP) algorithm, nodes to be scanned are always removed from the head of NOW, but no specific rules are required for entering nodes into $NEXT$. Glover et al. [10] have investigated six variants of the PSP algorithm, using the concept of *threshold value*. By contrast, our implementation uses FIFO, LIFO, PAPE and PRED to drive the addition of nodes into $NEXT$. Respectively, we call each of the resulting algorithms FIFO+FIFO, FIFO+LIFO, FIFO+PAPE and FIFO+PRED. Since FIFO is not affected by partitioning the candidate list, the results for FIFO+FIFO duplicate those for the one-list FIFO version.

As a starting point, we present in Tables 7–9 results for networks having $n = 200$ nodes. Similar data for networks of all sizes $100 \leq n \leq 1000$ can be found at the website [16]. Recall that LIFO goes from exponential complexity to polynomial complexity with FIFO+LIFO, and this is clearly reflected in the dramatically reduced number of node scans (Table 7). In fact, the two-list version FIFO+LIFO now clearly outperforms FIFO for densities $\delta < 20$ and is generally comparable for higher densities. This phenomenon persists

for larger networks as well. The two-list FIFO+PAPE slightly degrades the performance of the one-list PAPE at lower densities ($\delta \leq 5$), but significantly improves PAPE at higher densities. As a result, FIFO+PAPE now consistently outperforms FIFO at all densities; this result continues to hold throughout the range of network sizes $100 \leq n \leq 1000$.

Table 7. Number of node scans ($n = 200$).

Density	FIFO+FIFO	FIFO+LIFO	FIFO+PAPE	FIFO+PRED
3	286.7	263.7	265.1	265.5
4	305.4	290.4	286.8	286.9
5	318.2	302.1	298.3	297.7
10	385.3	378.2	366.3	362.5
15	407.9	403.0	386.4	383.9
20	437.6	438.5	421.2	411.2
25	444.6	445.4	429.0	424.4
30	445.8	438.6	427.8	423.2

The number of node scans per node (Table 8) increases with network density, ranging here from $1.3-2.2$ for all two-list algorithms except FIFO+FIFO =FIFO. In general, the number of node scans per node increases with n and typically lies in the range $[1.2, 2.6]$, implying that on average a node gets scanned at most 2.6 times even at the highest densities. In particular, for larger networks ($n \geq 500$) FIFO+PAPE reduces the average number of node scans of PAPE from $[1.3, 3.2]$ to $[1.4, 2.6]$, clearly a benefit in dense networks where the one-list PAPE is less efficient.

Table 8. Number of node scans per node ($n = 200$).

Density	FIFO+FIFO	FIFO+LIFO	FIFO+PAPE	FIFO+PRED
3	1.43	1.32	1.33	1.33
4	1.53	1.45	1.43	1.43
5	1.59	1.51	1.49	1.49
10	1.93	1.89	1.83	1.81
15	2.04	2.02	1.93	1.92
20	2.19	2.19	2.11	2.06
25	2.22	2.23	2.14	2.12
30	2.23	2.19	2.14	2.12

All partitioning shortest path variants are polynomial but are nonsharp. Even PRED, which guarantees local sharpness, may become locally nonsharp as a result of partitioning the candidate list. Comparing Table 1 for PRED with Table 7 for FIFO+PRED shows that partitioning degrades performance

at all densities, but especially at lower densities. However, as seen in Table 7 the partitioned version of PRED still dominates the other partitioned versions for all but the sparsest networks. Moreover, as the network size increases ($n \geq 500$), FIFO+PRED emerges as the best two-list algorithm at *all* densities.

We also carried out a regression analysis to summarize the behavior of the four two-list algorithms FIFO+FIFO, FIFO+LIFO, FIFO+PAPE and FIFO+PRED. Again we found that the power law model $S = \alpha n^{\beta} m^{\gamma}$ provided a very good fit to the data and a single regression model (with dummy variables for the different algorithms) was run, resulting in an R^2 value of 0.999. The estimated regression coefficients are shown in Table 10. Notice that FIFO+FIFO achieves the smallest value of $\hat{\gamma}$ followed by FIFO+PRED, indicating that for dense networks these two algorithms should perform well. Since the coefficients $\hat{\alpha}$ and $\hat{\beta}$ for FIFO+PRED are uniformly smaller than those for FIFO+FIFO, we anticipate that FIFO+PRED should dominate in practice (as we have previously observed). By comparing Table 4 and Table 10, we see that partitioning has the effect of substantially reducing the coefficient $\hat{\gamma}$ for PAPE, again indicating the improvement afforded by FIFO+PAPE in denser networks.

Table 9 displays the maturity values (number of updates per node scan) for $n = 200$. Recall that large maturity values appear to be desirable. Partitioning has the effect of drastically increasing the maturity of FIFO+LIFO compared to LIFO (Table 6), whereas the maturity values of PRED are decreased. In the case of FIFO+PAPE, the two-list version has smaller maturity values than PAPE for lower densities but larger maturity values for higher densities (where the standard PAPE algorithm is less competitive). In general, over the range $100 \leq n \leq 1000$, the breakpoint for PAPE occurs around $\delta = 15$. The largest maturity values are typically achieved by FIFO+PRED, especially at higher densities. At lower densities, comparable maturity values are achieved by FIFO+LIFO, FIFO+PAPE and FIFO+PRED.

Table 9. Number of updates per scan ($n = 200$).

Density	FIFO+FIFO	FIFO+LIFO	FIFO+PAPE	FIFO+PRED
3	1.19	1.26	1.24	1.24
4	1.30	1.36	1.36	1.36
5	1.38	1.46	1.45	1.45
10	1.71	1.77	1.80	1.81
15	1.89	1.98	1.99	2.01
20	1.99	2.06	2.06	2.09
25	2.14	2.18	2.22	2.24
30	2.25	2.28	2.34	2.37

Notably, our computational study of two-list algorithms shows that partitioning *dramatically* improves the performance of LIFO; indeed, it now outper-

Table 10. Regression models for two-list algorithms.

Algorithm	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\gamma}$
FIFO+FIFO	0.7084	0.8887	0.2052
FIFO+LIFO	0.6683	0.8569	0.2309
FIFO+PAPE	0.6585	0.8679	0.2223
FIFO+PRED	0.6738	0.8744	0.2136

forms FIFO, especially at lower densities. Also, partitioning enhances PAPE at higher densities, precisely where the standard implementation of PAPE needs improvement compared to FIFO. Overall, PRED remains the best two-list algorithm.

4 Conclusions

In this paper we have considered several simple variants of the basic label-correcting algorithm for solving single-source shortest path problems. Our objective was not to identify the most efficient label-correcting algorithm and implementation. Rather, this study has concentrated on fairly standard one-list disciplines (FIFO, LIFO, PAPE), with the aim of obtaining *insights* into the relative effectiveness of these common list processing variants, using the ideas of sharpness and maturity. Based on the idea of maintaining local sharpness, a new variant PRED has been proposed. In addition, we studied two-list (partitioning) algorithms, implementing the earlier described strategies to handle the list $NEXT$. In our computational study, all algorithms were coded in MATLAB to provide a consistent programming framework.

Our empirical study was based on generating random directed networks with specified numbers of nodes and arcs. In contrast to many previous computational studies, we insist that all nodes be accessible from the source node s (which was ensured using a random walk methodology [5]). This requirement guarantees that all instances generated at a fixed network size will have comparable work to accomplish, namely connecting *all* nodes to the source node by a shortest path tree. We evaluated the algorithms' empirical behavior on networks with the number of nodes n ranging from 100 to 1000 and the density δ ranging from 3 to 30. Rather than focusing on CPU time, we collected data on representative operation counts — which are independent of computing platform and can better facilitate understanding the relative strengths of the algorithms.

The computational study verified that among the one-list label-correcting algorithms, PAPE was better than FIFO for sparse networks, while FIFO was better than PAPE for denser networks. We offered several explanations for this behavior, based on examining the average number of node scans as well as the average number of updates per node scan (maturity). We also observed a

surprising stability in the average maturity of node labels as n increases and a high (negative) correlation between maturity and algorithm performance. The PRED variant was superior to the other algorithms for all but the sparsest networks, an observation that we attribute to the large maturity values it produces. In a certain sense, PRED interpolates between FIFO and PAPE in an attempt to capture the benefits of both. Namely, rather than employing a very rigid criterion for entering nodes at the head of *LIST* (as done in PAPE), a more relaxed criterion is adopted by PRED that allows nodes to be entered at the tail of *LIST* (as done in FIFO) except when local nonsharpness would be compromised.

Partitioning the candidate list was in general beneficial for the LIFO and PAPE variants; it had no effect on FIFO and mildly degraded PRED. Most strikingly, LIFO was transformed from being an order of magnitude slower than all algorithms to being fairly competitive — indeed, it consistently outperformed FIFO on the larger networks. The two-list FIFO+PAPE was improved significantly at higher densities, precisely where the performance of PAPE becomes inferior to that of FIFO. Overall, FIFO+PRED (which has guaranteed polynomial complexity) emerged as the superior algorithm among the two-list variants.

Future research would need to expand the empirical study to randomly generated grid networks, which are both structured and sparse. Also we have purposely focused on fairly straightforward modifications of the basic label-correcting approach for this initial study. As a result we have observed the benefits of having local sharpness and high maturity values, as well as the effects of partitioning on these standard algorithms. Such considerations should prove useful in devising other, more sophisticated shortest path variants. Developing and testing production versions of such codes is yet another activity for future investigation.

Acknowledgments

We are appreciative of the constructive comments of the referee, which have helped to improve the presentation of this paper.

References

1. R.K. Ahuja, T.L. Magnanti, and J.B. Orlin. *Network Flows: Theory, Algorithms, and Applications.* Prentice Hall, Englewood Cliffs, NJ, 1993.
2. R.K. Ahuja and J.B. Orlin. Use of representative operation counts in computational testing of algorithms. *INFORMS Journal on Computing*, 8:318–330, 1996.
3. R. Bellman. On a routing problem. *Quarterly of Applied Mathematics*, 16:87–90, 1958.

4. D.P. Bertsekas. A simple and fast label correcting algorithm for shortest paths. *Networks*, 23:703–709, 1993.
5. A. Broder. Generating random spanning trees. *30th Annual Symposium on the Foundations of Computer Science*, 442–447, 1989.
6. R. Dial, F. Glover, D. Karney, and D. Klingman. A computational analysis of alternative algorithms and techniques for finding shortest path trees. *Networks*, 9:215–248, 1979.
7. J. Gilsinn and C. Witzgall. A performance comparison of labeling algorithms for calculating shortest path trees. *NBS Technical Note 772*, U.S. Department of Commerce, 1973.
8. F. Glover and D. Klingman. New sharpness properties, algorithms and complexity bounds for partitioning shortest path procedures. *Operations Research*, 37:542–546, 1989.
9. F. Glover, D. Klingman, and N. Phillips. A new polynomially bounded shortest path algorithm. *Operations Research*, 33:65–73, 1985.
10. F. Glover, D.D. Klingman, N.V. Phillips, and R.F. Schneider. New polynomially bounded shortest path algorithms and their computational attributes. *Management Science*, 31:1106–1128, 1985.
11. B.L. Hulme and J.A. Wisniewski. A comparison of shortest path algorithms applied to sparse graphs. *Sandia Technical Note 78-1411*, Sandia Laboratories, 1978.
12. M.S. Hung and J.J. Divoky. A computational study of efficient shortest path algorithms. *Computers and Operations Research*, 15:567-576, 1988.
13. A. Kershenbaum. A note on finding shortest path trees. *Networks*, 11:399–400, 1981.
14. J.-F. Mondou, T.G. Crainic, and S. Nguyen. Shortest path algorithms: A computational study with the C programming language. *Computers and Operations Research*, 18:767–786, 1991.
15. U. Pape. Implementation and efficiency of Moore-algorithms for the shortest path problem. *Mathematical Programming*, 7:212–222, 1974.
16. D.R. Shier and M.G. Bardossy. Computational results for one-list and two-list label-correcting shortest path algorithms, December 2005. Available at the URL http://www.math.clemson.edu/~shierd/Shier/.
17. D.R. Shier and C. Witzgall. Properties of labeling methods for determining shortest path trees. *Journal of Research of the National Bureau of Standards*, 86:317–330, 1981.

The Ubiquitous Farkas Lemma

Rakesh V. Vohra

Department of Managerial Economics and Decision Sciences
Kellogg Graduate School of Management
Northwestern University
Evanston, IL 60208
r-vohra@kellogg.northwestern.edu

Summary. Every student of linear programming is exposed to the Farkas lemma, either in its original form or as the duality theorem of linear programming. What most students don't realize is just how ubiquitous the lemma is. I've managed to make a good living by simply trying to express problems in the form of linear inequalities and then examining the Farkas alternative. It stared with my dissertation (under Saul's supervision) and continues even today. So I could think of no better gift on the occasion of Saul's birthday than a compilation of some applications of the Farkas lemma. These applications have been pulled together from a variety of different sources.

Key words: Farkas lemma, linear programming, duality.

1 Introduction

As every operations research student knows, associated with every $m \times n$ matrix A of real numbers is a problem of the following kind:

> Given $b \in \Re^m$ find an $x \in \Re^n$ such that $Ax = b$, or prove that no such x exists.

Convincing someone that $Ax = b$ has a solution (when it does) is easy. One merely exhibits the solution, and the other can verify that the solution does indeed satisfy the equations. What if the system $Ax = b$ does not admit a solution. Is there an easy way to convince another of this? Stating that one has checked all possible solutions convinces no one.

By framing the problem in the right way, one can apply the machinery of linear algebra. Specifically, given $b \in \Re^m$, the problem of finding an $x \in \Re^n$ such that $Ax = b$ can be stated as: is b in the span of the column vectors of A? This immediately yields the **Fundamental Theorem of Linear algebra** first proved by Gauss.

Theorem 1. *Let A be an $m \times n$ matrix, $b \in \Re^m$, and $F = \{x \in \Re^n : Ax = b\}$. Then* **either** $F \neq \emptyset$ **or** *there exists $y \in \Re^m$ such that $yA = 0$ and $yb \neq 0$,* **but not both.**

The proof is not hard. If $F \neq \emptyset$, we are done. So suppose that $F = \emptyset$. Then, b is not in the span of the columns of A. If we think of the span of the columns of A as a plane, then b is a vector pointing out of the plane. Thus, any vector y orthogonal to this plane (and so to every column of A) must have a non-zero dot product with b. To verify the 'not both' portion of the statement, suppose not. Then there is a x such that $Ax = b$ and a y such that $yA = 0$ and $yb \neq 0$. This implies that $yAx = yb$, a contradiction, since $yAx = 0$ and $yb \neq 0$.

The next level up in difficulty is the following question:

> Given $b \in \Re^m$, find a non-negative $x \in \Re^n$ such that $Ax = b$, or prove that no such x exists.

The requirement that x be non-negative is what causes all the difficulty. It is to Gulya Farkas we are indebted for providing us with a generalization of the fundamental theorem of Linear algebra.[1]

Theorem 2 (The Farkas Lemma). *Let A be an $m \times n$ matrix, $b \in \Re^m$, and $F = \{x \in \Re^n : Ax = b, x \geq 0\}$. Then* **either** $F \neq \emptyset$ **or** *there exists $y \in \Re^m$ such that $yA \geq 0$ and $y \cdot b < 0$* **but not both.**

I'll assume that the reader is familiar with the result and its variants.[2] Now to the applications.

2 The Duality Theorem of Linear Programming

A staple of many a homework set in a class on linear programming (LP) is to prove the Farkas lemma from LP duality or the reverse. Here is one example of how to get from the Farkas lemma to the duality theorem. I'll take the following as the form of the primal problem (P):

$$Z_P = \max\{cx : \text{ s.t. } Ax = b, x \geq 0\}.$$

The dual of course is (D):

$$Z_D = \min\{yb : yA \geq c\}.$$

[1] Gulya Farkas (1847-1930) was a Hungarian theoretical physicist. The lemma that bears his name was announced by him in 1894 and has its roots in the problem of specifying the equilibrium of a system. The associated proof was incomplete. A complete proof was published in Hungarian by Farkas in 1898. It is more common to refer to the German version, "Theorie der Einfachen Ungleichungen," which appeared in the *J. Reine Angew Math.*, 124, 1901, 1-27. For an exposition in English, see [3].

[2] However, I will note that the many proofs of the lemma via the separating hyperplane theorem 'cheat'. These proofs do not bother to show that the cone generated by the columns of A is a closed set.

Lemma 1. *If problem (P) is infeasible, then (D) is either infeasible or unbounded. If (D) is unbounded, then (P) is infeasible.*

Proof. Let (P) be infeasible, and suppose for a contradiction that (D) has a finite optimal solution, say y^*. Infeasibility of (P) implies by the Farkas lemma a vector $\hat{y}$ such that $\hat{y}A \geq 0$ and $\hat{y} \cdot b < 0$. Let $t > 0$. The vector $y^* + t\hat{y}$ is a feasible solution for (D), since $(y^* + t\hat{y})A \geq y^*A \geq c$. Its objective function value is $(y^* + t\hat{y}) \cdot b < y^* \cdot b$, contradicting the optimality of y^*. Since (D) cannot have a finite optimal, it must be infeasible or unbounded.

Now suppose (D) is unbounded. Then, we can write any solution of (D) as $y + r$, where y is a feasible solution to the dual (i.e., $yA \geq c$) and r is a ray (i.e., $rA \geq 0$). Furthermore $r \cdot b < 0$, since (D) is unbounded. By the Farkas lemma, the existence of r implies the primal is infeasible. □

Theorem 3 (Duality Theorem). *If a finite optimal solution for either the primal or dual exists, then $Z_P = Z_D$.*

Proof. By Lemma 1, if one of Z_P and Z_D is finite, so is the other. Let x^* be an optimal solution to the primal and y^* an optimal solution to the dual. By weak duality,

$$Z_D = y^* \cdot b = y^*Ax^* \geq c \cdot x^* = Z_P.$$

To complete the proof, we show that $Z_D \leq Z_P$. Pick an $\epsilon > 0$ and consider the set

$$\{x : -cx \leq Z_P, Ax = b, x \geq 0\}.$$

By the definition of Z_P, this set is empty. So, by the Farkas lemma, there is a solution to the following system:

$$\begin{aligned} -\lambda c + yA &\geq 0, \\ \lambda(-Z_P - \epsilon) + yb &< 0, \\ \lambda &\geq 0. \end{aligned}$$

Let that solution be (λ^*, y^*). We show that $\lambda^* > 0$. Suppose not. Since $\lambda^* \geq 0$, it follows that $\lambda^* = 0$. This implies that

$$\begin{aligned} y^*A &\geq 0, \\ y^*b &< 0. \end{aligned}$$

By the Farkas lemma, this implies that the system $Ax = b$, with $x \geq 0$, is infeasible, which violates the initial assumption.

Let $y' = y^*/\lambda^*$. Since $\lambda^* > 0$, this is well defined. Also

$$y'A = \frac{y^*A}{\lambda^*} \geq c,$$

making y' a feasible solution for the dual problem. Further $y'b < Z_P + \epsilon$. Since y' is feasible in the dual, it follows that

$$Z_P \leq Z_D \leq y'b < Z_P + \epsilon.$$

Since $\epsilon > 0$ is arbitrary, it follows that $Z_P = Z_D$. □

3 Arbitrage Free Pricing

The word **arbitrage** comes from the French *arbitrer* and means to trade in stocks in different markets to take advantage of different prices. The absence of arbitrage opportunities is the driving principle of financial theory.[3]

Suppose there are m assets, each of whose payoffs depends on a future state of nature. Let S be the set of possible future states of nature, with $n = |S|$. Let a_{ij} be the payoff from one share of asset i in state j. A portfolio of assets is represented by a vector $y \in \Re^m$, where the i^{th} component, y_i, represents the amount of asset i held. If $y_i > 0$, one holds a **long** position in asset i, whereas $y_i < 0$ implies a **short** position in asset i.[4] Let $w \in \Re^n$ be a vector whose j^{th} component denotes wealth in state $j \in S$. We assume that wealth (w) in a future state is related to the current portfolio (y) by

$$w = yA.$$

This assumes that assets are infinitely divisible, returns are linear in the quantities held, and the return of the asset is not affected by whether one holds a long or short position. Thus, if one can borrow from the bank at 5%, then one can lend to the bank at 5%.

The no arbitrage condition asserts that a portfolio that pays off non-negative amounts in every state must have a non-negative cost. If $p > 0$ is a vector of asset prices, we can state the no arbitrage condition algebraically as follows:

$$yA \geq 0 \Rightarrow y \cdot p \geq 0.$$

Equivalently, the system $yA \geq 0, y \cdot p < 0$ has no solution. From the Farkas lemma, we deduce the existence of a non-negative vector $\hat{\pi} \in \Re^m$ such that

$$p = A\hat{\pi}.$$

Since $p > 0$, it follows that $\hat{\pi} > 0$. Scale $\hat{\pi}$ by dividing through by $\sum_j \hat{\pi}_j$. Let $p^* = p/\sum_j \hat{\pi}_j$ and $\pi = \hat{\pi}/\sum_j \hat{\pi}_j$. Notice that π is a probability vector. As long as relative prices are all that matter, scaling the prices is of no relevance. After the scaling, $p^* = A\pi$. In words, there is a probability distribution under which the expected value of each security is equal to its buying/selling price. Such a distribution is called a **risk neutral probability distribution**. A risk-neutral investor using these probabilities would conclude that the securities are fairly priced.

[3] The exposition in this section is based in part on [4].

[4] To hold a long position is to acquire the asset in the hope that its value will increase. To hold a short position is to make a bet that the asset will decline in value. On the stock market this is done by selling a stock one does not own now and buying it back at a later date, presumably when its price is lower. In practice, one's broker will 'borrow' the stock from another client and sell it in the usual way. At some time after the sale, one tells the broker stop and buys the stock borrowed at the prevailing price and returns them to the 'lender'. The strategy yields a profit if the price of the stock goes down.

4 Markov Chains

A Markov chain is a system that can be in one of n states.[5] At each time period there is a chance of the system of moving from one state to another (or remaining in its current state). The probability that the system transitions from state j to state i is denoted a_{ij}. These probabilities are called transition probabilities. Call the $n \times n$ matrix of transition probabilities A. If p_k is the probability that the system is in state k at period t, say, then the probability that it will be in state i in the next period is $\sum_{j=1}^{n} p_j a_{ij}$. In matrix notation, this would be Ap. A probability vector x is called a **stationary distribution** if $Ax = x$. In words, the probability of being in any particular state remains unchanged from one period to the next.

The goal is to show that there is a non-negative x such that $(A - I)x = 0$ and $\sum_{j=1}^{n} x_j = 1$. In matrix notation:

$$\begin{pmatrix} A - I \\ e \end{pmatrix} x = \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

where e is an n-vector of all 1's. Let $B = \begin{pmatrix} A \\ e \end{pmatrix}$ and $b = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$. We seek a non-negative solution to $Bx = b$. Suppose such a solution does not exist. By the Farkas lemma, there is a $y = (z_1, \ldots, z_n, -\lambda)$ such that $yB \geq 0$ and $yb < 0$. Now

$$yB \geq 0 \Rightarrow z(A - I) - \lambda \cdot e \geq 0$$

and

$$yb < 0 \Rightarrow -\lambda < 0 \Rightarrow \lambda > 0.$$

Hence for all j,

$$\sum_i a_{ij} z_i - z_j \geq \lambda > 0.$$

Choose an i for which z_i is largest and call it k. Since $\sum_i a_{ij} = 1$, we have that $z_k \geq \sum_i a_{ij} z_i$. Hence

$$0 = z_k - z_k \geq \sum_i a_{ij} z_i - z_k \geq \lambda > 0,$$

a contradiction. Hence there must be a probability vector x such that $Ax = x$.

[5] Named in honor of Andrei Markov. Markov begat Voronoi (of the diagrams). Voronoi begat Sierpinski (of space filling curve fame). Sierpinski begat Neyman (of the Neyman-Pearson lemma). Neymann begat Dantzig who in turn begat Saul. For a more extensive discussion of Markov chains, see [2].

5 Exchange Economies

The standard stylized model of a competitive market, due to Leon Walras, makes the following assumptions:

1. A finite set A of m agents (consumers).
2. A finite number, n, of goods that are divisible.
3. The endowment of agent i is denoted $e^i \in \Re^n_+$ and her utility function by $U^i : \Re^n_+ \to \Re$. Each U^i is assumed continuous, concave and monotone.
4. Each agent is aware of the price of every good.
5. The transaction costs of a sale, purchase, etc., are zero.
6. Agents can buy and sell as much and as little as they want at the going price. Their transactions do not affect the price.(They are price-takers.)

Suppose a price p_j for each good $j = 1, \ldots, n$, is announced at which trade will take place. At those prices, agent i has a budget of $p \cdot e^i$. Suppose agents report the demands that maximize their utility subject to their budgets. Will the demands of the agents exceed the available resources? Is there are price vector $p \in \Re^n_+$ where the demands would balance the available supply? Remarkably, yes. Under the right prices, agents acting *independently* to maximize their utility will specify demands so that supply balances demand. These prices are called **equilibrium prices**. The equilibrium prices along with the resulting allocation of goods is called a **Walrasian Equilibrium** in honor of Leon Walras.

An allocation X is a collection of vectors $(x^1, x^2, \ldots, x^m)$, where x^i denotes the vector of goods consumed by agent i. An **equilibrium** is a price vector $p \in \Re^n_+$ such that

1. $x^i \in \Re^n_+$,
2. $\sum_{i \in A} x^i \leq \sum_{i \in A} e^i$,
3. $x^i \in \arg\max\{U^i(x) : px \leq pe^i, \quad x \geq 0\}$.

The first condition says that each agent consumes a non-negative amount of each good. The second condition says that the total amount of each good consumed cannot exceed the available supply. The last condition requires that the amount consumed by agent i maximize their utility. Under these conditions, proving the existence of an equilibrium requires the use of Kakutani's fixed point theorem. We consider a special case which will allow us to use the Farkas lemma to turn the problem of finding an equilibrium into a convex program.[6] Specifically, the utility of each agent is linear in x^i, i.e., $U^i(z) = \sum_{j=1}^n v_{ij} z_j$.

Given a price vector, agent i solves problem (P_i):

$$\max \sum_{j=1}^{n} v_{ij} x^i_j$$

[6] This result is due to [5].

$$\text{s.t.} \sum_{j=1}^{n} p_j x_j^i \leq p \cdot e^I,$$
$$x_j^i \geq 0 \ \forall j.$$

Here x_j^i is the amount of good j that agent i consumes. Let $x^i(p)$ be the optimal solution to this program. The price vector p is an equilibrium if

$$\sum_{i=1}^{m} x_j^i(p) \leq \sum_{i=1}^{m} e_j^i \ \forall j. \tag{1}$$

The dual to problem (P_i) is (D_i):

$$\min \lambda^i (p \cdot e^i)$$
$$\text{s.t. } \lambda^i p_j \geq v_{ij} \ \forall j,$$
$$\lambda^i \geq 0.$$

Appending dual to primal, and using the fact that their respective objective function values must coincide, we obtain the following system:

$$\sum_{j=1}^{n} p_j x_j^i \leq p \cdot e^i, \tag{2}$$
$$\sum_{j=1}^{n} v_{ij} x_j^i = \lambda^i p \cdot e^i, \tag{3}$$
$$\lambda p_j \geq v_{ij} \ \forall j, \tag{4}$$
$$x_j^i, \lambda^i \geq 0 \ \forall j. \tag{5}$$

We show that the budget constraint (2) is redundant, given the feasibility constraints (1) and $\lambda^i > 0$.

From (3) and (4) of (D_i), we have:

$$\lambda^i (p \cdot e^i) \leq \lambda^i \sum_{j=1}^{n} p_j x_j^i.$$

Therefore

$$\lambda^i \sum_{i=1}^{m} p \cdot e^i \leq \lambda^i \sum_{i=1}^{m} \sum_{j=1}^{n} p_j x_j^i$$
$$\Rightarrow \sum_{i=1}^{m} p \cdot e^i \leq \sum_{i=1}^{m} \sum_{j=1}^{n} p_j x_j^i = \sum_{i=1}^{m} p \cdot e^i.$$

Since this last inequlity holds at equality, the initial inequality must also hold at equality, i.e., $\sum_{j=1}^{n} p_j x_j^i = p \cdot e^i$.

Therefore, the problem of finding an equilibrium reduces to finding $x^1, x^2, \ldots, x^m$, $\lambda^1, \ldots, \lambda^m$, and p such that:

$$\sum_{j=1}^{n} v_{ij} x_j^i = \lambda^i p \cdot e^i \ \forall i, \tag{6}$$

$$\lambda^i p_j \geq v_{ij} \ \forall j, \tag{7}$$

$$\sum_{i=1}^{m} x_j^i = \sum_{i=1}^{m} e_j^i \ \forall j, \tag{8}$$

$$x_j^i, \lambda^i \geq 0 \ \forall i, j. \tag{9}$$

We can use (6) and (7) in the system above to eliminate the λ's :

$$\frac{\sum_{j=1}^{n} v_{ij} x_j^i}{p \cdot e^i} \geq \frac{v_{ij}}{p_j} \ \forall i, j, \tag{10}$$

$$\sum_{i=1}^{m} x_j^i = \sum_{i=1}^{m} e_j^i \ \forall j, \tag{11}$$

$$x_j^i \geq 0 \ \forall i, j. \tag{12}$$

It is convenient to introduce a normalization. Each agent i has an endowment of 1 unit of good i, i.e., one good per agent and no two agents have the same good. In this case $p \cdot e^i = p_i$. Then (10)-(12) reduces to

$$\frac{\sum_{j=1}^{n} v_{ij} x_j^i}{p_i} \geq \frac{v_{ij}}{p_j} \ \forall i, j,$$

$$\sum_{i=1}^{m} x_j^i = \sum_{i=1}^{m} e_j^i = 1 \ \forall j,$$

$$x_j^i \geq 0 \ \forall i, j.$$

Since p is not fixed, this system is non-linear. However, by taking logs, we can turn the problem of finding the equilibrium solution with price vector p and allocation $\{x^1, \ldots, x^m\}$ into the following convex program:

$$\ln(p_i) - \ln(p_j) \leq \ln\left(\frac{\sum_r v_{ir} x_{ir}}{v_{ij}}\right) \ \forall i, j,$$

$$\sum_{i=1}^{m} x_j^i = 1 \ \forall j,$$

$$x_j^i \geq 0 \ \forall i, j.$$

6 Auctions

Auctions are a venerable and popular selling institution. The word auction comes from the Latin *auctus* meaning to increase. We illustrate how the Farkas lemma can be used in the design of an auction.[7] An auction can take various forms, but for our purposes it consists of two steps. In the first, bidders announce how much they are willing to pay (bids) for the object. In the second, the seller chooses, in accordance with a previously announced function of the bids, who gets the object and how much each bidder must pay. This choice could be random.[8]

The simplest set up involves two risk neutral bidders and one seller. The seller does not know how much each bidder is willing to pay for the object. Each bidder is ignorant of the valuations of other bidders. It is the uncertainty about valuations that makes auctions interesting objects of study. If the seller knew the valuations, she would approach the bidder with the highest valuation and make him a take-it-or-leave-it offer slightly under their valuation and be done with.[9]

The uncertainty in bidder valuations is typically modeled by assuming that their monetary valuations are drawn from a commonly known distribution over a finite set W.[10]

In some contexts it is natural to suppose that the valuations are drawn independently. This captures the idea of values being purely subjective. In this case, the value that one bidder enjoys from the consumption of the good does not influence the value that the other bidders will enjoy. Here we suppose that the valuations are correlated. One context where such an assumption makes sense is in bidding for oil leases. The value of the lease depends on the amount of oil under the ground. Each bidder's estimate of that value depends on seismic and other surveys of the land in question. It is reasonable to suppose that one bidder's survey results would be correlated with another's, because they are surveying the same plot of land.

Denote by v^i the value that bidder i places on the object. For any two $a, b \in W$, let $p_{ab} = Pr[v^2 = b|v^1 = a] = Pr[v^1 = b|v^2 = a]$. The important assumption we make is that no row of the matrix $\{p_{ab}\}$ is a non-negative linear combination of the other rows. We refer to this as the **cone assumption.** Were the values drawn independently, the rows of this matrix would be identical.

Each bidder is asked to report their value. Let Q^1_{ab} be the probability that the object is assigned to agent 1 when he reports a and bidder 2 reports b. Notice that $Q^2_{ab} = 1 - Q^1_{ab}$. Let T^1_{ab} be the payment that bidder 1 makes if he

[7] This section is based on [1].

[8] In other books you can learn about the revelation principle, which explains why this formulation is general enough to encompass all auctions.

[9] Making an offer slightly under their valuation ensures the buyer strictly prefers to buy.

[10] This is known as the **common prior** assumption.

reports a and bidder 2 reports b. Similarly define T^2_{ab}.[11] The variable T^j_{ab} could be positive (a transfer from bidder to seller) or negative (a transfer from seller to bidder). Notice a bidder may end up making a payment even if it does not win the object. Given this, a bidder may not participate in the auction. We shall deal with this possibility later.

Two constraints are typically imposed on the auction design. The first is called **incentive compatibility**. The expected payoff to each bidder from reporting truthfully (assuming the other does so as well) should exceed the expected payoff from bidding insincerely. Supposing bidder 1's valuation for the object is a, this implies that

$$\sum_{b \in W} p_{ab}[Q^1_{ab}a - T^1_{ab}] \geq \sum_{b \in W} p_{ab}[Q^1_{kb}a - T^1_{kb}] \quad \forall k \in W \setminus a.$$

The left hand side of this inequality is the expected payoff (assuming the other bidder reports truthfully) to a bidder with value a who reports a. The right hand side is the expected payoff (assuming the other bidder reports truthfully) to a bidder with value a who reports k as their value. This constraint must hold for each $a \in W$ and a similar one must hold for bidder 2.

The incentive compatibility constraint does not force any bidder to bid sincerely. Only if all other bidders bid sincerely is it the case that one should bid sincerely. Furthermore, the inequality in the incentive compatibility constraint means that it is possible for a bidder to be indifferent between bidding sincerely or lying. At best the incentive compatibility constraint ensures that bidding sincerely is mutually rational. One could demand that the auction design offer greater incentives to bid sincerely than the ones considered here, but that is a subject for another paper.

The second constraint, called **individual rationality**, requires that no bidder should be made worse off by participating in the auction. This ensures the participation of each bidder. It is not obvious how to express this constraint as an inequality, since the act of participation does not tell us how a bidder will bid. This is where the incentive compatibility constraint is useful. With it, we can argue that if a bidder participates, she will do so by bidding sincerely. Hence, if bidder 1's valuation is $a \in W$ and he reports this, which follows from incentive compatibility, we can express individual rationality as:

$$\sum_{b \in W} p_{ab}[Q^1_{ab}a - T^1_{ab}] \geq 0.$$

This constraint must hold for each $a \in W$ and for bidder 2 as well.

The goal of the auctioneer is to design the auction so as to maximize her expected revenue subject to incentive compatibility and individual rationality. Notice that her expected revenue is maximized when the expected profit to

[11] We could make the payments random as well. However a risk neutral bidder would focus on the expected payments anyway, so we can, without loss of generality, suppose payments are deterministic.

all bidders is 0. Given incentive compatibility, bidder 1's expected profit when he values the object at a is

$$\sum_{b \in W} p_{ab}[Q^1_{ab}a - T^1_{ab}].$$

A similar expression holds for bidder 2. So, the auctioneer maximizes expected revenue if she can choose Q^j and T^j so that for all $a \in W$ bidder 1's expected profit is zero, i.e.,

$$\sum_{b \in W} p_{ab}[Q^1_{ab}a - T^1_{ab}] = 0,$$

and bidder 2's expected profit for all $b \in W$ is zero, i.e.,

$$\sum_{a \in W} p_{ab}[Q^2_{ab}b - T^2_{ab}] = 0.$$

Substituting this into the incentive compatibility and individual rationality constraints, the auctioneer seeks a solution to:

$$\begin{aligned}
\sum_{b \in W} p_{ab}[Q^1_{kb}a - T^1_{kb}] &\leq 0, \quad \forall k \in W \setminus a,\ a \in W,\\
\sum_{a \in W} p_{ab}[Q^2_{ak}b - T^2_{ak}] &\leq 0, \quad \forall k \in W \setminus b,\ b \in W,\\
\sum_{b \in W} p_{ab}[Q^1_{ab}a - T^1_{ab}] &= 0, \quad \forall a \in W,\\
\sum_{a \in W} p_{ab}[Q^2_{ab}b - T^2_{ab}] &= 0, \quad \forall b \in W,\\
Q^1_{ab} + Q^2_{ab} &= 1 \quad \forall a, b \in W,\\
Q^1_{ab}, Q^2_{ab} &\geq 0 \quad \forall a, b \in W.
\end{aligned}$$

It turns out to be convenient to fix the value of Q^j's in the inequalities above and ask if there is a feasible T^j. Rewriting the above inequalities by moving terms that are fixed to the right hand side (with a change in index on the last two to make the Farkas alternative easier to write out):

$$-\sum_{b \in W} p_{ab}T^1_{kb} \leq -\sum_{b \in W} p_{ab}Q^1_{kb}a, \quad \forall k \in W \setminus a,\ a \in W, \tag{13}$$

$$-\sum_{a \in W} p_{ab}T^2_{ak} \leq -\sum_{a \in W} p_{ab}Q^2_{ak}b, \quad \forall k \in W \setminus b,\ b \in W, \tag{14}$$

$$\sum_{b \in W} p_{kb}T^1_{kb} = \sum_{b \in W} p_{kb}Q^1_{kb}k, \quad \forall k \in W, \tag{15}$$

$$\sum_{a \in W} p_{ak}T^2_{ak} = \sum_{a \in W} p_{ak}Q^2_{ak}k, \quad \forall k \in W. \tag{16}$$

Let y^1_{ak} be the variable associated with the first inequality, y^2_{kb} be associated with second inequality, z^1_k with the third, and z^2_k with the fourth set of inequalities.

The Farkas lemma asserts that there is no solution to (13)-(16) if there is a solution to the system:

$$-\sum_{a\neq k} p_{ab}y^1_{ak} + p_{kb}z^1_k = 0, \forall k, b \in W,$$
$$-\sum_{b\neq k} p_{ab}y^2_{kb} + p_{ak}z^2_k = 0, \forall a, k \in W,$$
$$y \geq 0,$$

such that

$$-\sum_{a\in W}\sum_{k\neq a}\Big[\sum_{b\in W} p_{ab}Q^1_{kb}a\Big]y^1_{ak} - \sum_{b\in W}\sum_{k\neq b}\Big[\sum_{b\in W} p_{ab}Q^2_{ak}a\Big]y^2_{kb} +$$
$$\sum_{k\in W}\sum_{b\in W} p_{kb}Q^1_{kb}kz^1_k + \sum_{k\in W}\sum_{b\in W} p_{kb}Q^2_{kb}kz^2_k < 0.$$

Using the first equation and non-negativity of the p's and the y's, we conclude that the z's must be non-negative as well. The last inequality, which must hold strictly, prevents all of the y variables being zero. Given this, the first equation contradicts the cone assumption made earlier. Thus, the Farkas alternative has no solution, implying that (13)-(16) has a solution.

Acknowledgments

The author's research was supported in part by NSF grant ITR IIS-0121678. The paper has benefited from the comments of Bruce Golden.

References

1. R. Cremer and R. McLean. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* , 56:1247-1258, 1988.
2. W. Feller. *An Introduction to Probability Theory and its Applications.* Volume 1, third edition, Wiley, New York, 1968.
3. O. Mangasarian. *Nonlinear Programming.* SIAM, Philadelphia, 1994.
4. R. Nau and K. McCardle. Arbitrage, rationality, and equilibrium. *Theory and Decision* 31:199-240, 1992.
5. M. Primak. An algorithm for finding a solution of the linear exchange model and the linear Arrow-Debreu model. *Kibernetika*, 5:76-81, 1984.

Parametric Cardinality Probing in Set Partitioning

Anito Joseph and Edward K. Baker

Department of Management Science
School of Business
University of Miami
Coral Gables, Florida 33124
ajoseph@miami.edu, ebaker@miami.edu

Summary. In this work, we investigate parametric probing methods based on solution cardinality for set partitioning problems. The methods used are inspired by the early work of Gass and Saaty on the parametric solution to linear programs, as well as the later work of Joseph, Gass, and Bryson that examined the duality gap between the integer and relaxation solutions to general integer programming problems. Computational results are presented for a collection of set partitioning problems found in the literature.

Key words: Set partitioning; parametric cardinality probing.

1 Introduction

The set partitioning problem (SPP) assumes a set of m elements that are to be partitioned into mutually exclusive and collectively exhaustive subsets. In the enumeration of n possible subsets, one may define a matrix $[a_{ij}]$ where $a_{ij} = 1$ if the i-th element of the set is contained in the j-th subset, and $a_{ij} = 0$ otherwise. A set of binary decision variables associated with the n subsets is used in the model where $x_j = 1$ if subset j is used in the solution, and $x_j = 0$ otherwise. If the cost of creating each subset, j, is given as c_j, then the minimal cost set partitioning problem may be specified as the following binary linear program:

Minimize $\quad \Sigma\, c_j\, x_j,$

Subject to: $\quad \Sigma\, a_{ij}\, x_j = 1,\ i = 1, 2, \ldots, m$, and $a_{ij} = [0,1].$

$\qquad\qquad x_j = [0,1],\ j = 1, 2, \ldots, n.$

The set partitioning model has been used in a wide variety of successful applications. The vehicle routing problem, for example, was first formulated as a set partitioning problem in 1964 [1]. Since that time, many researchers, see [21] for example, have used the model. Similarly, the airline crew scheduling problem was formally posed as a set partitioning problem in 1969 [1]. Subsequent solution approaches to the problem have included linear programming-based methods [22],

heuristic procedures [3], and column generation approaches [25]. Garfinkel and Nemhauser [6] considered the use of the set partitioning problem in the solution of political districting problems. This formulation has recently been solved effectively using branch and cut methods [24].

The cardinality of the solution of the set partitioning problem is equal to the number of subsets used in the solution. In the mathematical formulation of the problem given above, the cardinality of the solution may be determined by the summation of the x_j binary decision variables. This is simply:

$$\Sigma x_j = y,$$

where y is the cardinality of the solution. In cases where non-integer values of the decision variables are allowed, for example in the linear programming relaxation of the SPP, the value y will be called the global cover of the solution.

Set partitioning problems of a specified cardinality occur frequently in practice. In vehicle routing, for example, it may be desired to partition a set of m delivery customers into exactly k delivery routes. Similarly, in various crew scheduling applications where the number of crews is fixed, the m tasks to be scheduled are to be partitioned among the k crews. This is the case in the crew scheduling problems faced by many European airlines where the number of crews operating each month is fixed.

The closeness, numerically and structurally, of the integer and the associated linear relaxation solution to integer programs is frequently observed in the literature [12]. Consequently, many solution algorithms proposed for integer programming problems use the linear programming relaxation solution as their starting point. Typically, numerical closeness has been measured in terms of the gap between objective function values. Joseph, Gass and Bryson [19], examined the relationships between both the objective function values and the structure of the solution vectors for general integer programming problems and their respective relaxations. In the case of the SPP, it has been noted by various researchers, e.g. [11] and [23], that when the number of rows is small, linear programming, or linear programming with branch-and-bound, can provide integer solutions quickly. Further, it has been shown that if an application can be modeled as a set partitioning problem with a totally unimodular structure, then the linear programming solution is guaranteed to be integer [14], [15], and [16]. Finally, a decomposition approach, proposed by Joseph [18], was successful in quickly identifying optimal solutions to difficult set partitioning problems. The success of this approach was due in part to the inherent closeness of the relaxed and integer solution vectors so that enumerating a very small number of fractional variables gave rise to sub-problems that were solved in a fraction of the time it would have taken to solve the original set partitioning problem.

In this paper, we investigate the formal and explicit consideration of cardinality in the solution of minimal cost set partitioning problems. Specifically, we probe solution cardinality using linear programming techniques to search the global cover contours of the convex polytope as an approximation of the cardinality contours of the integer problem. The results of the cardinality probing are used to obtain

bounds on the optimal integer solution and on the possible cardinality of that solution. Such bounds are then used within a conventional branch and bound algorithm to reduce the effort required to identify the optimal integer solution. Our approach starts by appending the cardinality constraint to the mathematical model of the set partitioning problem and allowing the cardinality to vary. In this analysis, we vary the cardinality in the manner of the parametric dual approach of Gass [7, p.157]. This parametric analysis was first proposed in the seminal papers of Gass and Saaty [8], [9], and [28], and has since become part of the standard post-optimality analysis of linear programming.

Other researchers have also explored the relationship between cardinality and the solution to the set partitioning problem. The use of cardinality to provide bounds on the enumeration within a traditional cost minimization branch and bound procedure was explored in [17] and [18]. Similarly in [20], the performance of standard branch and bound solvers in the solution to set partitioning problems found in the literature was investigated. The authors found in a number of cases that simply adding the "optimal" cardinality constraint to the problem formulation significantly reduced the number of iterations required to find the optimal integer solution.

Exploration of solution cardinality has also been found to be an effective methodology in the field of constraint programming. Cardinality has been used as one of the multi-threaded avenues of solution space search in a hybrid approach to solving combinatorial problems [31]. Several authors, for example [27] and [30], have proposed constraint programming approaches for various scheduling problems classically formulated as set covering or set partitioning problems. Finally in [2], the general approach to set constraint solving is discussed and a specific constraint solver based on cardinality inferences is proposed. The author presents several applications, including the set covering model, and indicates significant computational efficiencies over other methods.

Probing techniques applied to a solution space are, in general, expedient partial explorations used to obtain information of value in directing a more complete exploration of the space. Several authors, for example [5] and [10], consider various methods for reducing the size and solution space of large problems. A recent survey paper [29], considers various probing techniques, and their implementation, in the solution of mixed integer programming problems.

2 Parametric Probing of Solution Cardinality

The solution to the linear programming relaxation of the SPP has been a traditional starting point in the search for the optimal integer solution for set partitioning problem. If one augments the relaxed problem with the constraint, $\Sigma x_j - y = 0$, one can monitor the global cover, y, of the structural variables, i.e. the x_j variables, in the solution. Let the original set partitioning problem augmented with an additional $m+1^{st}$ constraint $\Sigma x_j - y = 0$, be defined as the problem **P0**. That is, let

P0: $\{\text{Min } Z_0 = \Sigma c_j x_j \mid \Sigma a_{ij} x_j = 1, \Sigma x_j - y = 0; x_j = 0 \text{ or } 1; j = 1,..., n, i = 1,..., m, y \geq 0\}$.

The solution to the relaxation of **P0** may result in either of two possible cases. In the first case, the structural variables are all integer and y is equal to the cardinality of the solution. In the second case, some or all of the structural variables are fractional. In this case, the solution cardinality is unresolved and y, as the global cover, gives an estimate of the possible solution cardinality. Since the partitioning problem is solved in the first case, it is the second case which is of interest.

When the solution of the linear programming relaxation of **P0** is fractional, the global cover, y, in the solution may be either fractional or integer. When the global cover is fractional, we begin the probing procedure by letting $w = floor(y)$, where *floor*(y) is the greatest integer value less than or equal to y. When the global cover is integer, we begin the probing procedure by setting w equal to y.

Using the value w as a starting point, we parametrically vary the global cover of the solution by θ integer units, i.e. $\Sigma x_j = w + \theta$, (where θ can be positive or negative) to obtain objective function bounds on the integer solution for different possible cardinality values. This is a special case of the parametric dual procedure since only integer values of θ apply in the search for an integer solution. By carrying out the parametric probing analysis for a range of possible cardinalities surrounding the starting point, we can determine bounds, and possibly feasible integer solutions, that may be used to curtail the search.

It is noted that when y is integer-valued, the above procedure for cardinality probing cannot result in improved lower bound information since initially $y = w$. To obtain improved bound information in these cases, we focus on the cardinality of specific sub-problems. A sub-problem in this instance can be defined for any subset of rows with a non-integer cover. Therefore, any fractional-valued column x_j such that $a_{ij} = 1$ in two or more rows can be used to identify subproblems with an unresolved cardinality. For example, let column x_h be fractional in the LP solution and let $a_{rh} = a_{ph} = 1$. We can then we can find Σx_k, where $k \in \{j: a_{ij} = 1; i = r, \text{ or } i = p\}$. Therefore, $\Sigma x_k = \upsilon$ gives the local cover for rows r and p. To illustrate let $x_1 = .67$ with $a_{21} = a_{41} = 1$, and let $x_3 = .33$ and $x_5 = .33$ where ($a_{23} = 1$, $a_{43} = 0$), and ($a_{25} = 0$, $a_{45} = 1$). Hence, the cover for both rows 2 and 4 are satisfied, and $\Sigma x_k = 1.33$.

For any sub-problem where υ is non-integer, two new problems **P1** and **P2** can be obtained by respectively appending the local cover constraint $\Sigma x_k \leq floor\ (\upsilon)$, and $\Sigma x_k \geq ceiling\ (\upsilon)$ to SPP or **P0**. Resolving the sub-problem cardinality is, therefore, equivalent to branching on the rows of the matrix A. Using the sub-problem approach provides an opportunity to search within the space where $y = w$ as the initial global cover by using local cover constraints. These constraints can be identified and appended to problem **P0** individually as the $m+2^{nd}$ constraint giving rise to problems **P1** and **P2**. In this way improved bound information may be obtained for integer-valued covers. The use of these two constraints, one in **P1** and the other in **P2**, eliminates any extreme points where $\Sigma x_j = w$ and floor $(\upsilon) \leq \upsilon \leq$ ceiling (υ), and ensures that only extreme points with integer-valued covers (w and υ) are considered.

3 Computational Results

We report results for a selection of test problems taken from the literature. The test problems are a subset of the pure set partitioning real world airline crew scheduling problems used in [12, p.676]. A description of the problems used in the current paper is presented in Table 1. In the table, the problem dimensions are given as the number of rows and the number of columns. The objective function values given are the minimal cost solution to the linear programming relaxation of the partitioning problem, Z_{LP}, and the optimal integer solution, Z_{IP}. (These values are may also be found in [12]. In the column titled Global Cover, the values of y for both the linear relaxation and the integer solution are given.

Table 2 shows the results of the parametric probing when the linear relaxation of **P0** is solved. The probing produced the optimal integer solution in eight of the 34 cases considered. In twenty-nine of the problems, an integer solution was found by parametric probing. Improved lower and/or upper bounds were obtained for thirty-one of the problems studied. For larger problems, e.g. 7292x646, the single cardinality cut of problem **P0** was not effective at improving solution bounds and suggests that a better definition of the cardinality is needed than can be obtained from the column cover constraint of problem **P0**.

3.1 Bound Information gained from Parametric Probing

A summary of the bound information obtained from the parametric cardinality probing using problem **P0** is presented in Table 3. The lower bound shown is the minimum objective function value over the range of cardinalities explored, while the upper bound is the minimum cost integer solution found over the range of cardinalities studied.

For thirteen of the problems the lower bound information could not be improved beyond the linear programming solution for **P0** because the global cover y was integer-valued. For two of these thirteen problems, 7479x55 and 36699x71, parametric probing showed that the only possible value for the solution cardinality was equal to the global cover y of the solution to the linear programming relaxation of **P0**.

To obtain improved bounds at the column cover y, we move to a local search and use the sub-problem concept to correct unresolved cardinality for a subset of columns. The simplest implementation of the cardinality correction is where we choose one fractional-valued column and form the local cover constraints over the remaining n-1 columns. For this implementation, the column selected belongs to the set F where $F = \{j: 0.5 \leq x_j < 1\}$ and has the cheapest cost per unit of row cover, i.e. $c_k/a_k = \min (c_j/a_j, j \in F)$. If the set F is empty, then the fractional column with the cheapest cost per unit of row cover is selected. Such columns are critical columns contributing the at least one-half the cover requirement in their associated rows and will be replaced by columns of the same or greater cost contribution. Therefore, to find a lower bound solution for the case where the global cover y is integer-valued, we can probe the cardinality of this subset of (n-1) columns.

Table 1. Description of test problems.

	Problem Dimensions		Global Cover (y)		Objective Function	
	Columns	Rows	LP	IP	Z_{LP}	Z_{IP}
1	197	17	4.5	5	10972.5	11307
2	294	19	6.333	7	14570.0	14877
3	404	19	3.75	4	10658.3	10809
4	577	25	6.5	7	7380	7408
5	619	23	6.5	7	6942.0	6984
6	677	25	6.5	5	9868.5	10080
7	685	22	5	5	16626.0	16812
8	711	19	8	8	12317.0	12534
9	770	19	4	4	9961.5	10068
10	771	23	5.5	6	6743.0	6796
11	899	20	4	4	10453.5	10488
12	1072	18	5.5	6	8897.0	8904
13	1079	23	4.5	4	7485.0	7656
14	1210	18	3.5	3	8169.0	8298
15	1217	20	5	5	5852.0	5960
16	1220	23	5	5	5552.0	5558
17	1355	22	4.5	5	9877.0	9933
18	1366	19	7	7	5843.0	6314
19	1709	23	6	6	7206.0	7216
20	1783	20	3.333	4	7260.0	7314
21	2540	18	3.5	4	4185.3	4274
22	2653	26	3.8	4	3726.8	3942
23	2662	26	4.5	4	7980.0	8038
24	3068	23	5	5	6484.0	6678
25	6774	50	8.667	8	7640	7810
26	7292	646	94.625	95	26977.2	27040
27	7479	55	13	13	1084	1086
28	8820	39	19	19	116254.5	116256
29	10757	124	40.625	41	338864.3	340160
30	16043	51	22	22	50132.0	50146
31	28016	163	23.333	24	17731.7	17854
32	36699	71	17	17	215.3	219
33	43749	59	11.5	13	24447.0	24492
34	118607	61	15	15	10875.8	11115

Table 2. Computational results of parametric probing.

		Value of θ for [w] + θ						
Problem	[w]	≤-3	-2	-1	0	1	2	≥3
197x17	4	Infeas	Infeas	Infeas	13444.5	11307**	12441*	13812*
294x19	6	Infeas	Infeas	17082*	14688	14877**	15285*	16998*
404x19	3	Infeas	Infeas	Infeas	Infeas	10803	11610*	12826.5
577x25	6	Infeas	Infeas	8266*	7436*	7408**	8128*	9112*
619x23	6	Infeas	9088*	7580*	6943	6963	Infeas	Infeas
677x25	4	Infeas	Infeas	Infeas	14517*	10056	10455*	11361*
685x22	5	Infeas	Infeas	16899	16626[a]	16970	17733	18843*
711x19	8	Infeas	15008*	12754	12317[a]	12851	13702*	14696*
770x19	4	Infeas	Infeas	Infeas	9961.5[a]	10336.5	10828.5	11721*
771x23	5	Infeas	Infeas	8691	6871	6796**	7658*	8664*
899x20	4	Infeas	Infeas	Infeas	10453.5[a]	11274*	12418.5	13962*
1072x18	5	Infeas	Infeas	Infeas	9323	8904**	9638*	10734*
1079x23	4	Infeas	Infeas	Infeas	7656**	7625	7946*	9138
1210x18	3	Infeas	Infeas	Infeas	8247	8658	9837*	11403*
1217x20	5	Infeas	7650*	6324.57	5852[a]	6334	6964*	8066*
1220x23	5	Infeas	Infeas	6552*	5552[a]	6096*	7242*	8586*
1355x22	4	Infeas	Infeas	Infeas	10321.5	9933**	11091*	12396*
1366x19	7	7978*	6688	6169	5843[a]	6205.5	6568*	7274*
1709x23	6	Infeas	Infeas	8385.33	7206[a]	7340*	8128*	9002
1783x20	3	Infeas	Infeas	Infeas	7364	7314**	7780*	8848*
2540x18	3	Infeas	Infeas	Infeas	5082	4250	4560*	5208*
2653x26	3	Infeas	Infeas	Infeas	5612.67	3737	3846	4224
2662x26	4	Infeas	Infeas	Infeas	8022	8094*	8484*	
3068x23	5	Infeas	9177.14	7102	6484[a]	7272	8298	9328*
6774x50	8	Infeas	8541.5	7863.1	7663.3	7646	8026*	8642*
7292x646	94	27178.5	27067.7	27010.1	26979.1	26980.9	27001.4	27064.3
7479x55	13	Infeas	Infeas	Infeas	1084[a]	Infeas	Infeas	Infeas
8820x39	19	116259*	116373*	116268*	116254.5[a]	116265*	116826*	117930*
10757x124	40	Infeas	Infeas	Infeas	338984	338919	339612	340746
16043x51	22	50556	50384	50212	50312[a]	50304*	50672*	51170*
28016x163	23	18644.2	17955	17777.7	17741.3	17753.7	17802	18129*
36699x71	17	Infeas	Infeas	Infeas	215.3[a]	Infeas	Infeas	Infeas
43749x59	11	Infeas	27837.9	25393.5	24666.4	24460.2	24492*	25395*
118607x61	15	11661.4	11101.44	10917.7	10875.75	10976	11359.4	11905.2
* integer solution , ** optimal integer solution, [a] LP bound from P0								

Table 3. A summary of parametric cardinality probing bound information.

Problem	Z_{LP}	Parametric Bounds		Z_{IP}
		Lower	Upper	
197x17	10972.5	11307	11307	11307
294x19	14570.0	14688	14877	14877
404x19	10658.3	10803	11610	10809
577x25	7380.0	7408	7408	7408
619x23	6942.0	6943	7580	6984
677x25	9868.5	10056	10455	10080
685x22	16626.0	-	18843	16812
711x19	12317.0	-	13702	12534
770x19	9961.5	-	11721	10068
771x23	6743	6796	6796	6796
899x20	10453.5	-	11274	10488
1072x18	8897.0	8904	8904	8904
1079x23	7485	7625	7656	7656
1210x18	8169.0	8247	9837	8298
1217x20	5852.0	-	6964	5960
1220x23	5552.0	-	6096	5558
1355x22	9877.5	9933	9933	9933
1366x19	5843.0	-	6568	6314
1709x23	7206.0	-	7340	7216
1783x20	7260.0	7314	7314	7314
2540x18	4185.3	4250	4560	4274
2653x26	3726.8	3737	4942	3942
2662x26	7980	8022	8094	8038
3068x23	6484	-	9328	6678
6774x50	7640	7646	8026	7810
7292x646	26977.2	26979.1	-	27040
7479x55	1084	-	-	1086
8820x39	116254.5	-	116265	116256
10757x124	338864.3	338919	-	340160
16043x51	50132.0	-	50304	50146
28016x163	17731.7	17741.3	18129	17854
36699x71	215.3	-	-	219
43749x59	24447.0	24460.2	24492	24492
118607x61	10875.8	10903	-	1115

Further, we fix the global cover at its original solution value y and conduct the search (solve **P1** and **P2**) within this y value. We find the best solution when the omitted column may be considered (**P1**) and when the omitted column may not be considered (**P2**). The lower of these two objective function values becomes the new lower bound at cover y.

Table 4. Bounds for LP relaxations with integer global covers.

Problem Size	Z_{LP}	Min (cj/mj)	P1	P2	Z_{IP}
685x22	16626.0	3015/7	16779	16772.2	16812
711x19	12317	198/2	18976*	12334	12534
770x19	9961.5	1650/8	10233*	10068**	10068
899x20	10453.5	2235/5	10701*	10488**	10488
1217x20	5852	240/3	6626	5941	5960
1220x23	5552	1152/10	5630*	5554	5558
1366x19	5843	190/2	5922	6320.67	6314
1709x23	7206	262/3	7368*	7216**	7216
3068x23	6484	240/2	6678**	6682*	6678
8820x39	116254.5	1104/4	116259*	116256**	116256
16043x51	50132	(1992/6)	50158*	50146**	50146

Table 4 shows results for eleven of the thirteen problems with integer values for the global cover. Significant improvements were obtained by removing only one column and searching on the remaining columns. Optimal integer solutions were found for six problems and bound improvements were significant for the remaining five. The remaining two of the set of thirteen problems with integer global covers, 36699 x 71 and 7479 x 55, showed minimal or no improvement based on selecting only one column. These are larger problems with initially very small duality gaps. Their structure was made up of subgroups of columns that were very similar in terms of cost and row cover. Thus the current strategy of omitting only one column to correcting the cardinality of the remaining columns meant that there were many alternatives that could replace the column at a similar cost contribution. This suggests that the criteria for the cardinality restriction should involve more than one column being omitted, or be focused on specific subsets of the problem columns

4 Conclusion

Using parametric probing of cardinality contours can lead to valuable information that can be used to bound and guide the search for an optimal solution to the SPP. For the test problems studied, simple cardinality cuts were shown to be effective in finding optimal solutions and for bounding the search. The results obtained point to the potential of including cardinality as a consideration in the search and exploiting yet another dimension of closeness in solution vectors. By considering only integer values for the column cover, the cardinality probing approach generates cuts that eliminate extreme points of the polytope that have non-integer covers, considering only those extreme points with the specified integer covers. Thus it is guaranteed that no feasible integer solution point will be eliminated. The cuts used in this study were straightforward; however, the approach needs to be further developed for dealing with larger and more complex problem structures.

The computational results of the cardinality probing suggest a possible row branching scheme for solving the set partitioning column. The row branching scheme would focus on the number of columns that will cover a chosen set of rows rather than having to make specific choices directly from among a number of eligible columns. The particular combination of rows selected for branching on is guided by information gathered from the solution of the linear programming relaxation of SPP. Overall, cover requirements are straightforward to implement and the information for guiding the search can be obtained at a small computational cost.

References

1. J.P. Arabeyre, J. Fearnley, F.C. Steiger, and W. Teather. The airline crew scheduling problem: a survey. *Transportation Science,* 3:140-163, 1969.
2. F. Azevedo. Constraint solving over multi-valued logics. In *Frontiers in Artificial Intelligence and Applications.* Vol. 91. IOS Press. Berlin. 2003
3. E.K. Baker, L.D. Bodin, W.F. Finnegan, and R.J Ponder. Efficient heuristic solution to an airline crew scheduling problem. *AIIE Transactions,* 11:79-85, 1979.
4. M. Balinski and R. Quandt. On an integer program for a delivery problem. *Operations Research,* 12:300-304, 1964.
5. A.L. Brearley, G. Mitra, and H.P. Williams. Analysis of mathematical programming problems prior to applying the simplex algorithm. *Math Programming,* 8:54-83, 1975.
6. R.S. Garfinkel and G.L. Nemhauser. Optimal political districting by implicit enumeration techniques. *Management Science,* 16:B495-B508, 1970.
7. S.I. Gass. *Linear Programming.* Fourth Edition. McGraw-Hill. New York. 1972.
8. S.I. Gass and T.L. Saaty. Parametric objective function part II: generalization. *Operations Research,* 3:395-410, 1955.
9. S.I. Gass and T.L. Saaty. The computational algorithm for the parametric objective function. *Naval Research Logistics Quarterly,* 2:112-126, 1955.
10. M. Guignard and K. Spielberg. Logical reduction methods in zero-one programming: minimal preferred inequalities. *Operations Research,* 29:49-74, 1981.
11. K. Hoffman and M. Padberg. Improving LP-representations of zero-one programs for branch-and-cut. *ORSA Journal on Computing,* 2:121-134, 1991.
12. K. Hoffman and M. Padberg. Solving airline crew scheduling problems by branch-and-cut. *Management Science,* 39:657-682, 1993.
13. A. Joseph. A parametric formulation of the general integer linear programming problem. *Computers and Operations Research,* 22:883-892, 1995.
14. A. Joseph and N. Bryson. Partitioning of sequentially ordered systems using linear programming. *Computers and Operations Research,* 24:679-686, 1997.
15. A. Joseph and N. Bryson. W-efficient partitions and the solution of the sequential clustering problem. *Annals of Operations Research: Nontraditional Approaches to Statistical Classification and Regression,* 74:305-319, 1997.

16. A. Joseph and N. Bryson. Parametric linear programming and cluster analysis. *European Journal of Operational Research*, 111:582-588, 1998
17. A Joseph. Cardinality corrections for set partitioning. Working paper, Department of Management Science, University of Miami. 1999.
18. A. Joseph. A concurrent processing framework for the set partitioning problem, *Computers and Operations Research*, 29:1375-1391, 2002.
19. A. Joseph, S.I. Gass, and N. Bryson. An objective hyperplane search procedure for solving the general integer linear programming problem, *European Journal of Operational Research*, 104:601-614, 1998.
20. A. Joseph, B. Rayco and E. Baker. Cardinality constrained set partitioning. University of Miami Working Paper. Coral Gables, Florida. 2004.
21. J.P. Kelly and J. Xu. A set-partitioning-based heuristic for the vehicle routing problem. *INFORMS Journal on Computing*, 11:161-172, 1999.
22. R.E. Marsten. An algorithm for large set partitioning problems. *Management Science*, 20:774-787, 1974.
23. R.E. Marsten and F. Shepardson. Exact solution of crew scheduling problems using the set partitioning model: recent successful applications. *Networks*, 11:165-177, 1981.
24. A. Mehrotra, G.L. Nemhauser and E. Johnson. An optimization-based heuristic for political districting, *Management. Science*, 44:1100 – 1114, 1998.
25. A. Mingozzi, M.A. Boschetti, S. Ricciadelli and L. Bianco. A set partitioning approach to the crew scheduling problem. *Operations Research*, 47:873-888, 1999.
26. G. Mitra. Investigation of some branch and bound strategies for the solution of mixed integer linear programs. *Mathematical Programming*, 4:155-170, 1973.
27. J. Regin. Minimization of the number of breaks in sports scheduling problems using constraint programming. In *Constraint Programming and Large Scale Optimization*. DIMACS, 57:115-130. 2001.
28. T.L. Saaty and S.I. Gass. Parametric Objective Function Part I, *Operations Research*, 2:316-319, 1955.
29. M.W.P. Savelsbergh. Preprocessing and probing techniques for mixed integer programming problems. *ORSA Journal on Computing*, 4:445-454, 1994.
30. B.M. Smith, C.J. Layfield and A. Wren. A constraint pre-processor for a bus driver scheduling system. In *Constraint Programming and Large Scale Optimization*. DIMACS, 57:131-148, 2001.
31. F. Zabatta. Multithreaded constraint programming: a hybrid approach. In *Constraint Programming and Large Scale Optimization*. DIMACS, 57:41-64, 2001.

A Counting Problem in Linear Programming

Jim Lawrence

Department of Mathematical Sciences
George Mason University
Fairfax, VA 22030–4444
lawrence@gmu.edu

Summary. Using a popular setup, in solving a linear programming problem one looks for a tableau of the problem that has no negative elements in the last column and no positive elements in the last row. We study a matrix whose (i, j)–th entry counts the tableaux for such a problem (here taken to be totally nondegenerate) having i negative elements in the last column and j positive elements in the last row. It is shown that this matrix possesses a certain symmetry, which is described.

Key words: Linear programming; oriented matroid; tableau enumeration; polytope; h–vector; f–vector.

1 Introduction

It is assumed that the reader has some familiarity with linear programming. If not, there are a great many books that were written to serve as textbooks on linear programming; they are all, to some extent, descendants of the first such book written by Saul Gass [2].

Suppose we have a linear programming problem (assumed to be "totally nondegenerate"), having s nonnegative variables and r additional inequality constraints. In order to solve the problem using the simplex method, we may construct a suitable tableau and, by pivoting, attempt to move to a tableau in which the last column has no negative entries and the last row has no positive entries (ignoring the entry they share in common). A crude measure of progress toward this solution tableau is indicated by the pair of numbers (a, b), where a is the number of negative entries in the last column and b is the number of positive entries in the last row of the current tableau. Of course, these numbers may go up and down during the trek; it's not at all clear what this information tells us about getting to the solution. Even so, when these numbers aren't too big, intuition seems to dictate that we are "getting warmer," and that a tableau with $a = b = 0$ may not be many steps away. Thus we are led to the question of what can be said about the

$(r+1) \times (s+1)$ matrix N whose (a,b)–th entry (where $0 \leq a \leq r$, $0 \leq b \leq s$) is the number of tableaux for the problem having a negative entries in the last column and b positive entries in the last row. The main theorem below gives a simple property of this matrix.

This question is related to certain enumerative results concerning convex polytopes, arrangements of hyperplanes, and oriented matroids. Since our linear programming problem is totally nondegenerate, and under the assumption that the feasible region for the problem is nonempty and bounded, this feasible region is a simple convex polytope P of dimension s. ("Simple" means that, at each vertex, exactly s of the $r+s$ inequality constraints are satisfied with equality — each vertex lies on exactly s facets. Equivalently, exactly s edges emanate from each vertex.) Letting f_k denote the number of k–dimensional faces of P, so that f_0 is the number of vertices, f_1 is the number of edges, f_{s-1} is the number of facets, and $f_s = 1$, the vector $(f_0, \ldots, f_s)$ is the *f–vector of P*.

Each vertex of P is on exactly s edges, since P is simple. The objective function is not constant on any edge of this polytope, again by total nondegeneracy. We say that an edge that emanates from a vertex is *leaving* if the value of the objective function at that vertex is less than its value at the other vertex of the edge (so that the objective function increases when one moves along the edge away from the given vertex). Let h_j $(0 \leq j \leq s)$ denote the number of vertices of P for which there are exactly j leaving edges. The vector $(h_0, \ldots, h_s)$ is called the *h–vector*.

The f–vector can be determined from the h–vector, by the following system of linear equations:

$$f_k = \sum_{j=0}^{s} \binom{j}{k} h_j \quad \text{for } 0 \leq k \leq s.$$

These equations result from the fact that each k–dimensional face of P possesses a unique vertex at which the objective function is minimized, and, if a given vertex has exactly j "leaving" edges, so that it contributes to h_j, then exactly $\binom{j}{k}$ faces of dimension k achieve their minima at the vertex.

Since the objective function factors into the determination of the h–vector in the above, one might guess that there would be lots of different h–vectors, depending upon the objective function chosen. In fact, there is only one, as can be seen by noting that the above system of equations is triangular with 1's on the diagonal, and therefore invertible. Upon inversion, it is clear that the h–vector can be determined from the f–vector. Since the f–vector does not depend upon the choice of objective function, any objective function (for which the nondegeneracy assumption is satisfied) yields the same h–vector. In particular, if the objective function is replaced by its negative, one sees that the h–vector is symmetric: $h_j = h_{s-j}$ $(0 \leq j \leq s)$. When these equations are written in terms of the f–vector, they are known as the *Dehn–Somerville equations.*

The tableaux for our problem, which has no negative entries in the last column, corresponds to the vertices of P. The number of outgoing edges at the vertex is the number of positive entries in the last row. It follows that the h–vector is the first row of N. (It can be shown that all entries of the last row of N are zero, in this case.) Similarly, if the dual of our problem has a feasible region that is nonempty and bounded, then the h-vector for the polytope that is the dual feasible region appears as the first column of N (and the last column of N consists of 0's).

It is the symmetry of the h–vector noted above that is generalized in this paper for the entire matrix N (with a modification that removes the need for the condition that the feasible region be nonempty and bounded).

The h–vector figured prominently in McMullen's proof in [6] of the "Upper Bound Conjecture" concerning the maximum number of vertices that a polytope of given dimension and with a given number of facets might have, and in the proof, jointly by Billera and Lee [1] and Stanley [9], of the characterization of the f–vectors of the simple polytopes, conjectured by McMullen in [7]. Related counting problems for arrangements of hyperplanes and, more generally, for oriented matroids have been considered, for example, in [4], [5]. The "mutation count matrices" of [5] are related to, but not the same as, the matrices N considered here.

In what follows, A denotes an $r \times s$ matrix of real numbers; B is a column vector of length r; C is a row vector of length s; D is a real number; and M is the $(r+1) \times (s+1)$ composite block matrix

$$M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}.$$

The matrix M, augmented by row and column labels indicating the variables associated with them (denoted by $x_1, \ldots, x_n$, where $n = r + s$), is the Tucker tableau for the linear programming problem (described by Goldman and Tucker, [3]):

$$\begin{aligned} \text{Maximize} \quad & CX_{(0)} - D \\ \text{subject to} \quad & AX_{(0)} \leq B, \\ & X_{(0)} \geq 0. \end{aligned}$$

Here $X_{(0)}$ denotes the column vector $X_{(0)} = [x_1, \ldots, x_s]^T$ of real variables.

The simplex method solves such a problem by performing "pivot" steps beginning with the matrix M and ending at a matrix from which the solution can be easily obtained. A *(Tucker) pivot* on the (nonzero) entry M_{i_0,j_0} of M consists of changing M by performing the following operations:

- Each entry $M_{i,j}$ not in the same row or column as M_{i_0,j_0} (so that $i \neq i_0, j \neq j_0$), is replaced by $M_{i,j} - M_{i,j_0} M_{i_0,j} / M_{i_0,j_0}$;
- Each entry M_{i,j_0} in the same column as the pivot entry M_{i_0,j_0} but different from it is replaced by $-M_{i,j_0} / M_{i_0,j_0}$;

- Each entry $M_{i_0,j}$ in the same row as the pivot entry but different from it is replaced by $M_{i_0,j}/M_{i_0,j_0}$;
- The entry itself is replaced by its inverse, $1/M_{i_0,j_0}$; and finally
- The labels of the row and column of the pivot element are exchanged.

For example, if

$$A = \begin{pmatrix} 1 & 2 & 3 \\ 6 & 3 & 1 \end{pmatrix}, \qquad B = \begin{pmatrix} 6 \\ 24 \end{pmatrix},$$

$$C = (1 \quad 1 \quad 1), \quad \text{and} \quad D = 0,$$

then we have the following Tucker tableau. (Here and later, T_{ijk}, where $1 \leq i < j < k \leq 5$, represents a tableau having x_i, x_j, x_k, and u as its column labels, not necessarily in that order.)

T_{123} :

	x_1	x_2	x_3	u
x_4	1	2	3	6
x_5	6	3	1	24
v	1	1	1	0

Upon performing a pivot on the entry whose row is labeled by x_5 and whose column is labeled by x_1, we get the following new tableau.

T_{235} :

	x_5	x_2	x_3	u
x_4	-1/6	3/2	17/6	2
x_1	1/6	1/2	1/6	4
v	-1/6	1/2	5/6	-4

The order of the rows is irrelevant, as is the order of the columns; the labels keep track of required information. This tableau is repeated again in the list below. Each tableau appears with rows and columns sorted according to their labels. Some labeling schemes include labels for dual variables. Here only primal variables are used as labels.

The simplex method performs Tucker pivots only for pivot elements that are in the portions of the matrices occupied by A (henceforth termed the "A part"). For the example, all the possible matrices obtainable from the initial tableau by sequences of such pivots are listed here:

T_{123} :

	x_1	x_2	x_3	u
x_4	1	2	3	6
x_5	6	3	1	24
v	1	1	1	0

T_{124} :

	x_1	x_2	x_4	u
x_3	1/3	2/3	1/3	2
x_5	17/3	7/3	-1/3	22
v	2/3	1/3	-1/3	-2

T_{125} :

	x_1	x_2	x_5	u
x_3	6	3	1	24
x_4	-17	-7	-3	-66
v	-5	-2	-1	-24

T_{134} :

	x_1	x_3	x_4	u
x_2	1/2	3/2	1/2	3
x_5	9/2	-7/2	-1/2	15
v	1/2	-1/2	-1/2	-3

T_{135} :

	x_1	x_3	x_5	u
x_2	1/3	2	1/3	8
x_4	-2/3	-3	7/3	-10
v	-1/3	-1	2/3	-8

T_{145} :

	x_1	x_4	x_5	u
x_2	17/7	-1/7	3/7	66/7
x_3	-9/7	3/7	-2/7	-30/7
v	-1/7	-2/7	-1/7	-36/7

T_{234} :

	x_2	x_3	x_4	u
x_1	2	3	1	6
x_5	-9	-17	-6	-12
v	-1	-2	-1	-6

T_{235} :

	x_2	x_3	x_5	u
x_1	1/2	1/6	1/6	4
x_4	3/2	17/6	-1/6	2
v	1/2	5/6	-1/6	-4

T_{245} :

	x_2	x_4	x_5	u
x_1	7/17	-1/17	3/17	66/17
x_3	9/17	6/17	-1/17	12/17
v	1/17	-5/17	-2/17	-78/17

T_{345} :

	x_3	x_4	x_5	u
x_1	-7/9	-1/3	2/9	10/3
x_2	17/9	2/3	-1/9	4/3
v	-1/9	-1/3	-1/9	-14/3

.

These are all the tableaux that can be obtained by pivoting in the A part, beginning from the original one of our example.

Let $X_{(1)} = [x_{1+s}, \ldots, x_{r+s}]^T$, and consider the vector space $\mathcal{W}$ consisting of all vectors

$$X = \begin{pmatrix} X_{(0)} \\ X_{(1)} \\ u \\ v \end{pmatrix} \in R^{r+s+2}$$

such that

$$X_{(1)} = uB - AX_{(0)}, \qquad v = -uD + CX_{(0)}.$$

The problem can now be written as one of maximizing v subject to $u = 1$, the nonnegativity of the x_i's, and $X \in \mathcal{W}$. In these equations, $X_{(1)}$ and v are determined linearly from $X_{(0)}$ and u; so $\mathcal{W}$ is the graph of a linear function from R^{s+1} to R^{r+1}. The tableau determines the linear function. The graph $\mathcal{W}$ is a linear subspace of $R^{r+s+2} = R^{n+2}$ having dimension $s+1$. Given another set of s of the x_i's, unless there is a linear relation relating them, it is possible to transform the equation system in such a way that v and the other r x_i's are determined linearly from those original s, and u. Indeed, Tucker pivoting gives the tableau of such a function. The subspace $\mathcal{W}$ is unchanged; we look at it differently, as the graph of a function of a different set of s variables.

There are only $\binom{n}{s}$ sets of s of the n variables, so there cannot be more than this number of Tucker tableaux. In the nondegenerate case of the example, we do indeed have all $\binom{5}{2} = 10$ distinct tableaux. There is sufficient nondegeneracy that no tableau for the problem has a zero entry in the A, B, or C part. We say that the problem exhibits *total nondegeneracy.* We refer to the set of tableaux in our list as the *complete set of tableaux* for the problem. In general, given an initial tableau for a problem, the complete set of tableaux for the problem consists of all the tableaux that can be derived from the initial one by sequences of pivots on elements in the A part of the tableaux.

By a basic theorem of linear programming applied to the totally nondegenerate case, exactly one of the following three possibilities holds for the complete set of tableaux for the problem:

- There is a unique tableau having only positive entries in the B part and only negative entries in the C part, from which the optimal solution can be obtained (primal and dual feasibility);
- There is a tableau having only positive entries in the B part and having another column for which the entry in the C part is positive and all other entries negative (primal feasibility, dual infeasibility);
- There is a tableau having only negative entries in the C part and having another row for which the entry in the B part is negative and all other entries are positive (primal infeasibility, dual feasibility).

It is not possible for the totally nondegenerate problem to be both primal and dual infeasible.

Let N be the $(r+1) \times (s+1)$ matrix of nonnegative integers whose (i,j)-th entry is the number of tableaux in the complete set that have exactly i negative elements in the B part of the tableaux and exactly j positive elements in the C part. Then $N(0,0)$ is either 0 or 1; it is 1 (in the presence, as we assume, of total nondegeneracy) if and only if our problem has a solution – if and only if the feasible set is nonempty and the objective function is bounded above on it. The *solution tableau* is the one, if it exists, having 0 negative elements in the B part and 0 positive elements in the C part. Let N' be the $(r+1) \times (s+1)$

matrix whose (i,j)-th entry is the number of tableaux in the complete set that have exactly i positive elements in the B part of the tableaux and exactly j negative elements in the C part. For our example, we have

$$N = \begin{pmatrix} 1 & 2 & 2 & 1 \\ 3 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \text{ and } N' = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 3 \\ 1 & 2 & 2 & 1 \end{pmatrix}.$$

Thanks to total nondegeneracy, there is no 0 entry in the B or C part of any tableau, so, for $0 \leq i \leq r$ and $0 \leq j \leq s$, $N'_{i,j} = N_{r-i,s-j}$. Therefore, the sum $Q = N + N'$ has some symmetry: $Q_{i,j} = Q_{r-i,s-j}$, for all i, j.

The sum is:

$$Q = \begin{pmatrix} 1 & 2 & 2 & 1 \\ 3 & 1 & 1 & 3 \\ 1 & 2 & 2 & 1 \end{pmatrix},$$

which actually has even more symmetry: $Q_{i,j} = Q_{r-i,j} = Q_{i,s-j} = Q_{r-i,s-j}$, for all i, j. Our objective is to show that this is true for all totally nondegenerate linear programming problems.

Mulmuley [8] has considered a matrix that is a submatrix of N, in a more special case. Let k be a nonnegative integer, and let X_k denote the set of points satisfying at least $n-k$ of the inequality constraints (so that X_0 is the feasible region). Under the assumptions that (i) X_k is bounded and (ii) X_0 is nonempty, Mulmuley's "h–matrix" is the submatrix of N consisting of its first $k+1$ rows, and he has shown that N satisfies the equations $N(i,j) = N(i,s-j)$ when $0 \leq i \leq k$ and $0 \leq j \leq s$. It is not difficult to show that, when (i) and (ii) are satisfied, the last $k+1$ rows of N consist of 0's, so that the first $k+1$ rows of N and Q coincide.

2 Derived Problems

There are several operations that can be performed on a tableau that lead to different problems having different solutions, but for which the complete set of tableaux for the original problem provide sufficient information. We consider some of these, one-by-one.

Suppose that the row or column of the tableau that is indexed by x_i is multiplied by -1. If it is a row, the new tableau corresponds to a problem identical to the original, except that the sense of the inequality constraint corresponding to that row is reversed; if it is a column, then the variable labeling that column is required to be nonpositive instead of nonnegative. This operation commutes with the operation of performing a pivot on the tableau: If we multiply the row or column indexed by x_i by -1 and then perform a pivot on the entry in row and column labeled by x_j and x_k, the resulting tableau is the same as the tableau that results when we first perform the pivot and then perform the multiplication by -1. This being the case, it

becomes clear that the complete set of tableaux for the new problem can be obtained by replacing each tableau for the original problem by the result upon multiplying its x_i–labeled row or column by -1.

Suppose next that x_i is a row label in our tableau and that row is deleted. This corresponds to removing the corresponding inequality from the constraint set. Pivoting again commutes with this operation, as long as we do not pivot in the row labeled x_i. The complete set of tableaux for the new problem is obtained by replacing each tableaux for the original, having the property that x_i labels a row, by the result upon deleting that row. The tableaux in which x_i labels a column are discarded.

Suppose that x_i is a column label in the original tableau, and that that column is deleted. This amounts to replacing the original problem by one in which x_i is set to zero: $x_i = 0$. The complete set of tableaux for the new problem is obtained by removing the column labeled x_i from each original tableau that has a column labeled x_i. The tableaux in which x_i labels a row are discarded.

Of course, linear programming duality amounts to replacing the original tableau by the negative of its transpose; and this also commutes with the pivot operations.

We term the problems obtained by the operations of multiplication of rows, columns by -1, deletion of rows, columns, the *derived problems* of the original. The duals of the derived problems are the derived problems of the dual. For each derived problem, the complete set of tableaux can be obtained from those of the original problem, as above. If P, Q, and R are pairwise disjoint sets of variables x_i, with $|P| \leq r$ and $|Q| \leq s$, then $\mathcal{D}(P, Q, R)$ is the derived problem for which the complete set of tableaux is obtained by starting with those tableaux in the original set for which the variables in P appear as row labels and those in Q appear among the column labels, and deleting from them the rows labeled by variables in P, deleting the columns labeled by variables in Q, and multiplying rows and columns labeled by elements of R by -1. (We leave it to the reader to sort out the meaning of those derived problems $\mathcal{D}(P, Q, R)$ for which $|P| = r$ or $|Q| = s$. These cases are useful in some of the proofs below.)

We consider the number $\mathcal{F}(p, q, t)$, which is the count of the derived problems $\mathcal{D}(P, Q, R)$ that are feasible and such that the objective function is bounded above on the feasible region, and having $|P| = p$, $|Q| = q$, and $|R| = t$. We call p, q, and t the *parameters* of the derived problem $\mathcal{D}(P, Q, R)$.

Theorem 1. *The number $\mathcal{F}(p, q, t)$ is given by a linear combination of the entries of the matrix N:*

$$\mathcal{F}(p, q, t) = \sum_{\substack{0 \leq i \leq r, \\ 0 \leq j \leq s}} \gamma(p, q, t; i, j) N(i, j).$$

The coefficients $\gamma(p, q, t; i, j)$ are nonnegative integers.

Proof. Given a tableau T, it is easy to see that the number of derived problems having parameters p, q, and t that have T as the solution tableau is a function only of the numbers i and j, which count the negative entries in the B part and positive entries in the C part of T; $\gamma(p, q, t; i, j)$ is this number. Since each derived problem that has a nonempty feasible set on which the objective function is bounded above has a unique solution tableau, the sum yields the total number. □

Suppose T is a tableau, and suppose its B part has i negative entries and its C part has j positive entries, so that it contributes to the number $N(i, j)$. How many derived problems $\mathcal{D}(P, Q, R)$ having parameters p, q, and t have T as the solution tableau? We may construct them all, as follows. Let k and ℓ be nonnegative integers whose sum is t. Choose a set R_1 of k variables from among the i variables, labeling the rows of T that have a negative entry in the B part; there are $\binom{i}{k}$ ways to do this. Next choose a set R_2 of ℓ variables from among the j variables, labeling the columns of T that have a positive entry in the C part. There are $\binom{j}{\ell}$ ways to do this. Let $R = R_1 \cup R_2$. From among the row labels, choose a set P having p variables. Include in P all the variables not in R_1 that label a row having a negative entry in the B part; choose the remaining $p - i + k$ elements of P from among the $r - i$ row labels having a positive entry in the B part. There are $\binom{r-i}{p-i+k}$ ways to do this. Finally, choose a set Q having q elements, including the labels of the $j - \ell$ columns having positive elements in the C part, with the other $q - j + \ell$ elements chosen from the $s - j$ labels of columns having negative elements in the C part. There are $\binom{s-j}{q-j+\ell}$ ways to choose Q in this fashion. It follows that

$$\gamma(p, q, t; i, j) = \sum_{\substack{k+\ell=t,\\ k,\ell\geq 0}} \binom{i}{k}\binom{j}{\ell}\binom{r-i}{p-i+k}\binom{s-j}{q-j+\ell}.$$

There are usually many more values of $\mathcal{F}(p, q, t)$ than entries in the matrix N, so there are linear relations among the values $\mathcal{F}(p, q, t)$.

The next theorem shows that N can be obtained from those values of $\mathcal{F}$ for which: $q = 0$, $0 \leq p \leq r$, and $0 \leq t \leq s$. For such p, q, and r, we have:

$$\gamma(p, q, t; i, j) = \binom{i}{t-j}\binom{r-i}{p-i+t-j}$$

and

$$\mathcal{F}(p, 0, t) = \sum_{\substack{0\leq i\leq r,\\ 1\leq j\leq s}} \binom{i}{t-j}\binom{r-i}{p-i+t-j} N(i, j).$$

Theorem 2. *In Theorem 1, upon restricting p, q, and t to satisfy $q = 0$, $0 \leq p \leq r$, and $0 \leq t \leq s$, and providing a suitable ordering, we obtain a triangular system of linear equations having 1's on the diagonal. Thus it is invertible, so that, given arbitrarily the values of the function $\mathcal{F}$ for p, q, and t so restricted, there is a unique matrix N for which the equations of Theorem 1 are satisfied.*

Proof. We linearly order the set of pairs (a, b) of integers: $(a, b) \prec (a', b')$ if $b < b'$, or $b = b'$ and $a < a'$. Consider $(p, t) = (i, j)$; then, in the expression yielding $\mathcal{F}(p, 0, t)$ above, the coefficient of $N(i, j)$ is 1. Suppose $(p, t) \prec (i, j)$; then this coefficient is 0. □

3 The Proof

Two lemmas will be useful in the proof of the main result.

Let Π denote the problem corresponding to the original tableau, T. Let Π' denote the problem corresponding to the result upon multiplying the column of T labeled u by -1; let Π'' denote the problem corresponding to the result upon multiplying the row of T labeled v by -1; and let Π''' denote the problem corresponding to the result upon performing both of those operations. Since the pivot operations commute with these operations, it is clear that, from the complete set of tableaux for Π, we obtain a complete set for each of the other problems by performing the same operation on all of the tableaux in the complete set Π.

For the corresponding derived problem $\mathcal{D}(P, Q, R)$, function $\mathcal{F}$, and matrix N, we shall now write $\mathcal{D}_\Pi(P, Q, R)$, $\mathcal{F}_\Pi$, and N_Π; and similarly for the other three problems.

Given the sets P, Q, R of variables, pairwise disjoint, let

$$\iota_\Pi(P, Q, R) = \begin{cases} 1 & \text{if } \mathcal{D}_\Pi(P, Q, R) \text{ is feasible and dual-feasible,} \\ 0 & \text{otherwise.} \end{cases}$$

The functions $\iota_{\Pi'}$, $\iota_{\Pi''}$, and $\iota_{\Pi'''}$ are similarly defined.

Lemma 1. *For each choice of P, Q, and R, we have*

$$\iota_\Pi(P, Q, R) + \iota_{\Pi'''}(P, Q, R) = \iota_{\Pi'}(P, Q, R) + \iota_{\Pi''}(P, Q, R).$$

Proof. We make the argument in the case $P = Q = R = \emptyset$; however it will be clear that the same argument applies in general. Note that $\mathcal{D}(\emptyset, \emptyset, \emptyset)$ is simply the original problem, Π.

We designate the following statement by $B+$: There is a tableau in the complete set of tableau for Π which that has only positive elements in the B part. Similarly define statements, $B-$, $C+$, and $C-$.

Either $B+$ and $B-$ both hold, or $C+$ and $C-$ both hold; for otherwise one of the four problems Π, Π', Π'', or Π''' is neither primal feasible nor dual feasible. Then it is easy to see that, depending upon whether neither, exactly one, or both of the other two statements hold, both sides in the above equation must be 0, 1, or 2. (Actually, 2 is not possible; but this fact is not needed here.) □

Lemma 2. *We have*

$$\mathcal{F}_\Pi(p, q, t) + \mathcal{F}_{\Pi'''}(p, q, t) = \mathcal{F}_{\Pi'}(p, q, t) + \mathcal{F}_{\Pi''}(p, q, t).$$

Proof. Observe that

$$\mathcal{F}_{\Pi}(p,q,t) = \sum_{(P,Q,R)} \iota_{\Pi}(P,Q,R),$$

where the summation extends over triples (P,Q,R) of pairwise disjoint sets of variables for which $|P| = p$, $|Q| = q$, and $|R| = t$. The same holds for the other three problems, so this lemma is a consequence of the preceding lemma. □

Theorem 3. *Letting* $Q(i,j) = N(i,j) + N(r-i, s-j)$, *we have* $Q(i,j) = Q(r-i,j) = Q(i,s-j) = Q(r-i,s-j)$.

Proof. Equivalently, we must show that $N_{\Pi}(i,j) + N_{\Pi}(r-i,s-j) = N_{\Pi}(r-i,j) + N_{\Pi}(i,s-j)$. It is clear that:

$$N_{\Pi'}(i,j) = N_{\Pi}(r-i,j),$$

$$N_{\Pi''}(i,j) = N_{\Pi}(i,s-j),$$

and

$$N_{\Pi'''}(i,j) = N_{\Pi}(r-i,s-j).$$

The equations of Theorem 1 also give

$$\mathcal{F}_{\Pi}(p,q,t) + \mathcal{F}_{\Pi'''}(p,q,t) - \mathcal{F}_{\Pi'}(p,q,t) - \mathcal{F}_{\Pi''}(p,q,t)$$

in terms of

$$N_{\Pi}(i,j) + N_{\Pi'''}(i,j) - N_{\Pi'}(i,j) - N_{\Pi''}(i,j).$$

However, by Lemma 2,

$$\mathcal{F}_{\Pi}(p,q,t) + \mathcal{F}_{\Pi'''}(p,q,t) - \mathcal{F}_{\Pi'}(p,q,t) - \mathcal{F}_{\Pi''}(p,q,t) = 0,$$

so by Theorem 2 we have

$$N_{\Pi}(i,j) + N_{\Pi'''}(i,j) - N_{\Pi'}(i,j) - N_{\Pi''}(i,j) = 0,$$

or equivalently,

$$N_{\Pi}(i,j) + N_{\Pi}(r-i,s-j) = N_{\Pi}(r-i,j) + N_{\Pi}(i,s-j).$$

□

4 Notes

An interesting open problem is that of characterizing the counting matrices N. Certainly, these matrices satisfy the following conditions: The entries are nonnegative integers; the sum of the entries is $\binom{n}{s}$; and, for the matrix Q having $Q(i,j) = N(i,j) + N(r-i,s-j)$, one has $Q(i,j) = Q(r-i,j) = Q(i,s-$

$j) = Q(r - i, s - j)$. These conditions are far from sufficient to characterize the possible matrices.

The results in this paper hold more generally for uniform oriented matroids. With minor changes, the proofs hold in the more general setting. It is unknown whether or not one obtains additional matrices N from the "linear programming problems" derived in the setting of uniform oriented matroids.

The problem of characterizing the analogous matrices N, when total non-degeneracy is not assumed, is also open.

References

1. L. J. Billera and C. W. Lee. A proof of the sufficiency of McMullen's conditions for f-vectors of simplicial convex polytopes. *J. Combinat. Theory, Ser. A*, 31:237-255, 1981.
2. S. Gass. *Linear Programming Methods and Applications*. McGraw-Hill, 1958. (Fourth edition, 1975.)
3. A. J. Goldman and A. W. Tucker. Theory of linear programming. In *Linear Inequalities and Related Systems, Annals of Mathematics Studies* **38**, eds. H. W. Kuhn and A. W. Tucker, Princeton University Press, 1956.
4. J. Lawrence. Total polynomials of uniform oriented matroids. *European Journal of Combinatorics*, 21:3-12, 2000.
5. J. Lawrence. Mutation polynomials and oriented matroids. *Discrete and Computational Geometry*, 24:365-389, 2000.
6. P. McMullen. The maximum numbers of faces of a convex polytope. *Mathematika*, 17:179-184, 1970.
7. P. McMullen. The numbers of faces of simplicial polytopes. *Israel J. Math.*, 9:559-570, 1971.
8. K. Mulmuley. Dehn–Sommerville relations, upper bound theorem, and levels in arrangements. *Proceedings of Ninth Annual Symposium on Computational Geometry*: San Diego, California, May 12–21, 240-246, 1993.
9. R. P. Stanley. The numbers of faces of a simplicial convex polytope. *Adv. Math.*, 35:236-238, 1980.

Toward Exposing the Applicability of Gass & Saaty's Parametric Programming Procedure

Kweku-Muata Osei-Bryson

Department of Information Systems &
The Information Systems Research Institute
Virginia Commonwealth University
Richmond, VA 23284, USA
kweku.muata@isy.vcu.edu

Summary. In this paper we discuss some applications of Gass & Saaty's parametric programming procedure (GSP^3). This underutilized procedure is relevant to solution strategies for many problem types but is often not considered as a viable alternate approach. By demonstrating the utility of this procedure, we hope to attract other researchers to explore the use of GSP^3 as a solution approach for important problems in operations research, computer science, information systems, and other areas.

Key words: Parametric programming; integer programming; Lagrangian relaxation; multi-objective programming; multi-criteria decision making; simplex method; clustering; data mining; database systems.

1 Introduction

In this paper we present an overview of some applications of the parametric programming procedure of Gass and Saaty [16]. This procedure is relevant to solution strategies for many problem types but is often underutilized. In several cases it appears that it was not considered as an alternate approach, and the lack of its use was not justified. Our initial interest in applications of this procedure involved the Lagrangian relaxation of a network problem with a side constraint. This problem is a special case of the parametric linear programming problem (P^3); yet up to the mid-1980's it was not usually recognized as such and so was usually addressed using other techniques such as sub-gradient optimization. Bryson [6] demonstrated that this problem could be effectively addressed by the parametric programming procedure of Gass and Saaty (GSP^3). Since then we have observed other problems that could be addressed by this procedure. For example, Stanfel [32] addressed the sequential partitioning problem, which is a special case of the clustering problem, but did not recognize that this problem could also be addressed by GSP^3 [19]. More recently in Osei-Bryson and Joseph [30], we demonstrated that some technical information systems problems could be effectively addressed by GSP^3. In general we have seen special cases of P^3 in management science (MS), operations research (OR), computer science (CS), and information systems (IS) that were addressed by sub-optimal techniques. While we are not certain as to why this

situation occurs, some possible reasons include the researcher's unawareness of GSP^3 and/or lack of recognition of their research problem as being just a special case of P^3. Although in an earlier work [7], we discussed some applications of GSP^3, this situation suggests to us that it might be appropriate to re-tell at least part of the story on the wide applicability of GSP^3.

2 Gass & Saaty's Parametric Programming Procedure (GSP^3)

The parametric programming procedure of Gass and Saaty (GSP^3) solves the following linear programming problem for every value of the scalar multiplier w:

$$P_1: \text{Min } \{Z_0(X) + wZ_1(X) \mid AX = b;\ X \geq 0\}$$

where $Z_0(X)$ (e.g., $CX = \sum_j c_j x_j$) and $Z_1(X)$ (e.g., $DX = \sum_j d_j x_j$) are linear functions of X, and w is a scalar multiplier. Using the simplex method, the parametric programming procedure first solves problem P_1 for $w = w_0$, where w_0 is some arbitrary small number. The result is that either there is some efficient extreme-point solution X_0, or that there is no finite minimum for $w = w_0$.

Let $J_{(k)}$ be the index set of basic variables associated with X_k, and let the reduced costs have the form $(\underline{c}_j + w\underline{d}_j)$, where $\underline{c}_j$ and $\underline{d}_j$ are the reduced costs in terms of $Z_0(X)$ and $Z_1(X)$, respectively.

For optimality we require that $(\underline{c}_j + w\underline{d}_j) \leq 0$. Therefore, X_k is the optimal solution of problem P_1 for the interval $[w_{k(L)}, w_{k(U)}]$ where $w_{k(L)}$, $w_{k(U)}$ are determined as follows:

$$w_{k(L)} = -(\underline{c}_{k(L)}/\underline{d}_{k(L)}) = \text{Max } \{-\underline{c}_j/\underline{d}_j: \underline{d}_j < 0;\ j \notin J_{(k)}\}; \text{ or}$$
$$= -\infty \quad \text{if} \quad \underline{d}_j \geq 0 \ \forall\ j \notin J_{(k)}.$$

$$w_{k(U)} = -(\underline{c}_{k(U)}/\underline{d}_{k(U)}) = \text{Min } \{-\underline{c}_j/\underline{d}_j: \underline{d}_j \geq 0;\ j \notin J_{(k)}; \text{ or}$$
$$= +\infty \quad \text{if} \quad \underline{d}_j < 0 \ \forall\ j \notin J_{(k)}.$$

The procedure is terminated if either of the following conditions occurs: (a) $w_{k(U)} = +\infty$, (b) $w_{k(U)}$ is finite but all the corresponding $a_{i,k(U)} \leq 0$. Otherwise, the simplex method introduces $x_{k(U)}$ into the basis and eliminates the basic variable in the usual manner. Gass and Saaty [16] established that the resulting basis yields a minimum for at least one value of w, and that if $[w_{k+1(L)}, w_{k+1(U)}]$ is the interval for which the resulting basis yields a minimum then $w_{k(U)} = w_{k+1(L)}$.

Thus, the parametric programming procedure generates the set of efficient, extreme points that solve the single parameter LP problem for all w such that $-\infty \leq w \leq +\infty$ and will identify the parametric interval $[w_{k(L)}, w_{k(U)}]$ that is associated with each such extreme point X_k.

The first step of GSP^3 involves solving the LP problem P_1 with $w = 0$. It is known that LP problems can be solved in polynomial time [21, 5], and so this first step can be done in polynomial time. Given this initial solution, the next solution

for problem P_1 can be obtained using Gass & Saaty's parametric programming rule for identifying the relevant variable to leave the basis [16]. Succeeding solutions would be obtained in a similar manner. Thus generation of all the solutions of P_1 can be done in polynomial time.

3 Sample Applications

3.1 Lagrangian Relaxation: Network Problem with a Side Constraint

Lagrangian relaxation is usually applied to a problem that would be easy to solve except for some complicating conditions. One such problem is the network problem with a side-constraint [6]. Examples of this problem include the shortest path problem and the constrained assignment problem [1]. This problem can be stated as

$$P_{Net\_SC}: \text{Min } \{\textstyle\sum_j c_j x_j \mid AX = I; \sum_j d_j x_j = v; x_j \in \{0,1\}\},$$

where A is a totally unimodular matrix because of the network structure of the problem. If the side constraint (i.e., $\sum_j d_j x_j = v$) was not present, then the problem Min $\{\sum_j c_j x_j \mid AX = I; x_j \in \{0,1\}\}$ could be easily formulated and solved as a linear programming problem Min $\{\sum_j c_j x_j : AX = I; 0 \le x_j \le 1\}$. This follows from the fact that A is a totally unimodular matrix [29]. Given the inconvenience of the side-constraint, Lagrangian relaxation is often applied as part of a solution strategy for the network problem with a single side-constraint. The Lagrangian relaxation problem is expressed as Max $\{L(w) \mid w > 0\}$, where for each $w > 0$, $L(w)$ is obtained by solving the following problem:

$$P_{LR\_IP}: L(w) = \text{Min } \{\textstyle\sum_j c_j x_j + w(v - \sum_j d_j x_j) \mid AX = I; x_j \in \{0,1\}\}$$

Because of the total unimodularity property of A, problem P_{LR_IP} can be restated as the following parametric linear programming problem:

$$P_{LR\_PP}: L(w) = \text{Min } \{\textstyle\sum_j c_j x_j + w(v - \sum_j d_j x_j) \mid AX = I; 0 \le x_j \le 1\}$$

The reader may observe that *GSP*[3] can be used to automatically generate the optimal solutions of problem P_{LR_PP}, and thus problem P_{LR_IP} for all positive values of w. The optimal Lagrangian multiplier that provides the maximum value of $L(w)$ would be associated with the pair of non-degenerate intervals $[w_{k-1}, w_k]$ and $[w_k, w_{k+1}]$ such that $L(w_{k-1}) \le L(w_k)$, $L(w_{k+1}) \le L(w_k)$, $w_{k-1} < w_k < w_{k+1}$. Since $L(w)$ is a concave function and w_k provides a local maximum, then it also provides the global maximum for $L(w)$, and so *GSP*[3] can be used to find the optimal solution of the Lagrangian relaxation of problem P_{Net_SC} [6, 7].

3.2 Simplex Method: Preventing Cycling

The simplex method is said to stall when one or more pivot steps result in no change in the value of the objective function. Magnanti and Orlin [25] and Dantzig [10] presented methods based on parametric programming for dealing with this difficulty.

Let $CX = \sum_j c_j x_j$ be the original objective function. Let $DX = \sum_j d_j x_j$ be an additional objective function that is used to parameterize the cost function with the coefficients defined as follows:

$$d_j = 0 \qquad \forall\, j \in J_{(0)}$$
$$d_j = \|A_{\cdot j}\|(1 + \varepsilon_j) \qquad \forall\, j \notin J_{(0)}$$

where $J_{(0)}$ is index set of basis columns in the initial solution, $\|A_{\cdot j}\|$ is a norm or any positive number, and ε_j is chosen from a table of random numbers such that $0 < \varepsilon_j < 1$. The parametric objective function is: $\sum_j c_j x_j + w\sum_j d_j x_j$.

Let $B^{(k-1)}$ be the basis at the start of iteration k, and $C_{B(k-1)}$ and $D_{B(k-1)}$ be the components of C and D corresponding to the basis columns in $J_{(k-1)}$. It follows that the corresponding reduced cost vectors are:

$$\underline{C} = (C_{B(k-1)}B^{(k-1)}A - C)$$
$$\underline{D} = (D_{B(k-1)}B^{(k-1)}A - D) = -D$$

So at iteration k the non-basic variable x_t that is chosen to enter the basis is the one that provides the following maximum:

$$-(\underline{c}_t/\underline{d}_t) = \text{Max}\ \{-\underline{c}_j/\underline{d}_j: \underline{c}_j \geq 0, \underline{d}_j < 0; j \notin J_{(k-1)}\}$$

where the reduced cost is of the form $(\underline{c}_j + w_k\underline{d}_j)$ with $w_k > 0$. If all $\underline{c}_j < 0$, the algorithm terminates with the current solution as optimal. Dantzig [10] stated that at iteration k, the choice of x_t will be unique with probability 1 because of the random selection of the ε_j's, and that the algorithm will converge in a finite number of steps.

3.3 Multi-Objective Programming: Three Objective LP Problem

The three objective linear programming problem can be expressed as:

$$\text{Min}\ \{\ [Z_1(X), Z_2(X), Z_3(X)]\ |\ X \in S\}$$

where S is the convex set of feasible solutions, and $Z_1(X)$, $Z_2(X)$ and $Z_3(X)$ are linear functions of X. This problem has a set of non-dominated solutions S*, where a solution X^* is said to be non-dominated if there is no other feasible solution X_k such that $[Z_1(X_k), Z_2(X_k), Z_3(X_k)] \leq [Z_1(X^*), Z_2(X^*), Z_3(X^*)]$ and for at least one p $\in \{1, 2, 3\}$ the inequality $Z_p(X_k) < Z_p(X^*)$ holds.

For a given three-objective LP problem there are three associated two-objective LP problems:

$$MP_{(p,q)} : \text{Min } \{ [Z_p(X), Z_q(X)] : X \in S\} \qquad (p,q) \in \{(1,2), (2,3), (3,1)\}$$

Each problem $MP_{(p,q)}$ has a set of non-dominated solutions $S^*_{(p,q)}$ where a solution X^* is said to be non-dominated if there is no other feasible solution X_k such that $[Z_p(X_k), Z_q(X_k)] \leq [Z_p(X^*), Z_q(X^*)]$ and either $Z_p(X_k) < Z_p(X^*)$ or $Z_q(X_k) < Z_q(X^*)$ holds.

A useful relationship between the three objective LP problem and the associated two-objective LP problem is that each unique non-dominated solution in each $MP_{(p,q)}$ is also non-dominated in the three-objective problem. For the case where multiple non-dominated solution of an $MP_{(p,q)}$ provides the same value of the vector $[Z_p(X_k), Z_q(X_k)]$, the vector that provides the smallest value of $Z_r(X_k)$, where $r = \{1, 2, 3\} - \{p, q\}$, is a non-dominated solution of the three-objective problem.

The non-dominated solutions of problem $MP_{(p,q)}$ can be obtained by solving the following parametric linear programming problem: $PLP_{(p,q)}(w_p)$: Min $\{w_pZ_p(X) + (1- w_p)Z_q(X)] \mid X \in S\}$ for all $0 \leq w_p \leq 1$. This problem can be easily reformulated as the parametric linear programming problem. It follows that for each problem $MP_{(p,q)}$, $S^*_{(p,q)}$ the corresponding set of non-dominated solutions can be obtained by applying *GSP*[3]. It should also be noted that once $S^*_{(p,q)}$ has been generated then the starting solution for $S^*_{(q,r)}$ is either already present in $S^*_{(p,q)}$ or can be easily obtained.

Bryson [8] therefore proposed the following approach for generating:

1. Use GSP^3 to generate first $S^*_{(1,2)}$, then $S^*_{(2,3)}$, and finally $S^*_{(3,1)}$.
2. Generate $X^* \in S^*_{(3)}$ that are adjacent to at least one $X^* \in \{S^*_{(1,2)} \cup S^*_{(2,3)} \cup S^*_{(3,1)}\}$ where $S^*_{(3)} = S^* - \{S^*_{(1,2)} \cup S^*_{(2,3)} \cup S^*_{(3,1)}\}$.
3. Generate the other $X^* \in S^*_{(3)}$ that are adjacent to non-dominated solutions in $S^*_{(3)}$.

3.4 Multi-Criteria Decision Making: Estimating Overall Preference Values

MCDM problems are said to involve the prioritization of a set of alternatives in situations that involve multiple, sometimes conflicting criteria. MCDM problems may be addressed informally or formally. Various formal techniques have been proposed including the weighing model, which is a popular formulation for MCDM. Another formulation [27] involves viewing the overall preference value for each alternative as the sum of its partial values, where the partial value of criterion k for alternative j is the contribution of criterion k to the overall value of alternative j. Thus alternative j is assumed to have some unknown partial value T_{jk} in criterion k and the overall preference value V_j of alternative j is the sum of the partial values, which can be expressed as:

$$(1) \qquad V_j - T_{j1} - T_{j2} - \ldots - T_{jn} = 0 \qquad \forall\ j \in J$$

Moy et al. [27] developed a linear goal programming model to determine estimates of the partial values T_{jk} and the overall values V_j based on selected sets of pairwise comparisons as described below.

Pairwise Comparisons of Overall Preferences of Selected Pairs of Alternatives:

Let Ω denote the set of selected pairs of alternatives that are compared with regards to their overall preferences. For each (i, j) in Ω, the evaluator provides t_{ij}, an estimate of the lower bound of ratio (V_i/V_j). If all such comparisons were consistent, then we would have $V_i - t_{ij}V_j \geq 0$. However, because of the possibility of inconsistencies, deviational variables ε^+_{ij} and ε^-_{ij} are introduced to give the following equation:

$$(2) \qquad V_i - t_{ij}V_j + \varepsilon^+_{ij} - \varepsilon^-_{ij} = 0 \qquad \forall\ (i, j) \in \Omega$$

Pairwise Comparisons of Overall Preferences of Selected Pairs of Partial Values:

Let Θ denote the set of selected pairs of partial values that are compared. For each (j, k, i, p) in Θ, the evaluator provides t_{jkip}, an estimate of the lower bound of ratio (T_{jk}/T_{ip}). If all such comparisons were consistent, then we would have $T_{jk} - t_{jkip}T_{ip} \geq 0$. However, because of the possibility of inconsistencies, deviational variables δ^+_{jkip} and δ_{jkip} are introduced:

$$(3) \qquad T_{jk} - t_{jkip}T_{ip} + \delta^+_{jkip} - \delta = 0 \qquad \forall\ (j, k, i, p) \in \Theta$$

$\sum_{(i,j)} (\varepsilon^+_{ij} + \varepsilon^-_{ij})$ is a measure of the poorness of fit of the pairwise comparisons in Ω. $\sum_{(j,k,i,p)\in\Theta} (\delta^+_{jkip} + \delta_{jkip})$ is a measure of the poorness of fit of the pairwise comparisons in Θ.

Generating Estimates of the Overall Preferences and Partial Values:

To prevent a trivial solution to the system of equations (1) – (3), Moy et al. [27] included the constraint:

$$(4) \qquad \sum_j \sum_k T_{jk} \geq 1$$

They also argued that e^-_{ij} should be penalized heavier than ε^+_{ij} because "ε^-_{ij} reflects the inconsistencies between the input and estimated data", and so proposed that $\sum_{(i,j)\in\Omega}(\varepsilon^+_{ij} + \varepsilon^-_{ij})$ be replaced by $\sum_{(i,j)\in\Omega}(e^+_{ij} + he^-_{ij})$, where $h > 1$ is specified by the researcher. Similarly they proposed that $\sum_{(j,k,i,p)\in\Theta}(\delta^+_{jkip} + \delta_{jkip})$ be replaced by $\sum_{(j,k,i,p)\in\Theta}(\delta^+_{jkip} + h\delta_{jkip})$, where $h > 1$.

Moy et al. [27] proposed that the unknown partial values could be identified by solving the following LP problem:

P_{P_V}: Min $\{C_1(\sum_{(i,j)\in\Omega}(\varepsilon^+_{ij} + h\varepsilon^-_{ij})) + C_2(\sum_{(j,k,i,p)\in\Theta}(\delta^+_{jkip} + h\delta_{jkip})) \mid (1) - (4)\}$
where all variables are non-negative.

The difficulty is that the pair values of (C_1, C_2) are also unknown, and so Moy et al. proposed that "Different values of C_1 and C_2 can be tried on the training sample and the pair of C_1 and C_2 values which gives good results can be chosen. Good results should have high correlations between respondent's input preferences and the estimated preferences or have the estimated first choice (the number one ranked alternative) matches with that of the respondent in the training sample". They thus used the following ratios for (C_2/C_1) in their experiments: (1/5), (1/1), (3/1), and (5/1). The reader may observe that if objective function of problem P_{P_V} is restated as:

$$\sum_{(i,j)\in\Omega}(\varepsilon^+_{ij} + h\varepsilon^-_{ij}) + (C_2/C_1)(\sum_{(j,k,i,p)\in\Theta}(\delta^+_{jkip} + h\delta_{jkip}))$$

then it is clear that P_{P_V} is just a special case of problem of P^3 and so the full set of weights $w = (C_2/C_1)$ would be generated by the GSP^3.

3.5 Clustering

Clustering involves the partitioning of a set of objects into a useful set of mutually exclusive clusters. There are numerous algorithms available for doing clustering. They may be categorized in various ways including: hierarchical [28, 33] or partitional [26], deterministic or probabilistic [4], hard or fuzzy [3, 11]. In some cases there is a natural ordering of the objects in the dataset while in the more general case there is no such natural ordering. For the former case sequential partitioning algorithms may be used to generate the optimal segmentation while in the latter case heuristics (e.g., k-Means) may be used to find 'good' segmentations.

3.5.1 Sequential Partitioning Problem

The sequential clustering problem involves the following assumptions:

- The objects are ordered sequentially and the final partition must maintain that sequential ordering. This sequential ordering may be based on physical or ordinal relationships among the objects.
- The overall clustering objective function can be formulated as the sum of the cost of the clusters in the partition.
- The set of relevant candidate clusters is limited due to admissibility constraints (e.g., maximum number of objects per cluster).

There are several practical examples of the sequential partitioning problem. Joseph and Bryson [18] presented a parametric programming-based approach for

addressing the computer program segmentation problem. In a recent work Osei-Bryson and Joseph [30] demonstrated that three technical Information Systems Problems (i.e., the attribute discretization for decision tree induction, the design of lock tables for database management systems, computer program segmentation) were special cases of the sequential partitioning problem that were previously addressed either by greedy sub-optimal heuristics [24, 9] or instance-specific procedures [22], and proposed a common procedure for addressing these problems that involve the use of *GSP*^{3}.

The fact that the clusters are restricted to a linear ordering of objects means that each cluster can be represented as a m-dimensional column vector $\mathbf{A_j} = \{a_{ij}\}$ where $a_{ij} = 1$ if object "i" is assigned to cluster "j" and $a_{ij} = 0$ otherwise, with the 1's in consecutive rows. Given this property, then a matrix $\mathbf{A} = \{\mathbf{A_j}\}$ would be said to have the total unimodularity property [29].

Let $\mathbf{X} = \{x_j\}$ be a binary m row column vector such that $x_j = 1$ if cluster "j" is in the optimal partition and $x_j = 0$ otherwise, then each of our three problem types can be formulated as the following linear integer programming problem:

$$\mathbf{P_1(g)}: Z_{Opt}(g) = \text{Min } \{\textstyle\sum_{j\in J} c_j x_j \mid \mathbf{AX} = \mathit{1}; \sum_{j\in J} x_j = g; x_j \in \{0,1\}\ \forall\, j \in J\}$$

where the matrix **A** consists of those column vectors that represent clusters that satisfied the relevant admissibility constraints (e.g., minimum and maximum number of objects per cluster), *1* is a column vector with the value 1 in each row, and g is the number of clusters. It should be noted that although the total number of possible columns for matrix **A** is given by $[(m^2 + m)/2]-1$, the admissibility constraints will usually significantly reduce this number.

While integer programming (IP) problems are in general difficult to solve optimally, linear programming (LP) problems can be solved optimally in polynomial time [5, 21]. Further based on the total unimodularity property of the matrix **A**, the following parametric LP problem is guaranteed to have integer solutions:

$$\mathbf{P_2(w)}: \text{ Min } \{\textstyle\sum_{j\in J}(c_j + w)x_j \mid \mathbf{AX} = \mathit{1}; x_j \geq 0\ \forall\, j \in J\}$$

where w (≥ 0) is a parameter the variations of which can be associated with different values of g [20]. This parametric linear programming problem can be solved by *GSP*^{3} and optimal partitions can be automatically generated for different values of g without requiring the user to specify a value.

Optimal solutions to $\mathbf{P_2}(w)$ are referred to as being *w-efficient partitions*, and are known to be integer-valued solutions of $\mathbf{P_1}(g)$ for the corresponding value of g. Solution of $\mathbf{P_2}(w)$ may not provide solutions of $\mathbf{P_1}(g)$ for all values of g. Those solutions of $\mathbf{P_1}(g)$ that are not generated in the solutions of $\mathbf{P_2}(w)$ for all $w \geq 0$ are referred to as being *w-inefficient partitions*. Given two consecutive w-efficient partitions with g_a and g_c intervals respectively, the decision-maker (e.g., knowledge engineer, database engineer) can find the lower bound on the value of the objective function of a partition with g_b intervals where $g_a < g_b < g_c$ using the following formula: $Z_{Opt}(g_b) > Z_{LB}(g_b) = Z_{Opt}(g_a) + (Z_{Opt}(g_c) - Z_{Opt}(g_a))*(g_b - g_a)/(g_c - g_a)$. If the

decision-maker is still interested in exploring the optimal partition of size g_b, then problem $\mathbf{P}_1(g_b)$ can be solved. It should be noted that in solving $\mathbf{P}_1(g_b)$ initial bounds would be based on properties of the immediate surrounding *w-efficient* partitions, and the starting solution vector would be that associated with one of the immediate surrounding *w-efficient* partitions. Thus, exploration and resolution of the sequential set partitioning problem can be successfully and efficiently framed within the context of identifying the *w-efficient* partitions. Therefore, the decision-maker is not required to make a premature decision on partition size (e.g., number of pages, number of intervals, number of hash tables) but is rather given the opportunity to explore the implications of different partition sizes.

Parametric Programming-based Procedure for Sequential Partitioning (P^3_SP):

Step 1:
a) Define the elementary objects and obtain the relevant data.
b) Specify the ordering of the items.
c) Define relevant admissibility constraints on clusters.

Step 2: Generate Potential Clusters
Generate clusters sequentially, and exclude those that do not satisfy the admissibility constraints from future consideration.

Step 3: Generate Optimal Partitions
a) Beginning with $w = 0$, apply *GSP*^{3} to solve problem $\mathbf{P_2(w)}$ **for** all $w \geq 0$.
b) For each solution obtained in part 3a, properties of the optimal *w-efficient* partition including its optimal objective function value $Z_{Opt}(g)$ and the corresponding value of g are recorded.
c) For each g_b that is associated with a *w-inefficient* partition, $Z_{LB}(g_b)$ the associated lower bound is computed.
d) For each *w-inefficient* partition of size g_b, the decision-maker examines the corresponding $Z_{LB}(g_b)$. If the decision-maker has an interest in exploring the properties of the optimal partition of size g_b, then problem $P_1(g_b)$ is solved.

Step 4: Analysis
The decision-maker conducts trade-off analysis using the results obtained in step 3.

The reader should note that even in the extreme pathological case where solving $\mathbf{P_2(w)}$ for all $w > 0$ results in w-efficient partitions only for sizes $g = 1$ and $g = m$, this procedure could still be used to generate partitions for other values of g that might be of interest to the decision-maker.

3.5.2 General Clustering Problem

For the general clustering problem there is no natural ordering of the objects, which is a fundamental assumption of the sequential clustering algorithms. In a recent work, Osei-Bryson and Inniss [31] proposed an approach for extending to P^3_SP address the general clustering problem.

1. Apply Procedure(s) for Generating Orderings of the set of Objects.
2. Use P^3_SP to generate Optimal Partitionings for each Ordering generated in Step 1.
3. Given the results of Step 2, select the 'best' Partitioning for each partition size (g).

The first step of this procedure requires one or more methods for generating orderings of the objects. Osei-Bryson and Inniss [31] presented two approaches: the FastMap algorithm of Faloutsos and Lin [15], and an approach based on spectral analysis [17, 2, 13, 14].

Osei-Bryson and Inniss [31] noted that this approach has several advantages including:

- Explicitly accommodating various cluster admissibility criteria including cluster size, maximum dissimilarity between any pair of objects in each cluster, and variance reduction. None of the traditional approaches accommodate all of these admissibility criteria, although some commercial implementations allow the user to specify the minimum cluster size.
- Explicitly accommodating any separable objective function rather than 'hoping' that feasibility and a 'good' solution will be achieved as a consequence of the given approach to grouping objects (e.g., k-Means only considers the distance to the closest cluster mean).
- Flexibility in the assignment of objects to clusters. So unlike agglomerative hierarchical algorithms (e.g., Ward), the assignments of objects to clusters in the partition of size $(g+1)$ are not restricted by the assignments made in the partition of size g.
- Not requiring the user to make a premature decision on partition size, but rather offering the opportunity to explore the implications of different partition sizes. It should be noted that while traditional hierarchical approaches also offer the user the opportunity to evaluate different partition sizes, there is no guarantee that for a given partition size each generated cluster satisfies the relevant admissibility criteria.

5 Conclusion

GSP^3 is an important solution resource for different problems in MS, OR, CS, and IS, but over 50 years after its initial presentation to the research community [16], it

remains relatively underutilized, and its applicability often unrecognized by even OR/MS researchers. In this paper we have presented a sample of problem types that can be addressed by approaches that involve the use of GSP^3. Its development in 1955 is a testament to the intellectual power of Gass and Saaty, at that time two young researchers. Its wide applicability is another piece of evidence of the richness of the legacy of these two researchers. We present this paper with the hope that it will attract other researchers to explore the use of GSP^3 as part of the solution approaches for on other important problems in MS, OR, CS, IS, and other areas.

References

1. V. Aggarwal. A Lagrangean-Relaxation Method for the Constrained Assignment Problem. *Computers & Operations Research,* 12:97-106, 1985.
2. C. Alpert and S. Yao. Spectral Partitioning: The More Eigenvectors, the Better. *32nd ACM/IEEE Design Automation Conference*, 195-200, 1995.
3. J. Bezdek. Pattern Recognition with Fuzzy Objective Function Algorithms. Plenum Press, New York, NY, 1981.
4. H. Bock. Probability Models in Partitional Cluster Analysis. *Computational Statistics and Data Analysis,* 23:5-28, 1996.
5. K.-H. Borgwardt. The Average Number of Steps Required by the Simplex Method is Polynomial. *Zeitchrift fur Operations Research,* 26:157-177, 1982.
6. N. Bryson. Parametric Programming and Lagrangian Relaxation: The Case of the Network Problem with a Single Side-Constraint. *Computers & Operations Research,* 18:129-140, 1991.
7. N. Bryson. Applications of the Parametric Programming Procedure. *European Journal of Operational Research,* 54:66-73, 1991.
8. N. Bryson. Identifying the Efficient Extreme-Points of the Three-Objective Linear Programming Problem. *Journal of the Operational Research Society,* 44:81-85, 1993.
9. M-S. Chen and P. Yu. Optimal Design of Multiple Hash Tables for Concurrency Control. *IEEE Transactions on Knowledge and Data Engineering,* 9:384-390, 1997.
10. G. Dantzig. Making Progress during a Stall in the Simplex Algorithm. Technical Report SOL 88-5, Stanford University, Stanford, CA, 1988.
11. R. Dave. Generalized Fuzzy C-Shells Clustering and Detection of Circular and Elliptic Boundaries. *Pattern Recognition*, 25:713-722, 1992.
12. P. Denning. Working Sets Past and Present. *IEEE Transactions on Software Engineering,* 6:64-84, 1980.
13. I. Dhillon. Co-Clustering Documents and Words Using Bipartite Spectral Graph Partitioning. *Proceedings of the 7th ACM SIGKDD*, 269-274, 2001.
14. C. Ding and X. He. Linearized Cluster Assignment via Spectral Ordering. *ACM Proceedings of the 21st International Conference on Machine Learning,* 30, 2004.

15. C. Faloutsos and K-I. Lin. FastMap: A Fast Algorithm for the Indexing, Data Mining, and Visualization of Traditional and Multimedia Datasets. *ACM SIGMOD Proceedings*, 163-174, 1995.
16. S. Gass and T. Saaty. The Computational Algorithm for the Parametric Objective Function. *Naval Research Logistics Quarterly,* 2:39-45, 1955.
17. L. Hagen and A. Kahng. New Spectral Methods for Ratio Cut Partitioning and Clustering. *IEEE Trans. on Computed Aided Design,* 11:1074-1085, 1992.
18. A. Joseph and N. Bryson. Partitioning of Sequentially Ordered Systems Using Linear Programming. *Computers & Operations Research,* 24:679-686, 1997a.
19. A. Joseph and N. Bryson. "Parametric Programming and Cluster Analysis", *European Journal of Operational Research* 111:582-588, 1999.
20. A. Joseph and N. Bryson. W-Efficient Partitions and the Solution of the Sequential Clustering Problem. *Annals of Operations Research: Nontraditional Approaches to Statistical Classification,* 74:305-319, 1997b.
21. N. Karmarkar. A New Polynomial-Time Algorithm for Linear Programming. *Combinatorica,* 4:373–395, 1984.
22. B. Kernighan. Optimal Sequential Partitions of Graphs. *Journal of the Association for Computing Machinery,* 18:34-40, 1971.
23. L. Kurgan and K. Cios. CAIM Discretization Algorithm. *IEEE Transactions on Knowledge and Data Engineering,* 16:145-153, 2004.
24. H. Liu and R. Setiono. Feature Selection by Discretization. *IEEE Transactions on Knowledge and Data Engineering*, 9:642-645, 1997.
25. T. Magnanti and J. Orlin. Parametric Linear Programming and Anti-Cycling Rules. *Mathematical Programming,* 41:317-325, 1988.
26. J. McQueen. Some Methods for Classification and Analysis of Multivariate Observations, In: Lecam, L.M. and Neyman, J. (Eds.): *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability*, 281-297, 1967.
27. J. Moy, K. Lam, and E. Choo. Deriving Partial Values in MCDM by Goal Programming. *Annals of Operations Research,* 74:277-288, 1997.
28. F. Murtagh. A Survey of Recent Advances in Hierarchical Clustering Algorithms which Use Cluster Centers. *Computer Journal,* 26:354–359, 1983.
29. G. L. Nemhauser and L. A. Wolsey. *Integer and Combinatorial Optimization.* Wiley, New York, 1988.
30. K.-M. Osei-Bryson and A. Joseph. Applications of Sequential Set Partitioning: Three Technical Information Systems Problems. *Omega,* 34:492-500, 2006.
31. K.-M. Osei-Bryson and T. Inniss. A Hybrid Clustering Algorithm. *Computers & Operations Research*, in press, 2006.
32. L. Stanfel. Recursive Lagrangian Method for Clustering Problems. *European Journal of Operational Research,* 27:332-342, 1986.
33. J. Ward. Hierarchical Grouping to Optimize an Objective Function. *Journal of the American Statistical Association,* 58:236–244, 1963.

The Noisy Euclidean Traveling Salesman Problem: A Computational Analysis

Feiyue Li[1], Bruce L. Golden[2], and Edward A. Wasil[3]

[1] Department of Mathematics
University of Maryland
College Park, MD 20742
fli@ umd.edu

[2] R.H. Smith School of Business
University of Maryland
College Park, MD 20742
bgolden@rhsmith.umd.edu

[3] Kogod School of Business
American University
Washington, DC 20016
ewasil@american.edu

Summary. Consider a truck that visits n households each day. The specific households (and their locations) vary slightly from one day to the next. In the noisy traveling salesman problem, we develop a rough (skeleton) route which can then be adapted and modified to accommodate the actual node locations that need to be visited from day to day. In this paper, we conduct extensive computational experiments on problems with n = 100, 200, and 300 nodes in order to compare several heuristics for solving the noisy traveling salesman problem including a new method based on quad trees. We find that the quad tree approach generates high-quality results quickly.

Key words: Traveling salesman problem; computational analysis; average trajectory.

1 Introduction

The Euclidean Traveling Salesman Problem (TSP) is a well-known combinatorial optimization problem that is easy to state – given a complete graph $G = \{N, E\}$, where N is the set of nodes, E is the set of edges, and the distances are Euclidean, find the shortest tour that visits every node in N exactly once – and difficult to solve optimally. Algorithmic developments and computational results are covered by Jünger et al. [6], Johnson and McGeoch [5], Coy et al. [2], Pepper et al. [9], and others.

Recently, Braun and Buhmann [1] introduced the following variant of the TSP which they refer to as the Noisy Traveling Salesman Problem (NTSP). "Consider a salesman who makes weekly trips. At the beginning of each week, the salesman has a new set of appointments for the week, for which he has to plan the shortest

round-trip. The location of the appointments will not be completely random, because there are certain areas which have a higher probability of containing an appointment, for example cities or business districts within cities. Instead of solving the planning problem each week from scratch, a clever salesman will try to exploit the underlying density and have a rough trip pre-planned, which he will only adapt from week to week." In this paper, we consider a salesman who makes daily trips.

Braun and Buhmann viewed each node in a TSP as being sampled from a probability distribution, so that many TSP instances could be drawn from the same distribution. They used the sampled instances to build an *average trajectory* that was not forced to visit every node. For Braun and Buhmann, the average trajectory was "supposed to capture the essential structure of the underlying probability density." The average trajectory would then be used as the "seed" to generate an actual tour for each new day of appointments. Braun and Buhmann applied their average trajectory approach to a problem with 100 nodes.

To make the problem more concrete, consider the following. Each day, companies such as Federal Express and United Parcel Service send thousands of trucks to make local deliveries to households all across the United States. Let's focus on one of these trucks. Each day, it visits approximately the same number of households, in the same geographic region. The specific households may change from one day to the next, but the basic outline of the route remains the same. For example, if the truck visits the household located at 10 Main Street today, it might visit 15 Main Street instead (across the street) tomorrow. In the noisy traveling salesman problem, we develop a rough (skeleton) route which can then be adapted and modified to accommodate the actual node locations that need to be visited, from day to day.

We point out that the NTSP is similar to, but different from, the Probabilistic Traveling Salesman Problem (PTSP). In the PTSP, only a subset k ($0 \leq k \leq n$) out of n demand points needs to be visited on a daily basis. The demand point locations are known with certainty. See Jaillet [4] for details. Connections between the PTSP and the NTSP are discussed by Li [8].

In this paper, we conduct extensive computational experiments using three different data sets with different underlying structures to test Braun and Buhmann's approach, a simple convex hull, cheapest insertion heuristic, and a new heuristic (called the *quad tree approach*) that we develop for generating an average trajectory.

In Section 2, we describe Braun and Buhmann's approach. We show how they generate an average trajectory and then use it to produce an actual tour. We present their limited computational results. In Section 3, we conduct our extensive computational experiments. In Section 4, we develop the quad tree approach and test its performance. In Section 5, we present our conclusions and directions for future research.

2 Average Trajectory Approach

First, we provide the background that is needed to develop an average trajectory. We then generate an average trajectory for a small problem with seven nodes. Second, we give the details of the average trajectory approach and present Braun and Buhmann's computational results.

2.1 Generating an Average Trajectory

We demonstrate how to generate an average trajectory for a problem with seven nodes. The coordinates of the seven nodes are given in Table 1. Consider the following three trajectories that pass through all seven nodes:

$\phi_1 = [1, 2, 3, 4, 5, 6, 7]$, $\phi_2 = [1, 2, 4, 5, 3, 7, 6]$, $\phi_3 = [2, 3, 6, 5, 4, 7, 1]$.

Each of these (ϕ_1, ϕ_2, and ϕ_3) is a tour. As we will see, not all trajectories are tours. The distance between two specific trajectories is illustrated below.

$$\text{distance}(\phi_1, \phi_2) = ||\phi_1 - \phi_2|| =$$
$$\max\{|x_1 - x_1|, |y_1 - y_1|\} + \max\{|x_2 - x_2|, |y_2 - y_2|\}$$
$$+ \max\{|x_3 - x_4|, |y_3 - y_4|\} + \max\{|x_4 - x_5|, |y_4 - y_5|\}$$
$$+ \max\{|x_5 - x_3|, |y_5 - y_3|\} + \max\{|x_6 - x_7|, |y_6 - y_7|\}$$
$$+ \max\{|x_7 - x_6|, |y_7 - y_6|\}.$$

Note that $||\phi_1 - \phi_2||$ is defined as the L_1 norm.

Of course, we observe that ϕ_2 is equivalent to a set of trajectories. Given symmetric distances, the tour $1 - 2 - 3$ may be represented by $3 \times 2 = 6$ equivalent trajectories: [1, 2, 3], [1, 3, 2], [2, 1, 3], [2, 3, 1], [3, 1, 2], [3, 2, 1]. These are equivalent in terms of total length. More generally, given a sequence of n nodes to visit, we can do the following: start at any of these nodes, visit each node in the sequence by moving from left to right (with wrap-around, if necessary), and then we can repeat, moving from right to left. This yields $2n$ equivalent trajectories.

Let TR be the set of all $2n$ trajectories equivalent to ϕ_2. Then, using ϕ_1 as a basis of comparison, find the specific trajectory that solves $\min_{\phi \in TR} ||\phi - \phi_1||$. Let ϕ_2^* be that trajectory. $||\phi_2^* - \phi_1||$ represents the shortest distance between ϕ_1 and ϕ_2. These calculations are presented in Table 2. In particular, the shortest distances between ϕ_1 and ϕ_2 and between ϕ_1 and ϕ_3 are computed and marked accordingly. The average trajectory of ϕ_1 and ϕ_2 becomes $(\phi_1 + \phi_2^*)/2$. Similarly, the average trajectory of ϕ_1, ϕ_2, and ϕ_3 becomes $(\phi_1 + \phi_2^* + \phi_3^*)/3$.

In Table 2, we show how to compute $||\phi_1 - \phi_2||$. For example, for $\phi_1 = [1, 2, 3, 4, 5, 6, 7]$ and $\phi_2 = [1, 2, 4, 5, 3, 7, 6]$, we have

$$||\phi_1 - \phi_2|| = 0 + 0 + 0.5 + 1.0 + 1.5 + 1.25 + 1.25 = 5.50.$$

In particular, consider the calculation of the final component in $||\phi_1 - \phi_2||$. The final stop in ϕ_1 is node 7 with a coordinate of (4.875, 2.625). The final stop in ϕ_2 is node 6 with a coordinate of (5.125, 3.875). The difference in x coordinates is 0.25 and the difference in y coordinates is 1.25, so that the maximum difference is 1.25.

Table 1. Coordinates of a seven-node TSP.

	Coordinates	
Node	x	y
1	3.375	2.375
2	2.875	3.375
3	3.375	4.625
4	3.875	4.875
5	4.875	4.625
6	5.125	3.875
7	4.875	2.625

Table 2. Computing shortest distances between ϕ_1 and ϕ_2 and between ϕ_1 and ϕ_3.

ϕ_2	$\|\| \phi_1 - \phi_2 \|\|$	ϕ_3	$\|\| \phi_1 - \phi_3 \|\|$
[1,2,4,5,3,7,6] →	5.50	[2,3,6,5,4,7,1]	8.75
[2,4,5,3,7,6,1]	8.00	[3,6,5,4,7,1,2]	11.75
[4,5,3,7,6,1,2]	11.25	[6,5,4,7,1,2,3]	13.00
[5,3,7,6,1,2,4]	13.50	[5,4,7,1,2,3,6]	13.25
[3,7,6,1,2,4,5]	13.75	[4,7,1,2,3,6,5]	11.75
[7,6,1,2,4,5,3]	11.25	[7,1,2,3,6,5,4]	8.00
[6,1,2,4,5,3,7]	5.75	[1,2,3,6,5,4,7] →	2.50
[6,7,3,5,4,2,1]	9.50	[7,4,5,6,3,2,1]	11.00
[7,3,5,4,2,1,6]	9.25	[4,5,6,3,2,1,7]	10.50
[3,5,4,2,1,6,7]	8.50	[5,6,3,2,1,7,4]	11.75
[5,4,2,1,6,7,3]	11.50	[6,3,2,1,7,4,5]	12.00
[4,2,1,6,7,3,5]	11.75	[3,2,1,7,4,5,6]	9.75
[2,1,6,7,3,5,4]	10.50	[2,1,7,4,5,6,3]	6.00
[1,6,7,3,5,4,2]	8.00	[1,7,4,5,6,3,2]	8.00

For all 14 trajectories ϕ that are from the same tour as ϕ_2, we select the closest trajectory, that is, we select [1, 2, 4, 5, 3, 7, 6] since it has the minimum norm (5.50). We average the coordinates of ϕ_1 and ϕ_2 to obtain the average trajectory. Similarly, for all 14 trajectories ϕ that are from the same tour as ϕ_3, we select the closest trajectory, that is, we select [1, 2, 3, 6, 5, 4, 7] since it has the minimum norm (2.50), and average the coordinates. The coordinates of the average trajectories are given in Table 3. In Figure 1, we show ϕ_1, ϕ_2, and ϕ_3 and the average trajectories $(\phi_1 + \phi_2)/2$ and $(\phi_1 + \phi_2 + \phi_3)/3$. We see that the two average trajectories do not pass through all seven nodes.

Table 3. Coordinates of two average trajectories for a seven-node TSP.

	Coordinates	
Node	x	y
1	3.375	2.375
2	2.875	3.375
3	3.625	4.750
4	4.375	4.750
5	4.125	4.625
6	5.000	3.250
7	5.000	3.250

(a) Average of ϕ_1 and ϕ_2.

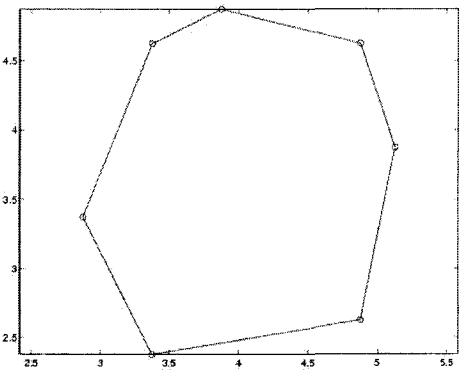
	Coordinates	
Node	x	y
1	3.375	2.375
2	2.875	3.375
3	3.542	4.708
4	4.625	4.458
5	4.375	4.625
6	4.625	3.792
7	4.958	3.042

(b) Average of ϕ_1, ϕ_2, and ϕ_3.

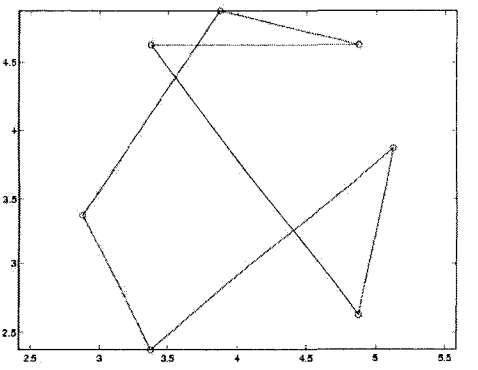
2.2 Braun and Buhmann's Approach

In practice, we face two questions in finding an average trajectory: (1) How do we generate sample trajectories that will produce a good average trajectory? (2) How do we use the average trajectory to generate a tour for a new problem instance? In this section, we describe the approach of Braun and Buhmann [1] to sampling and then generating a tour.

Braun and Buhmann start with one TSP instance and construct a Markov chain whose state space contains all permutations of the nodes. They sample from a Markov chain using random two-opt as the transition between two states and update the Markov chain using the Metropolis algorithm (see Kirkpatrick et al. [7]). Specifically, let x be the current state of the Markov chain, x' be the next state, x'' be the random two-opt of x, $\delta L = L(x'') - L(x)$ where $L(x)$ is the length of x, and U is a uniform (0, 1) random variable. If $\delta L < 0$, then $x' = x''$ with probability one, that is, always accept a downhill move. If $\delta L > 0$, then $x' = x''$ if $\exp(-\delta L / T) > U$ where T is the temperature; otherwise $x' = x$. This is referred to as a Markov Chain Monte Carlo simulation (denoted by MCMC). MCMC is, essentially, an iteration of simulated annealing with the temperature fixed to a single value.



(a) ϕ_1

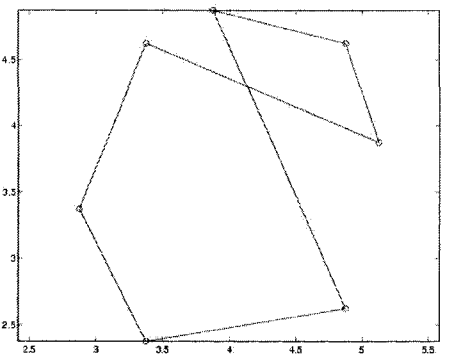


(b) ϕ_2

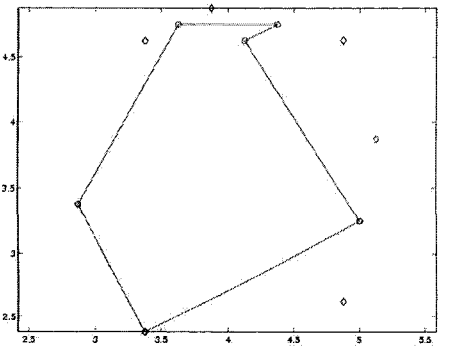
Fig. 1. Three trajectories (ϕ_1, ϕ_2, ϕ_3) that pass through all seven nodes of the TSP and two average trajectories (($\phi_1 + \phi_2$)/2, ($\phi_1 + \phi_2 + \phi_3$)/3) that pass through two nodes of the TSP and do not pass through five nodes of the TSP (denoted by the unconnected diamonds).

Braun and Buhmann draw one thousand samples from the Markov chain in order to generate the average trajectory. Between two consecutive samples, 100 transitions are performed to decouple the samples, that is, to reduce the dependency between the samples.

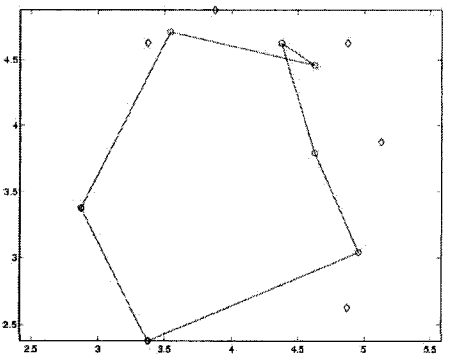
We summarize Braun and Buhmann's approach below. In Figure 2, we show an average trajectory for a 100-node problem that was generated by their approach.



(c) ϕ_3



(d) $(\phi_1 + \phi_2)/2$



(e) $(\phi_1 + \phi_2 + \phi_3)/3$

Fig. 1. (Continued)

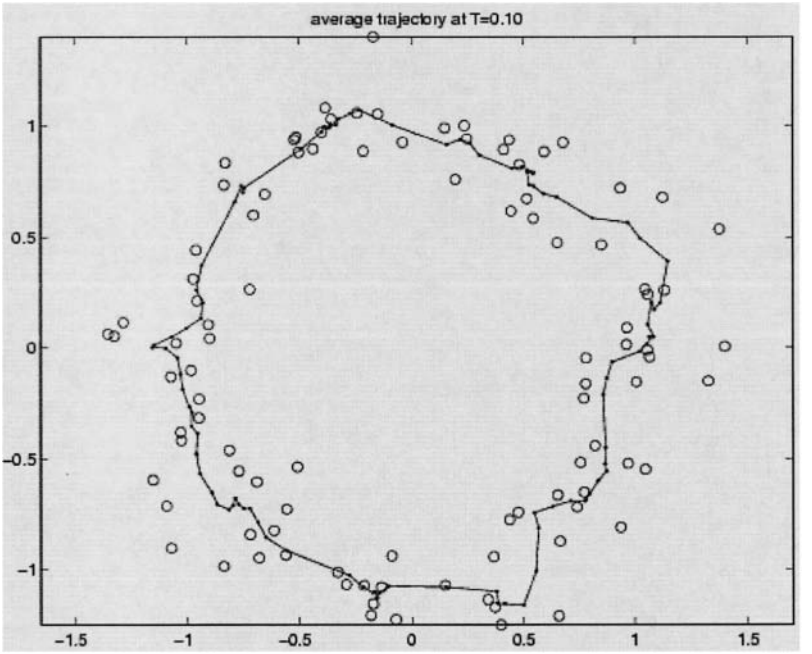


Fig. 2. Average trajectory for a 100-node problem generated by the approach of Braun and Buhmann.

Step 1. Select one instance, find an initial tour, and then run MCMC to generate sample tours.

Step 2. Average all sample tours to produce the average trajectory.

Step 3. For each new instance, apply the finite-horizon adaptation technique followed by a local search post-processor to generate the final tour.

Step 3 requires a bit of explanation. In order to generate a tour for a new TSP instance, Braun and Buhmann used a *finite-horizon adaptation* technique. First, the domain of the mapping for the average trajectory is extended from $\{1, \ldots, n\}$ to the interval $[1, \ldots, n + 1)$ which is called the *passing time*. For each node v_i in the new TSP instance, a point on the average trajectory with minimum distance to v_i is identified, and the passing time t_i is computed by linear interpolation. The permutation that sorts t_i, $i = 1, \ldots, n$ gives the initial solution for the new instance. A post-processor performs local optimization to remove intersections. We illustrate the finite-horizon adaptation technique in Figure 3.

Nodes 1 to 7 (in black) are taken from the average trajectory and they are mapped to points 1 to 7 on the passing-time axis. Nodes a to g (in white) are taken from the new instance. For each node v_i in the new instance, we find the nearest

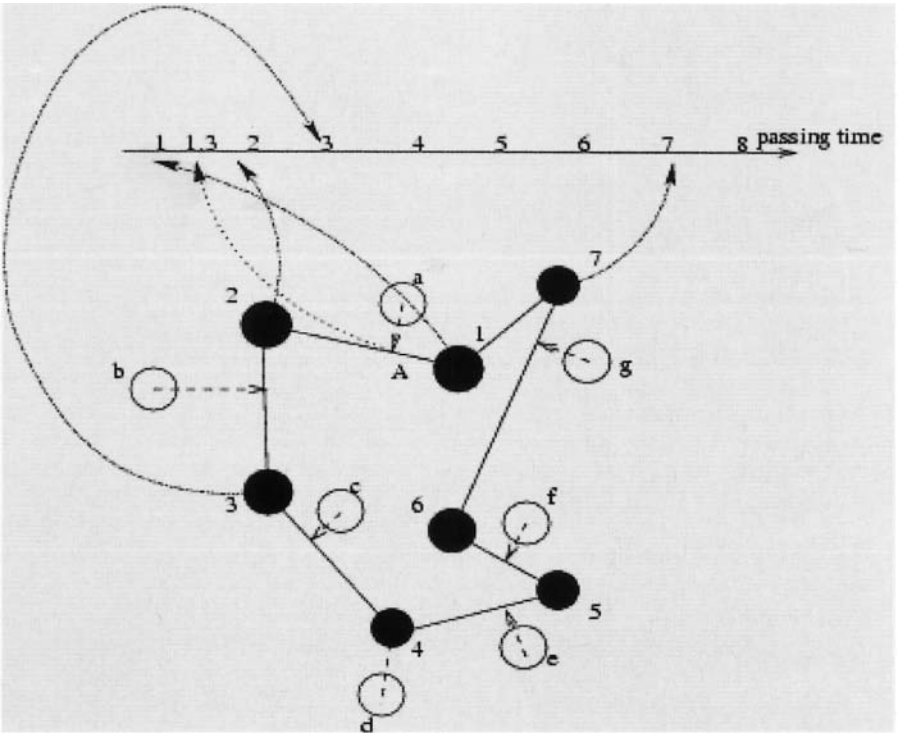


Fig. 3. Finite-horizon adaptation technique to generate a tour for a new instance.

point on the average trajectory (this is indicated by the straight, dashed lines with open arrow heads). Linear interpolation is then used to find the passing time t_i for a nearest node. Finally, we sort in increasing order of t_i to produce the ordering of the nodes in the new instance. For example, in Figure 3, the nearest point on the average trajectory to node a is A and its t value is 1.3.

We point out that there are several key limitations to Braun and Buhmann's approach. There are three parameters – temperature, number of samples, and interval between samples – whose values need to be set. Results are sensitive to the value of the temperature parameter. For example, average trajectories computed at low temperatures tend to overfit to the mean demand locations in the data.

It is not clear how many samples to select in MCMC. Braun and Buhmann used a sample size of 1,000 which requires lots of computation time. In our preliminary experiments, a sample size of 100 performed nearly as well and required much less computation time.

There is very limited computational experience with the three-step approach. Braun and Buhmann reported results for a single 100-node problem.

3 Computational Experiments

In this section, we conduct extensive computational experiments on Braun and Buhmann's approach. We provide the details of our experimental design, report computational results, and develop a new procedure for solving the NTSP.

3.1 Descriptions of Data Sets

In this section, we describe the data sets for our experiments. We use three data sets with different topologies and different sizes (100, 200, and 300 nodes).

Data Set One

Data set one is the same one used by Braun and Buhmann [1]. The mean demand locations are uniformly located on the unit circle, that is, $m_i = (\cos(2\pi i/n), \sin(2\pi i/n))$ for $i = 1,\ldots, n$ and the noise is normally distributed, that is, $N(0, \sigma^2)$. The coordinates of node $v_i = (x_i, y_i)$ are given by $x_i = \cos(2\pi i/n) + \zeta_i$ and $y_i = \sin(2\pi i/n) + \eta_i$, where ζ_i and η_i are $N(0, \sigma^2)$. In Figure 2, we show an instance from this data set with $n = 100$ and $\sigma^2 = 0.01$. In Figure 4, we show the same instance with average trajectories generated by four different temperatures ($T = 0.05$, 0.10, 0.15, 0.20). We see that temperature acts as a smoothing parameter. When the value of the temperature is low, then the average trajectory tends to overfit the data. When the value of the temperature is high, MCMC does not work well and the average trajectory tends to underfit the data. In Figure 4, the two middle temperatures seem to work the best.

Data Set Two

In the second data set, we have a hierarchical structure. At the first level, there are m clusters of means located around $(\cos(2\pi i/m), \sin(2\pi i/m))$ for $i = 1, \ldots, m$. For each cluster i, there are n means $(r \cos(2\pi j/n) + \cos(2\pi i/m), r \sin(2\pi j/n) + \sin(2\pi i/m))$ with independent Gaussian noise $N(0, \sigma^2)$. In Figure 5, we show an instance from this data set with $m = 6$, $n = 25$, $r = 0.25$, $\sigma^2 = 0.001$, and four different temperatures ($T = 0.02$, 0.05, 0.10, 0.20). The third temperature ($T = 0.10$) seems to work the best.

Data Set Three

In the third data set, we also have a hierarchical structure. There is a single mean. The x and y coordinates of the mean are randomly sampled from $N(0, 1)$. Each data point is then sampled around the mean with distribution $N(0, \sigma^2)$. In Figure 6, we show an instance from this data set $n = 200$, $\sigma^2 = 0.01$, and nine different temperatures. In this case, it is not clear that any of the temperatures work well.

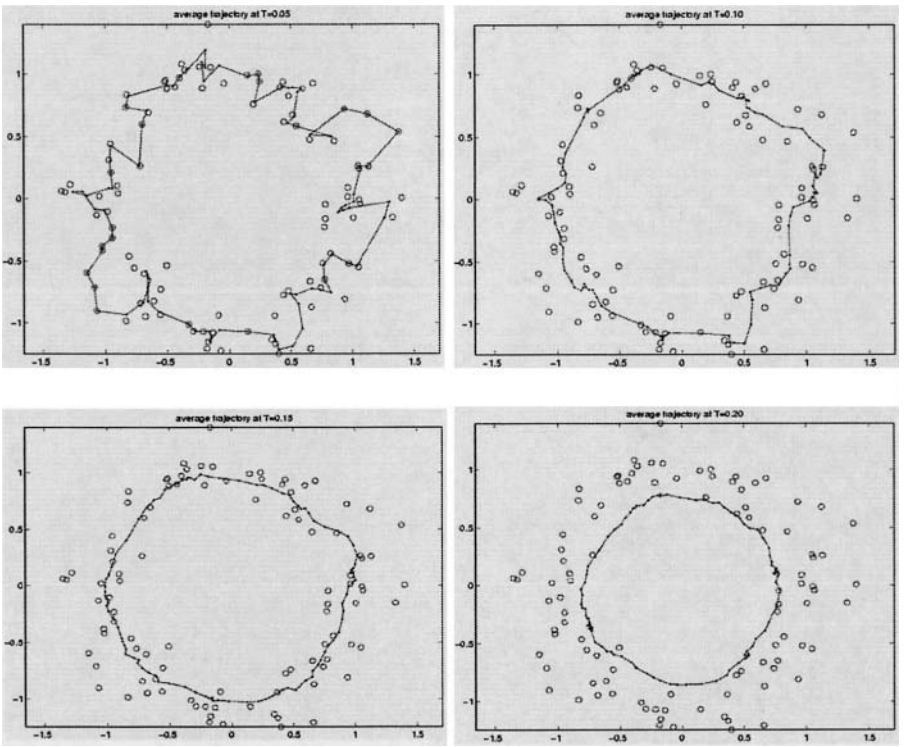
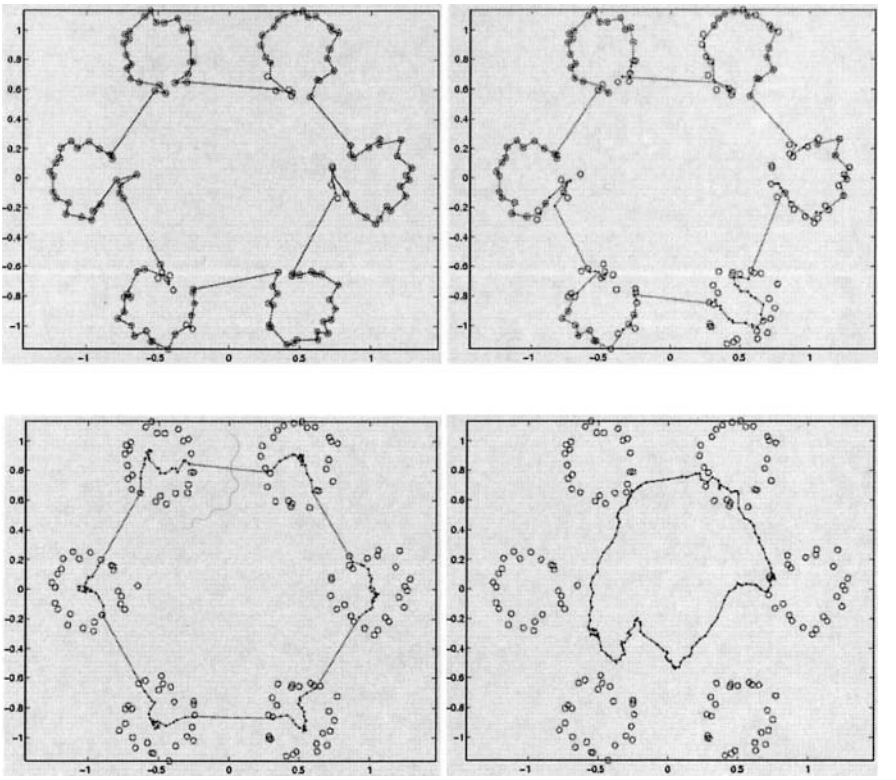


Fig. 4. Four average trajectories for a problem with $n = 100$ and $\sigma^2 = 0.01$ from data set one.

3.2 Experiments with Three Data Sets

For each of the three data sets, we conducted the following experiments. We used three sets of nodes ($n = 100, 200, 300$), three variances ($\sigma^2 = 0.01, 0.005, 0.0025$), and 30 temperatures ($T = 0.01$ to 0.30). At each temperature, we selected a sample of 100 instances in order to compute an average trajectory, generated the initial tour for a new instance using the finite-horizon adaptation technique, and applied the two-opt algorithm to the initial tour. For comparison purposes, we generated a tour for each new instance using a simple convex hull, cheapest insertion heuristic (denoted by CHCI). CHCI works in the following way: (1) Generate the convex hull of the instance – this becomes the initial subtour; (2) For each node that is not a vertex of the current subtour, insert it into the current subtour in the lowest-cost way (see Jünger et al. [6]).



(c) $T = 0.10$ (d) $T = 0.20$

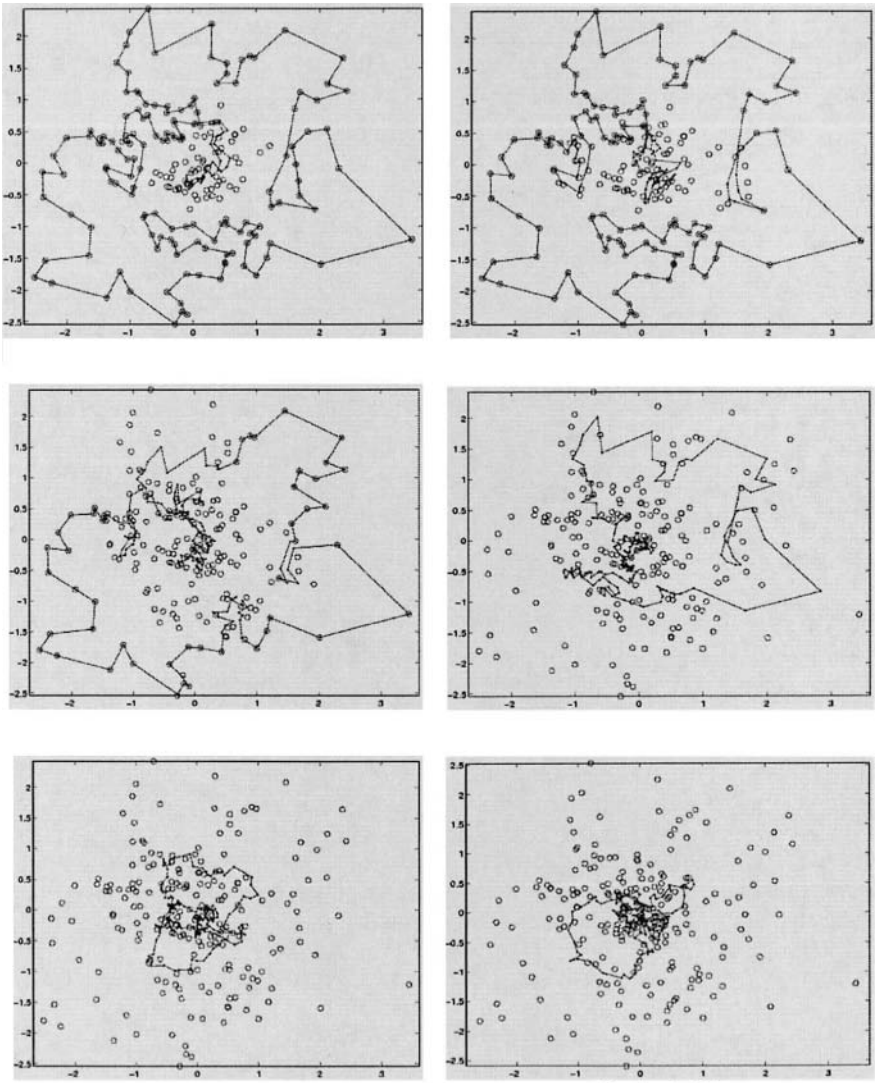
Fig. 5. Four average trajectories for a problem with $m = 6$, $n = 25$, $r = 0.25$, and $\sigma^2 = 0.001$ from data set two.

In Figure 7, we show the comparisons of the average trajectory approach versus CHCI for data set two for the three sets of nodes and the three variances. The x-axis gives the temperature and the y-axis gives the sample means for CHCI and the average trajectory approach (the figures for data sets one and three are given in [8].

In examining the figures for all three data sets, we make several observations.

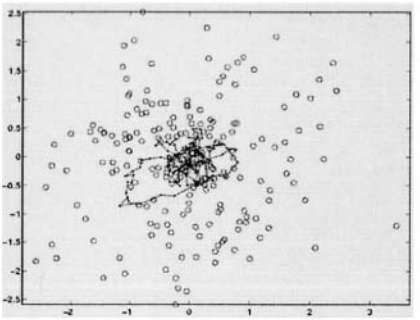
1. When the variance is *large*, the average trajectory approach does not perform as well as CHCI.
2. When the variance is *small*, the average trajectory approach performs about as well as CHCI in data sets one and three.
3. The topology of the data set plays a role in the performance of the average trajectory approach. CHCI is computationally faster on many instances.
4. There is a need to tune the temperature parameter (T) in the average trajectory approach in order to produce good results.

5. Problem size does not play much of a role in the performance of either method.
6. The tour produced by the average trajectory approach may not be visually appealing.

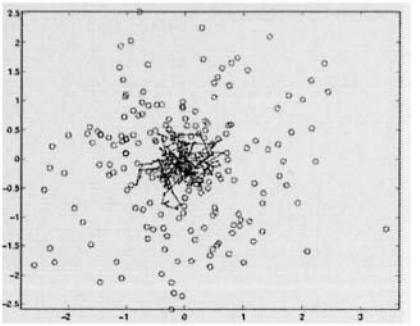


(e) $T = 0.30$ (f) $T = 0.40$

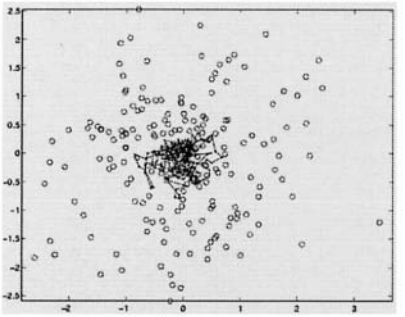
Fig. 6. Nine average trajectories for a problem with $n = 200$ and $\sigma^2 = 0.01$ from data set three.



(g) $T = 0.50$

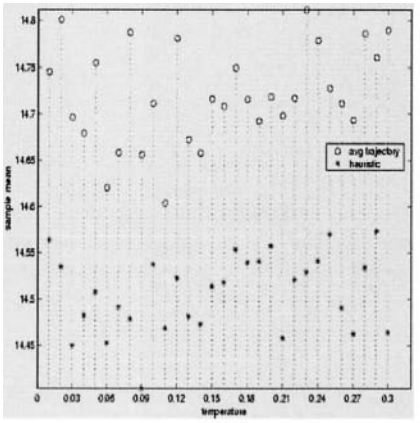


(h) $T = 0.60$

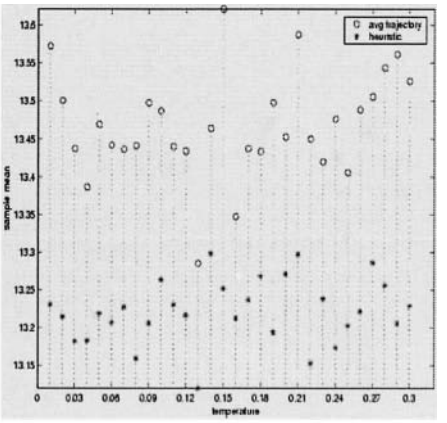


(i) $T = 0.70$

Fig. 6. (Continued)

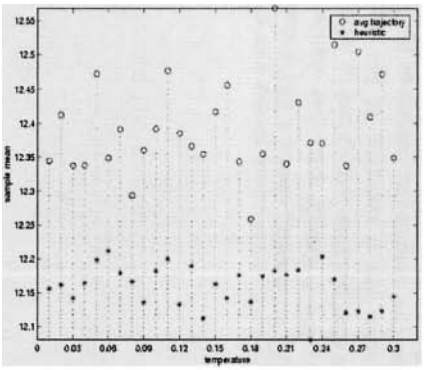


(a) $n = 100$, $\sigma^2 = 0.01$

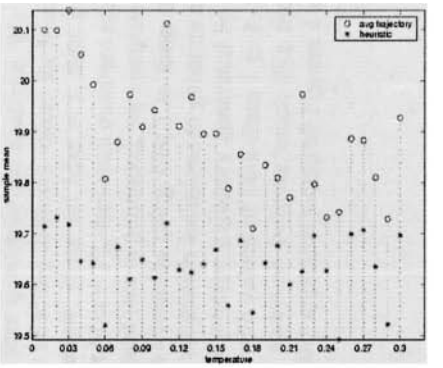


(b) $n = 100$, $\sigma^2 = 0.005$

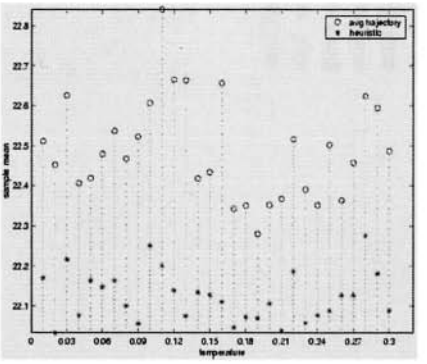
Fig. 7. Computational results for data set two for $n = 100$, 200, and 300, $\sigma^2 = 0.01$, 0.005, and 0.0025, and $T = 0.01$ to 0.30.



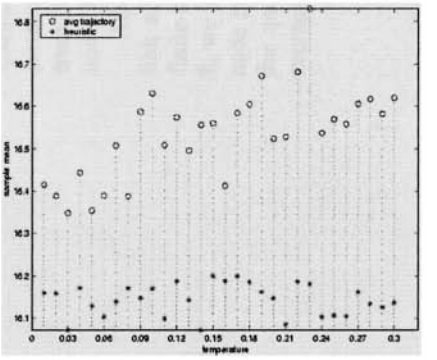
(c) $n = 100$, $\sigma^2 = 0.0025$



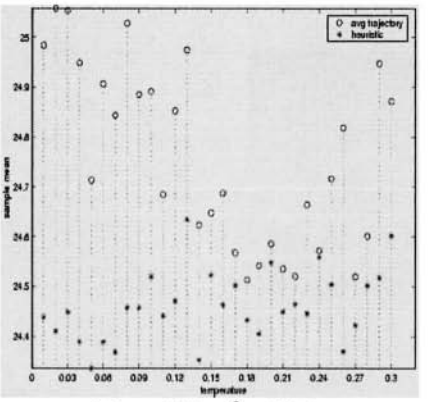
(d) $n = 200$, $\sigma^2 = 0.01$



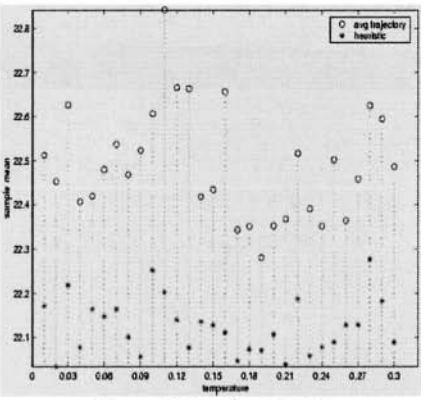
(e) $n = 200$, $\sigma^2 = 0.005$



(f) $n = 200$, $\sigma^2 = 0.0025$

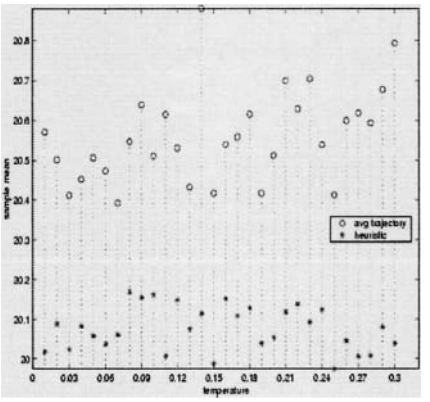


(g) $n = 300$, $\sigma^2 = 0.01$



(h) $n = 300$, $\sigma^2 = 0.005$

Fig. 7. (Continued)



(i) $n = 300$, $\sigma^2 = 0.0025$

Fig. 7. (Continued)

4 New Heuristic for the NTSP: Quad Tree Approach

In this section, we develop a new heuristic for the NTSP based on the quad tree and compare its performance to the average trajectory approach.

4.1 Quad Tree Approach

A quad tree is the extension of a binary tree in two-dimensional space (see Finkel and Bentley [3] for details). Given a TSP instance, we find a rectangle enclosing it and subdivide the rectangle into four regions called quadrants with the same geometry. We then generate a quad tree as follows. Select a quadrant and check the number of nodes that it contains. If the number of nodes does not exceed a specified value (we denote the parameter by Max), then we stop processing this quadrant. If the number of nodes is greater than Max, then we subdivide the quadrant into four quadrants and continue. We stop when all quadrants have been processed. By selecting different values for Max, we can generate a family of quad trees. For example, if we set Max = 1, then each quadrant contains at most one node. If we set Max = 10, then each quadrant contains at most 10 nodes.

After the quad tree is generated, we compute the centroid of each quadrant. We link all of the centroids using CHCI to form the average trajectory. We use the finite-horizon adaptation technique to generate a tour for a new instance. In Figure 8, we show the average trajectory generated by the quad tree approach for a 100-node problem with values of Max from 1 to 8. We see that the parameter Max in the quad tree approach behaves much like the temperature parameter T in the average trajectory approach (compare Figures 4 and 8).

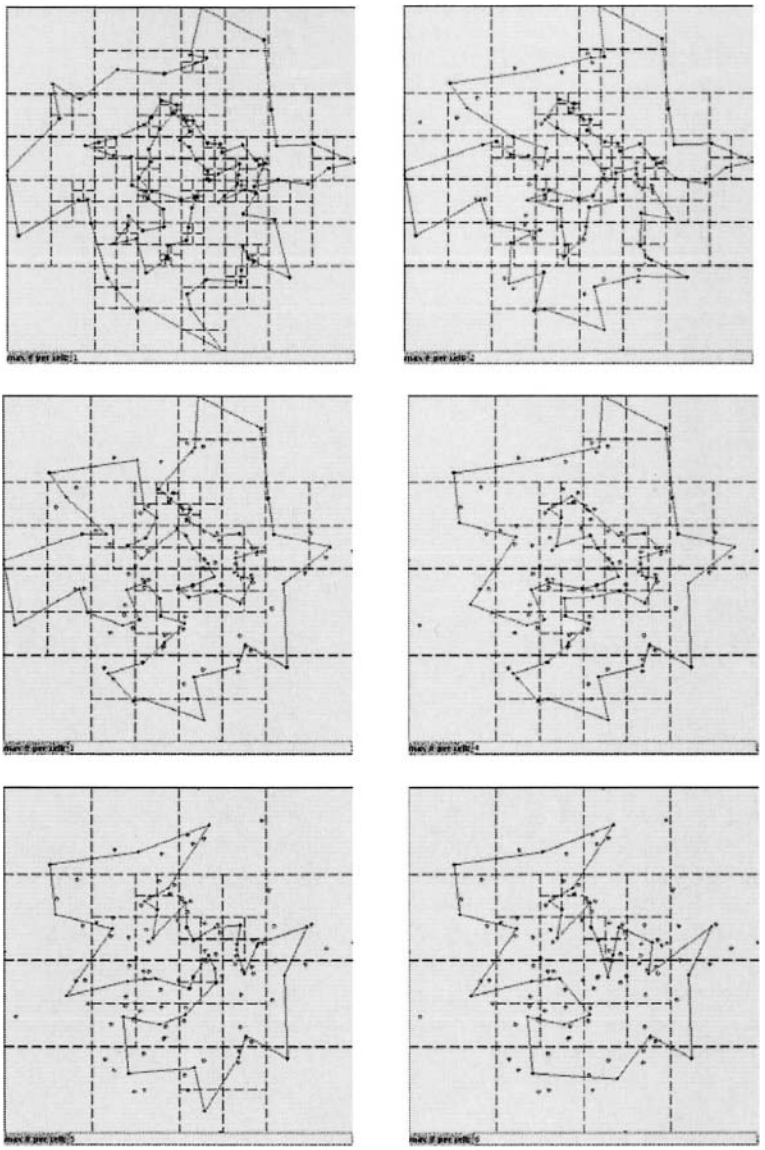


Fig. 8. Eight average trajectories generated by the quad tree approach for a problem with $n = 10$ and Max = 1, 2, ..., 8.

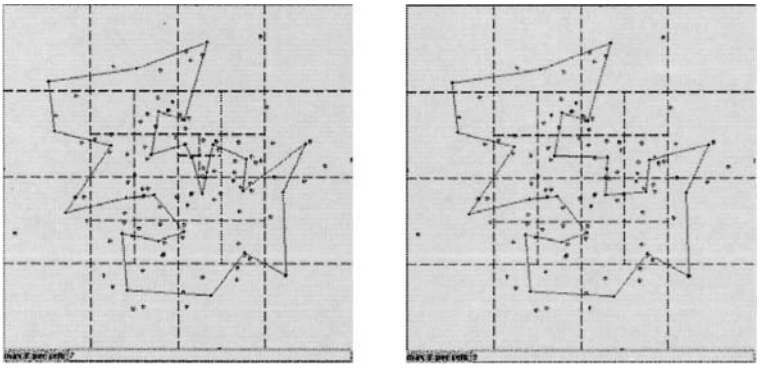


Fig. 8. (Continued)

4.2 Computational Experiments with the Quad Tree Approach

We selected a 100-node problem from data set three and conducted the following experiments with the average trajectory approach and the quad tree approach. We used three variances (σ^2 = 0.01, 0.005, 0.0025), 30 temperatures for the average trajectory approach (T = 0.01 to 0.30), and four values of Max for the quad tree approach (Max = 2, 4, 6, 8). We sampled 100 instances (at each temperature) in order to compute an average trajectory and generated a tour for each new instance by both methods.

In Figures 9, 10, and 11, we present our computational results. In each figure, the left-hand panel gives the average tour length over 100 instances, while the right-hand panel gives the percentage that the quad tree approach generates a lower-cost tour than the average trajectory approach over the 100 instances. In examining the figures, it is clear that, as Max increases in value (especially when Max = 6 and 8), the quad tree approach gives much better results. Apparently, limiting the number of nodes in a quadrant to 6 or 8 enables the average trajectory to capture the essential (general) shape of a good route. If the limit is smaller, the average trajectory mimics a specific (detailed) solution to a specific problem instance.

We point out that the average time to construct the average trajectory using MCMC in Braun and Buhmann's approach is about 60 ms (milliseconds) on an Athlon 1 GHz computer. The quad tree approach takes an average of about 160 ms when Max = 1 and about 20 ms when Max = 8 to construct the average trajectory. CHCI does not construct an average trajectory.

To generate the final tour for a new TSP instance, Braun and Buhmann's approach takes about 12 ms on average, while the quad tree approach takes about 10 ms and CHCI takes about 16 ms. Clearly, all three procedures are very quick.

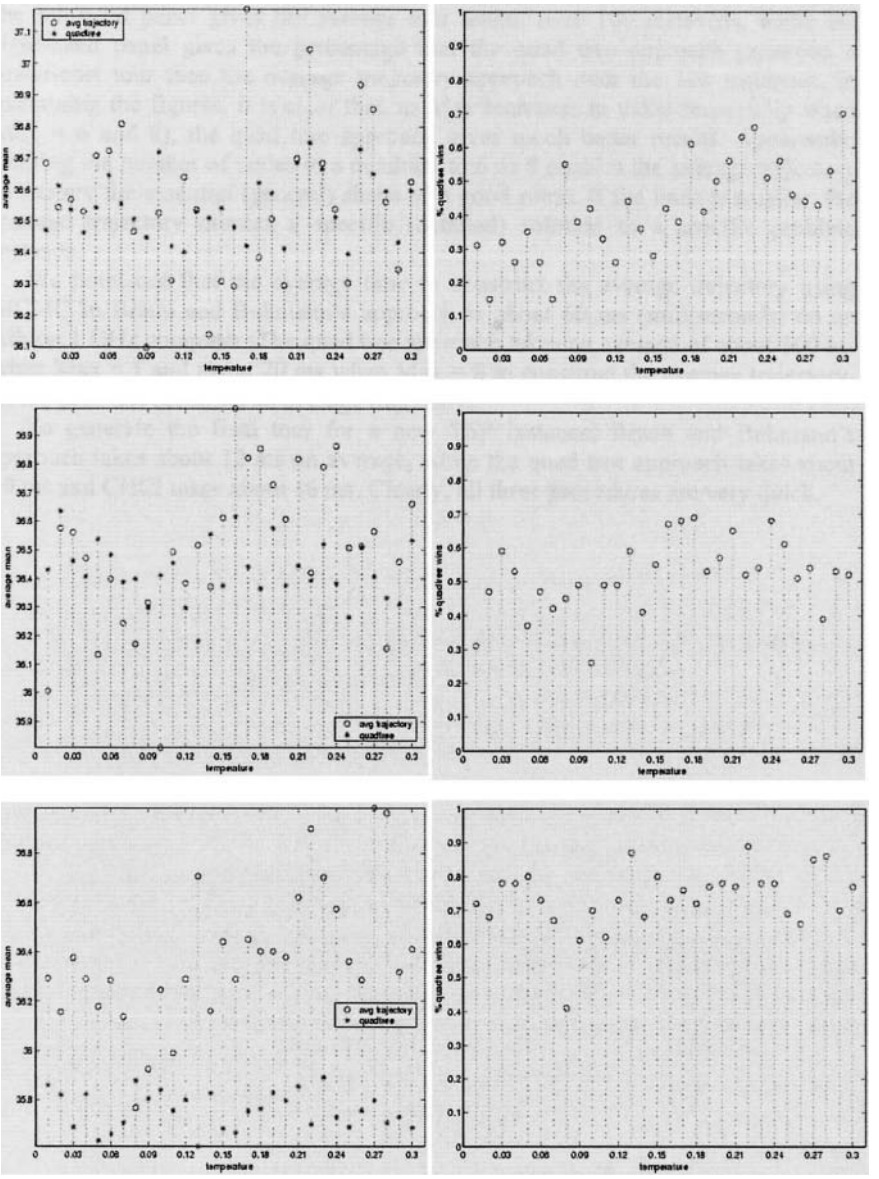


Fig. 9. Computational results for data set three for $n = 100$, $\sigma^2 = 0.01$, $T = 0.01$ to 0.30, and Max = 2, 4, 6, and 8 from top to bottom.

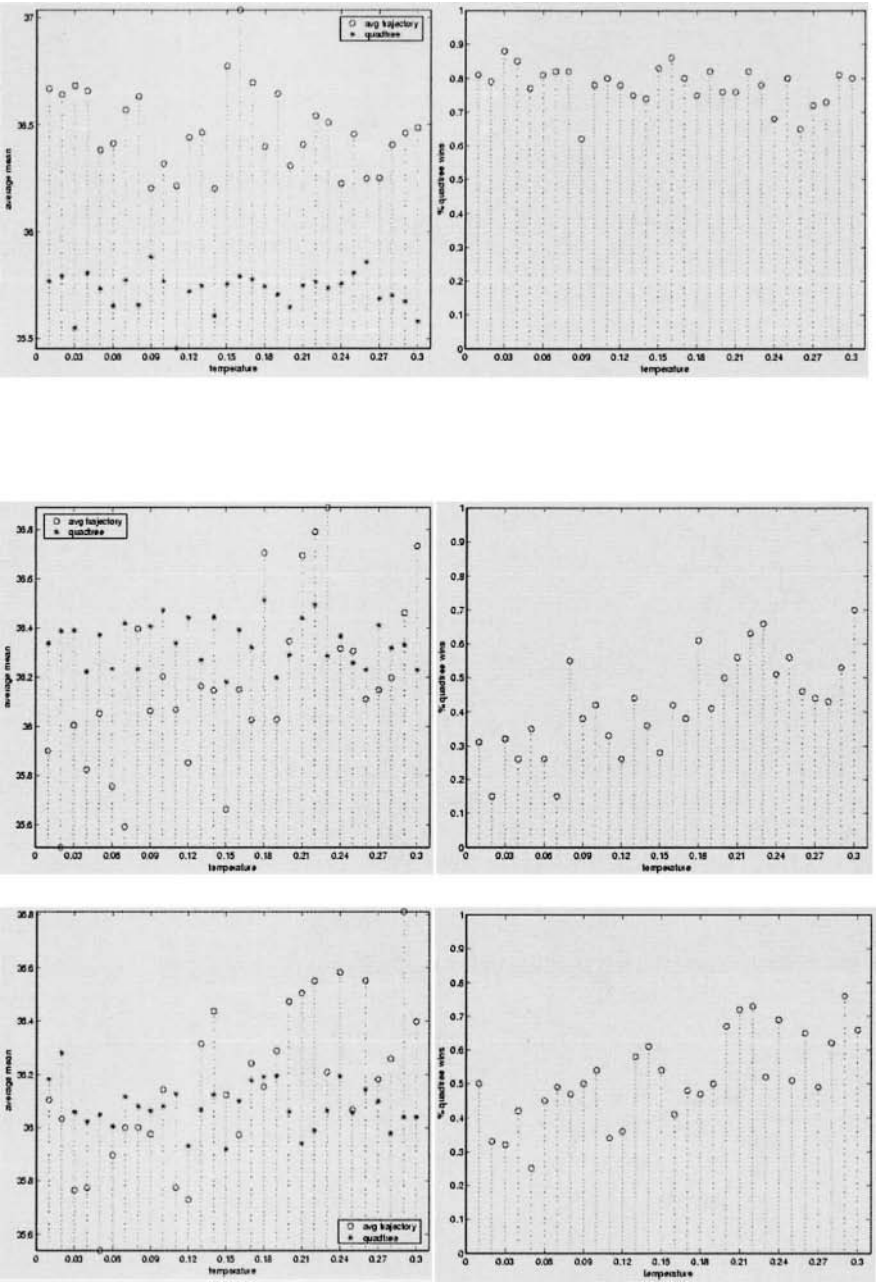


Fig. 10. Computational results for data set three for $n = 100$, $\sigma^2 = 0.005$, $T = 0.01$ to 0.30, and Max = 2, 4, 6, and 8 from top to bottom.

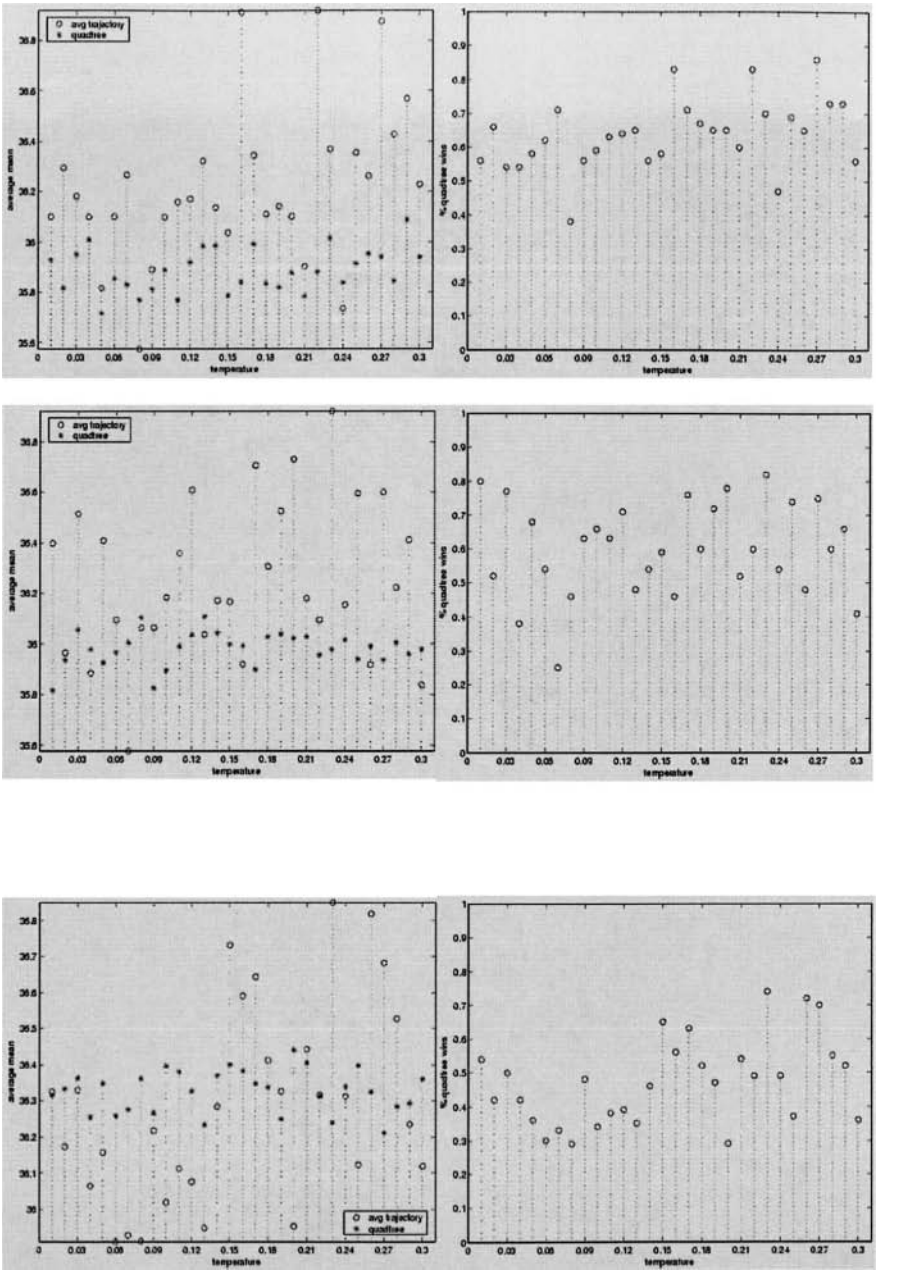


Fig. 11. Computational results for data set three for $n = 100$, $\sigma^2 = 0.0025$, $T = 0.01$ to 0.30, and Max = 2, 4, 6, and 8 from top to bottom.

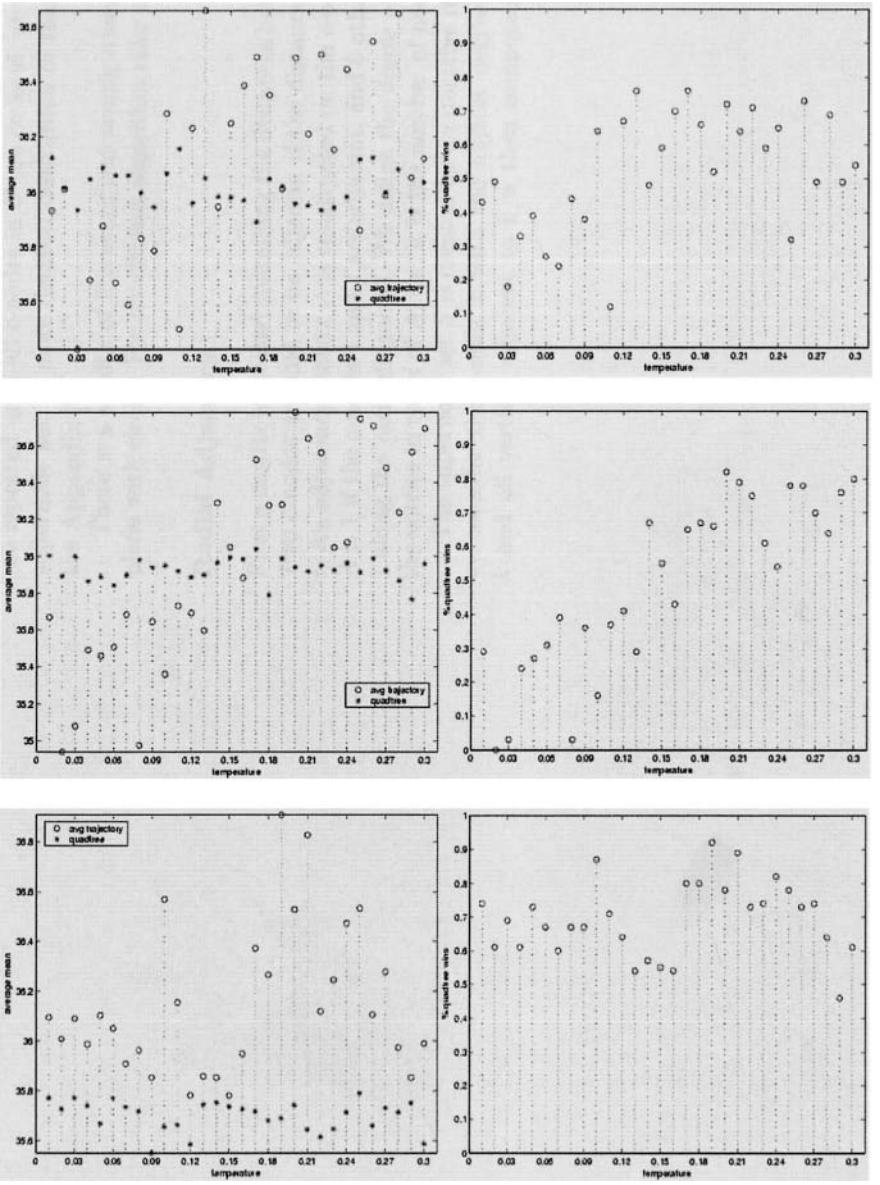


Fig. 11. (Continued)

5 Conclusions and Future Work

In this paper, we conducted extensive computational experiments using Braun and Buhmann's approach, a new quad tree approach that we developed, and a simple, convex hull, cheapest insertion heuristic to generate average trajectories and final tours for the noisy traveling salesman problem. We used problems that had 100, 200, and 300 nodes. All three procedures quickly generated average trajectories (160 ms to 600 ms on average) and final tours (10 ms to 16 ms on average).

In Braun and Buhmann's average trajectory approach, we needed to set the values of three parameters including the sample size and temperature. In the quad tree approach, we needed to set the value of one parameter (Max). When the value of Max was large, the quad tree approach generated final tours that were much lower in cost than the final tours generated by the average trajectory approach. When the problem variability was large, the average trajectory approach did not perform as well as the simple, convex hull, cheapest insertion heuristic. For the most part, when the problem variability was small, the quad tree approach performed better than the average trajectory approach. We recommend using the quad tree approach to generate average trajectories and high-quality final tours for the noisy traveling salesman problem with small variability and using the convex hull, cheapest insertion heuristic for a problem with large variability. In other words, as long as variability is not too large, an average trajectory can be computed in advance and adapted each day, rather than having to solve a new routing problem from scratch each day.

In future work, we will examine a variant of the NTSP in which some percentage of the nodes change regularly and the remaining nodes are fixed and always serviced. We call this the probabilistic noisy traveling salesman problem (PNTSP). The PNTSP is encountered frequently in practice. For example, drivers know that they will service certain high-demand customers every day and service low-demand customers every few days.

References

1. M. L. Braun and J. M. Buhmann. The noisy Euclidean traveling salesman problem and learning. in *Advances in Neural Information Processing Systems 14*, T. Dietterich, S. Becker, and Z. Ghahramani, eds., 251-258, MIT Press, Cambridge, MA, 2002.
2. S. Coy, B. Golden, G. Runger, and E. Wasil. See the forest before the trees: Fine-tuned learning and its application to the traveling salesman problem. *IEEE Transactions on Systems, Man, and Cybernetics*, 28A: 454-464, 1998.
3. R. A. Finkel and J. L. Bentley. Quad trees, a data structure for retrieval on composite keys. *Acta Informatica*, 4: 1-9, 1974.
4. P. Jaillet. A priori solution of a traveling salesman problem in which a random subset of the customers are visited. *Operations Research*, 36: 929-936, 1988.

5. D. Johnson and L. McGeoch. The traveling salesman problem: A case study in local optimization. in *Local Search in Combinatorial Optimization*, E. Aarts and J. K. Lenstra, eds., 215-310, Wiley, London, 1997.
6. M. Jünger, G. Reinelt, and G. Rinaldi. The traveling salesman problem. in *Network Models, Volume 7, Handbooks in Operations Research and Management Science*, M. Ball, T. Magnanti, C. Monma, and G. Nemhauser, eds., 225-330, North-Holland, Amsterdam, 1995.
7. S. Kirkpatrick, C. Gelatt, and M. Vecchi. Optimization by simulated annealing. *Science*, 220: 671-680, 1983.
8. F. Li. Modeling and solving variants of the vehicle routing problem: Algorithms, test problems, and computational results. Ph.D. dissertation, University of Maryland, College Park, Maryland, 2005.
9. J. Pepper, B. Golden, and E. Wasil. Solving the traveling salesman problem with annealing-based heuristics: A computational study. *IEEE Transactions on Systems, Man, and Cybernetics*, 32A: 72-77, 2002.

The Close Enough Traveling Salesman Problem: A Discussion of Several Heuristics

Damon J. Gulczynski, Jeffrey W. Heath, and Carter C. Price

Applied Mathematics Program
Department of Mathematics
University of Maryland
College Park, MD 20742
damon@math.umd.edu, jheath@math.umd.edu, price@math.umd.edu

Summary. The use of radio frequency identification (RFID) allows utility companies to read meters from a distance. Thus a meter reader need not visit every customer on his route, but only get within a certain radius of each customer. In finding an optimal route – one that minimizes the distance the meter reader travels while servicing each customer on his route – this notion of only needing to be close enough changes the meter reading problem from a standard Traveling Salesperson Problem (TSP) into a variant problem: Close Enough TSP (CETSP). As a project for a graduate course in network optimization various heuristics for finding near optimal CETSP solutions were developed by six groups of students. In this paper we survey the heuristics and provide results for a diverse set of sample cases.

Key words: Traveling salesman problem; radio frequency identification; electronic meter reading.

1 Introduction

Historically when a utility company measures the monthly usage of a customer, a meter reader visits each customer and physically reads the usage value at each site. Radio frequency identification (RFID) tags at customer locations can remotely provide data if the tag reader is within a certain radius of the tag. This changes the routing problem from one of a standard traveling salesman problem (TSP) to what we call a "Close Enough" TSP (CETSP). Thus the route lengths of the meter readers can be drastically reduced by developing heuristics that exploit this "close enough" feature.

We consider such a meter reading routing problem where each customer is modeled as a point in the plane. Additionally there is a point that represents the depot for the meter reader. A CETSP tour must begin and end at the depot and travel within the required radius, r, of each customer. For simplicity in the cases tested here the meter reader was not restricted to a road network.

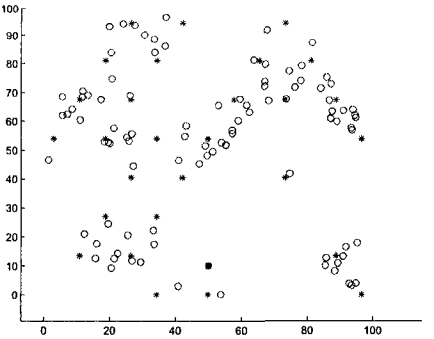


Fig. 1. An example of a supernode set on 100 nodes, with radius 9, and the depot located at (50,10). The circles represent the customer nodes, and the asterisks are the supernodes.

All distances are Euclidean and the objective is to minimize the total distance traveled. The solution to a standard TSP on the customer nodes obviously provides an upper bound for the CETSP.

Essentially the CETSP is a TSP with a spatial window. Thus it is conceptually similar to the TSP with a time window. Several heuristics are discussed in the work of Solomon [9]. In spite of the similarities the heuristics for the TSP problem with a time window do not directly apply to the CETSP. This is because we do not simply change the order of the points visited, we actually change the location of the points themselves. So the CETSP has an uncountable solution space.

As a class project six teams developed heuristics to produce solutions to this problem. This paper highlights the innovative developments from these projects. In Section two we discuss the proposed heuristics. In the third Section we present the numerical results across a test bed of cases. We conclude with some suggestions for further work in Section four. Psuedo code for some of the heuristics is provided in the appendix.

2 Heuristics

All of the heuristics developed have three distinct steps. Given an initial set C of customer nodes (c-nodes), the first step is to produce a *feasible supernode* set, S. A feasible supernode set is a set of points (containing the depot node) with the property that each customer node is within r units of at least one point in the set. In Figure 1 the set of asterisks $(*)$ represents a feasible supernode set since each customer (circle) is within r units ($r = 9$) of at least one asterisk. After producing S the second step is to find a near optimal TSP

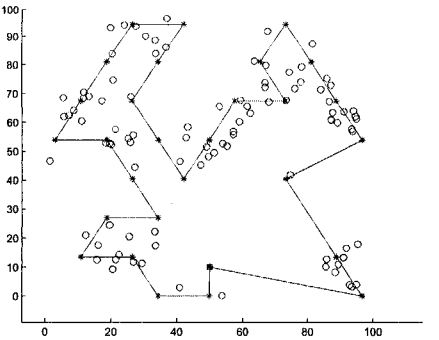


Fig. 2. An example CETSP tour on 100 nodes, with radius 9, and the depot located at (50,10). The circles represent the customer nodes, and the asterisks are the supernodes.

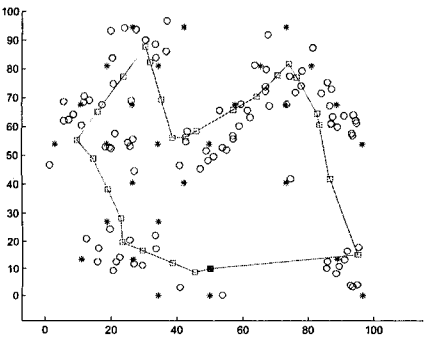


Fig. 3. The results of an economization of the tour in Figure 2.

tour, T, on the points in S, as seen in Figure 2. Since each customer node is within r units of a supernode, we are guaranteed that in traversing T we will pass within the required radius of each customer. We call T a feasible CETSP tour. Since the cardinality of the feasible supernode set is smaller than the number of customers, sometimes significantly, it is more efficient to generate the tour on S. Thus performing step 1 prior to step 2 requires much less computational time than starting by generating the TSP on C. The final step is an economization routine that reduces the distance traveled in T while maintaining feasibility, thus generating a shorter CETSP tour T'. The results of this economization can be seen in Figure 3.

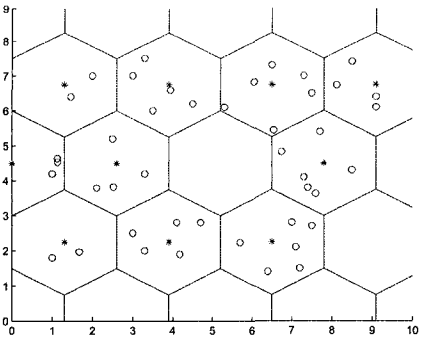


Fig. 4. An example of tiling the plane with regular hexagons of side length $r = 1.5$ units. We ensure that all customers (represented as small circles) in a given tile are within r units of the center of that tile (*).

2.1 Producing a Feasible Supernode Set

Four heuristics with variations were developed for producing a feasible supernode set: tiling with three variations, Steiner zone, sweeping circle, and radial adjacency. Each of these techniques is based on the assumption that it is desirable to have as few supernodes as possible. Examples can be constructed in which the fewest number of supernodes does not result in the shortest CETSP tour; however, empirical tests show that in general reducing the number supernodes is a good strategy.

Tiling Methods

For the tiling methods of producing a feasible supernode set, the plane is tiled with polygons that can be inscribed in a circle of radius r. This ensures that each c-node in the polygonal tile is within a distance r of its center. Thus by letting the supernodes be the centers of the tiles which contain at least one c-node we create a feasible supernode set. In implementation regular hexagons with side length r were chosen because over all regular polygons because they minimize the area of overlap between adjacent supernodes (See Figure 4). Once a feasible supernode set S has been constructed there are several techniques to reduce the cardinality of S, including shifting, merging, and circular extension.

One can translate or shift the tile centers in an attempt to reduce the total number of supernodes. The translation procedure works by performing a series of small shifts (vertically and/or horizontally) on the centers of the tiles. This process creates a collection of feasible supernode sets and from this collection we choose the set with the smallest cardinality (see Figure 5) [5].

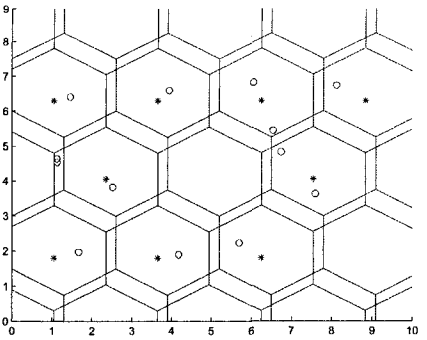


Fig. 5. The tiles from Figure 4 are translated as shown to reduce the total number of supernodes from ten to nine.

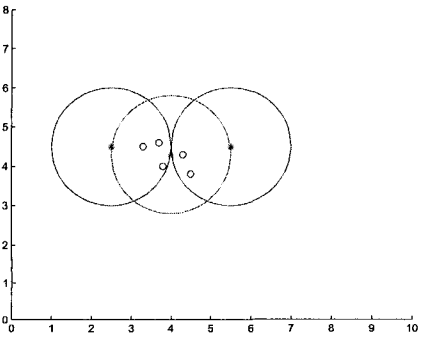


Fig. 6. Two supernodes are merged into one.

Merging works by considering the c-nodes in two adjacent tiles and determining if they can be covered with one circle. In this way it might be possible to merge two adjacent supernodes into one (see Figure 6). Specifically we take the maximal and minimal x and y values of the c-nodes in the two adjacent tiles: x_{min}, x_{max}, y_{min} and y_{max}. The midpoint of the extreme x's and y's, $(\frac{x_{min}+x_{max}}{2}, \frac{y_{min}+y_{max}}{2})$, is then considered as a potential supernode. If all the constituent c-nodes of the two tiles are within r units of this midpoint the merger is accepted, and the number of supernodes is reduced by one. This process continues until no mergers are possible [2].

Given the supernode set, S, obtained from the centers of hexagonal tiles it might be possible to eliminate some elements of S by considering the intersection of their corresponding circumscribed circles. Let us define the *degree* of a supernode as the number of c-nodes that lie within r units of this supernode,

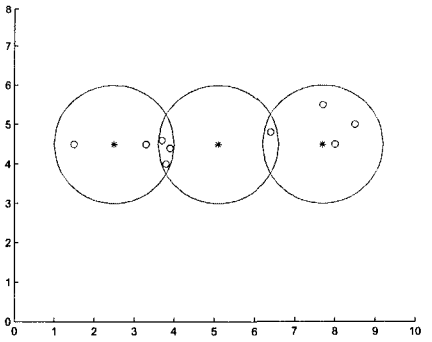


Fig. 7. By applying circular extension the middle supernode can be omitted.

and do not lie within r units of any other center. By associating c-nodes that lie in an overlap (i.e. an intersection of two circles) with the center of the largest degree it might be possible to eliminate superfluous centers reducing the size of the feasible supernode set (see Figure 7) [3].

Steiner Zone

From the nodes' perspective, the meter reader can visit any point within r units of the node. Thus there is a circle of radius r around the node, representing the node's service region, through which the meter reader must pass. If two c-nodes are less than $2r$ units apart, then their corresponding service region circles will overlap. Thus any node in this overlap region represents a candidate supernode covering these two nodes. We can minimize the cardinality of a supernode set S by choosing points that lie within the intersection of multiple circles. The Steiner zone method consists of finding these intersections first, and then choosing the supernode set in this manner.

Let $D(c_1, r), D(c_2, r), ..., D(c_k, r)$ be discs of radius r centered at c-nodes $c_1, c_2, ..., c_k$ respectively, then $D(c_1, r) \cap D(c_2, r) \cap ... \cap D(c_k, r)$ is the set of points within r units of each of $c_1, c_2, ..., c_k$. If this intersection is not empty we call it the *Steiner zone* of $c_1, c_2, ...c_k$ and denote it by $Z(c_1, c_2, ..., c_k)$. Furthermore we say that $Z(c_1, c_2, ..., c_k)$ covers $c_1, c_2, ..., c_k$, and k is called the *degree* of the Steiner zone (see Figure 8). Any point in the Steiner zone $Z(c_1, c_2, ..., c_k)$ can be a supernode covering the c-nodes $c_1, c_2, ..., c_k$.

Since out goal is to minimize the number of supernodes, it is advantageous to find Steiner zones of large degree. Ideally we would like to enumerate every Steiner zone and greedily choose those of largest degree. In practice this in unreasonable as the number of distinct Steiner zones on n c-nodes could be as large as 2^n. However, in order to quickly obtain a feasible supernode set yielding a good CETSP tour we need not consider all Steiner zones. Empirical

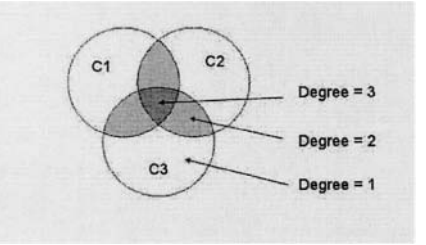


Fig. 8. The steiner zones of a three node group are displayed here.

tests have shown that it is sufficient to consider Steiner zones less than some fixed degree, and from those Steiner zones build zones of higher order. For our example results in Section 3 the fixed degree was four. Once all c-nodes are covered by at least one Steiner zone a feasible supernode set is obtained by choosing an arbitrary point in each respective zone. Psuedocode can be found in the Appendix [8].

Sweeping Circle

The sweeping circle heuristic covers the plane with overlapping circles with centers offset by $d_{min} = min(\sqrt{2}r, min(\{d_{ij}\}))$, where d_{ij} is the Euclidean distance between c-nodes i and j (see Figure 9). This heuristic is greedy because it chooses the center of the circle containing the most c-nodes and adds it to the set of supernodes S in each iteration (See Figure 10). All c-nodes within this circle are now covered. The sweeping and selection steps are repeated until all c-nodes in C are covered, at which point S is a feasible supernode set. A more detailed description of this procedure is provided in the Appendix [7].

There is a wealth of literature concerning covering a set of points on the plane with circles. For an efficient algorithm refer to Gonzalez [4].

Radial Adjacency

First a matrix is created containing the distances between each pair of c-nodes. Two c-nodes are said to be adjacent if the distance between them is at most r. An adjacency matrix A is constructed on the c-nodes, where entry (i, j) in A is 1 if the c-nodes i and j are adjacent, and 0 otherwise (by convention A is 1 along the main diagonal). We define the *degree* of a c-node i as the sum of the entries in row i of A, i.e., it is the number of nodes to which i is adjacent.

The supernode set S is created by an iterative method where at each step we consider the c-node k with the highest degree. The geometric mean of k and all vertices adjacent to k is then computed. If this geometric mean

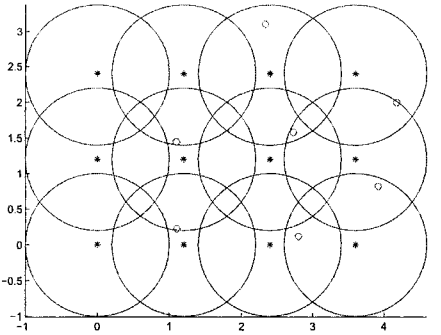


Fig. 9. Graphical depiction of initial sweeping circles.

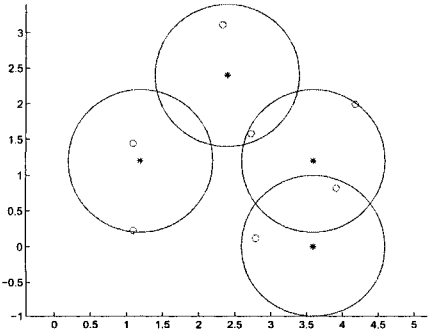


Fig. 10. Those circles containing the most c-nodes are selected until all c-nodes are covered.

is adjacent to more c-nodes than k, then it is designated as a supernode at this step. Otherwise, k is designated as a supernode (see Figure 11). A more detailed description of this procedure is provided in the Appendix.

In place of the greedy selection step, the following variation of this iteration was used for sparse data sets in the hope of minimizing the large travel distances needed to reach remote c-nodes. First select the c-node with the smallest degree. From all the c-nodes adjacent to the selected node, choose the one with the highest degree, k. Add k to S and remove k and all c-nodes adjacent to k from consideration. The heuristic ends when all nodes are considered [1].

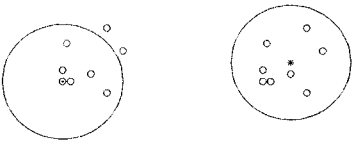


Fig. 11. In this case the geometric mean (* in right figure) of the points adjacent to k has a higher degree than k (the solid circle in left figure).

2.2 TSP solver

Once a feasible supernode set is produced the second step is to find a TSP tour, T, on the supernodes. There is a wealth of literature on this subject and we will not go into the topic here [6]. In practice the groups used a variety of TSP heuristics resulting in near-optimal tours. Thus we expect only minimal variation in the tour lengths due to the software selection.

2.3 Economization

The third and final step of the heuristics is the economization algorithm. This algorithm is based on minimizing the marginal cost of visiting a given node in the tour T. We first enumerate the supernodes on the tour in the order they are visited. If node i were not in T, the shortest path from node $i-1$ to node $i+1$ would be a straight line between the two. So we determine the line between $i-1$ and $i+1$ and move node i as close to this line as possible while still maintaining feasibility.

We repeat this routine for each node in T except for the depot node which must remain fixed. The resulting tour will have a total distance traveled no greater than the length of T. By repeating this technique on each new tour, a sequence of tours is created: $T_1, T_2, T_3, \ldots$ each one of length no greater than the one before it. The process terminates when the improvement between two tours drops below a specified tolerance. The resulting tour is our CETSP solution T' [5].

Table 1. Shortest tour lengths $\ell(T)$ generated by heuristics discussed.

Problem	Data Type	c-nodes	TSP length	Radius	Method	$\ell(T)$
1	clustered	100	655.09	9	Shifted Tiling	344.89
2	random	200	1154.06	20	Merging Tiling	288.16
3	clustered	300	1120.49	7	Merging Tiling	537.17
4	random	400	1533.95	5	Merging Tiling	797.04
5	clustered	500	1189.40	2	Merging Tiling	798.60
6	random	500	1627.91	27	Shifted Tiling	246.08
7	clustered	1000	2231.40	12	Steiner Zone	461.82

Table 2. Fewest supernodes generated by heuristics discussed.

Problem	Data Type	c-nodes	Radius	Method	supernodes
1	clustered	100	9	Steiner Zone	18
2	random	200	20	Steiner Zone	11
3	clustered	300	7	Steiner Zone	38
4	random	400	5	Steiner Zone	18
5	clustered	500	2	Steiner Zone	147
6	random	500	27	Steiner Zone	8
7	clustered	1000	12	Merging Tiling	30

Table 3. Tour lengths from each of the heuristics discussed.

Problem	Steiner Zone	Sweep	Radial Adjacency	Circular Extension	Merging	Shifting
1	375.56	378.65	512.19	410.56	377.87	344.89
2	288.39	300.83	448.33	342.34	288.16	327.31
3	560.26	562.64	758.95	651.01	537.17	597.94
4	838.72	849.42	1154.70	1141.80	797.04	870.68
5	819.90	1014.98	1040.50	1662.50	798.60	827.76
6	278.20	246.79	373.18	304.38	279.06	246.08
7	461.82	468.88	628.16	504.44	468.54	484.42

3 Numerical Results

These heuristics were tested on a diverse set of test cases. The distances for each method were compared and the best tour lengths are reported below along with the method that produced the best result. Also reported are the methods that resulted in the fewest supernodes for each problem. Some the data sets have randomly generated c-nodes while others have clusters of c-nodes. The TSP length represents the total distance of a near optimal solution of the TSP on the set of c-nodes. The data sets have a variety of values for the radius to provide a representative sample for real world applications. All of this was done in a 100 by 100 grid.

4 Conclusion

These heuristics produce CETSP tours in the test bed of cases that have significantly shorter length than TSP tours on the c-nodes. The hexagonal tiling heuristics were the most successful, particularly if an extension such as the shifting or merging heuristics is used. The Steiner zone method also proved to be quite effective. While the methods that result in fewer supernodes generally have the shorter final tour lengths, the method with the fewest supernodes does not always produce the shortest tour. The Steiner zone method produced the fewest supernodes in most cases but the tiling heuristics generally produced the shortest tour.

Clearly the use of RFID technology can reduce the travel distance of meter readers substantially, though further work is required to port this method from Euclidean space to more realistic road networks.

Acknowledgments

This work would not have been possible without Dr. Bruce Golden's ability to develop such interesting problems. Furthermore, we would like to thank all of the students in Dr. Golden's BMGT831 course for their cooperation and insight. The credit for so aptly naming the Close Enough Traveling Salesman Problem belongs to Damon's father.

References

1. B. Davis, T. Retchless and N. Vakili. Radio-frequency adoption for home meter reading and vehicle routing implementation. Final project for BMGT831 R. H. Smith School of Business. University of Maryland. Spring 2005.
2. J. Dong, N. Yang and M. Chen. A heuristic algorithm for a TSP variant: Meter reading shortest tour problem. Final project for BMGT831 R. H. Smith School of Business. University of Maryland. Spring 2005.
3. T. Gala, B. Subramanian and C. Wu. Meter reading using radio frequency identification. Final Project for BMGT831 R. H. Smith School of Business. University of Maryland. Spring 2005.
4. T. Gonzalez. Covering a set of points in multidimensional space. *Proceedings of the Twenty-eighth Annual Allerton Conference on Communications, Control, and Computing*, Oct 1990.
5. D.J. Gulczynski, J.W. Heath and C.C. Price. Close enough traveling salesman problem: Close only counts in horseshoes, and hand grenades... and meter reading? Final Project for BMGT831 R. H. Smith School of Business. University of Maryland. Spring 2005.
6. G. Gutin and A.P. Punnen, *The Traveling Salesman Problem and Its Variations*. Springer, 2002.

7. S. Hamdar, A. Afshar and G. Akar. Sweeping circle heuristic. Final Project for BMGT831 R. H. Smith School of Business. University of Maryland. Spring 2005.
8. W. Mennell, S. Nestler, K. Panchamgam and M. Reindorp. Radio frequency identification vehicle routing problem. Final Project for BMGT831 R. H. Smith School of Business. University of Maryland. Spring 2005.
9. M. Solomon. Algorithms for the vehicle routing and scheduling problem with time window constraints. *Operations Research*, 35:254-265, 1987.

A Appendix: Psuedocode for Iterative Methods

Heuristic: Steiner Zone

1. Enumerate the c-nodes $c_1, c_2, ..., c_n$.
2. $C' := \{c_1, ..., c_n\}$ (c-nodes)
3. $I := \{1, ..., n\}$ (index set)
4. $S := \{depot\}$ (supernode set)
5. **repeat**
6. $k := \min\{i \in I\}$
7. $ZLIST := \emptyset$ (list of Steiner zones covering c_k)
8. **for** all $j \in I$ with $j > k$
9. **if** $D(c_k, r) \cap D(c_j, r) = \emptyset$ (only non-empty intersections are considered)
10. **if** $D(c_j, r) \cap Z = \emptyset$ for all $Z \in ZLIST$ (no bigger Steiner zones covering c_k and c_j can be built)
11. Add $Z(c_k, c_j)$ to $ZLIST$
12. **else** (bigger zones can be built)
13. Let $Z := Z(c_k, c_{i(1)}, ..., c_{i(r-1)})$ be the Steiner zone of largest degree in $ZLIST$ such that $Z(c_k, c_j) \cap Z \neq \emptyset$
14. Add $Z(c_k, c_{i(1)}, ..., c_{i(r-1)}, c_j)$ to $ZLIST$ (a zone of degree $r+1$ is built)
15. **if** $r \leq 4$ (for reasonable runtime we only store and subsequently build off zones of at most degree four)
16. Add all sub-Steiner zones of $Z(c_k, c_{i(1)}, ..., c_{i(r-1)}, c_j)$ containing c_k and c_j to $ZLIST$ (i.e. if $r = 3$, $Z(c_k, c_{i(1)}, c_{i(2)}, c_j)$ then $Z(c_k, c_j)$, $Z(c_k, c_{i(1)}, c_j)$ and $Z(c_k, c_{i(2)}, c_j)$ are all added; this is the seedbed of small degree zones from which zones of larger degree are built)
17. **end-if**
18. **end-if-else**
19. **end-if**
20. **end-for**
21. Let $Z := Z(c_k, c_{i(1)}, ..., c_{i(m-1)})$ be the zone of largest degree in $ZLIST$
22. Add a point $z \in Z$ to S
23. Remove $c_k, c_{i(1)}, ..., c_{i(m-1)}$ from C'
24. Remove $k, i(1), ..., i(m-1)$ from I
25. **until** $C' = \emptyset$.

Heuristic: Sweeping Circle

1. Generate distance matrix, $D = [d_{ij}]$, on c-nodes
2. Initialize $C' := C$, $S := \{depot\}$.
3. **repeat**
4. Calculate $d_{min} = min(\sqrt{2}r, min(\{d_{ij}\}))$.
5. Cover the plane with circles of radius r translated vertically and horizontally by integer multiples of d_{min}.
6. Set P :=centers of the circles.
7. Assign each node in P a *degree* equal to the number of c-nodes within radius r of that node.
8. Add p to S where p is the node in P of highest degree.
9. Remove from C' the c-nodes within r of p.
10. Update the degrees of P by considering only those c-nodes in C'.
11. **until** $C' = \emptyset$.

Heuristic: Radial Adjacency

1. Construct the adjacency matrix on the c-nodes.
2. Calculate the degree of each c-node.
3. Set $C' := C$, $S := \{depot\}$.
4. **repeat**
5. Select $k \in C'$ with the largest degree.
6. Calculate k^*, the geometric mean of k and all c-nodes adjacent to k.
7. **If** degree(k^*) $\geq$ degree(k)
8. Add k^* to S and remove all vertices adjacent to k^* from C'.
9. **Else**
10. Add k to S and remove k and all vertices adjacent to k from C'.
11. **End If-Else**
12. **until** $C' = \emptyset$.

Twinless Strongly Connected Components

S. Raghavan

Robert H. Smith School of Business & Institute for Systems Research
Van Munching Hall
University of Maryland
College Park, MD 20742
raghavan@umd.edu

Summary. Tarjan [9] describes how depth first search can be used to identify Strongly Connected Components (SCC) of a directed graph in linear time. It is standard to study Tarjan's SCC algorithm in most senior undergraduate or introductory graduate computer science algorithms courses. In this paper we introduce the concept of a *twinless strongly connected component* (TSCC) of a directed graph. Loosely stated, a TSCC of a directed graph is (i) strongly connected, and (ii) remains strongly connected even if we require the deletion of arcs from the component, so that it does not contain a pair of *twin arcs* (twin arcs are a pair of bidirected arcs (i, j) and (j, i) where the tail of one arc is the head of the other and vice versa). This structure has diverse applications, from the design of telecommunication networks [7] to structural stability of buildings [8]. In this paper, we illustrate the relationship between 2-edge connected components of an undirected graph—obtained from the strongly connected components of a directed graph—and twinless strongly connected components. We use this relationship to develop a linear time algorithm to identify all the twinless strongly connected components of a directed graph. We then consider the augmentation problem, and based on the structural properties developed earlier, derive a linear time algorithm for the augmentation problem.

Key words: Digraph augmentation; strong connectivity; linear time algorithm.

1 Introduction

Let $D = (N, A)$ be a directed graph (digraph) with node set N and arc set A. A pair of nodes x and y are *twinless reachable* if there exists a directed path from node x to node y, and a directed path from node y to node x, such that for every arc (i, j) contained in the path from node x to node y, the path from node y to node x does not contain arc (j, i). The *Twinless Strongly Connected Components* (TSCCs) of a digraph are the equivalence classes of nodes under the "twinless reachable" condition (we will show later that the twinless reachable condition defines an equivalence relationship). We

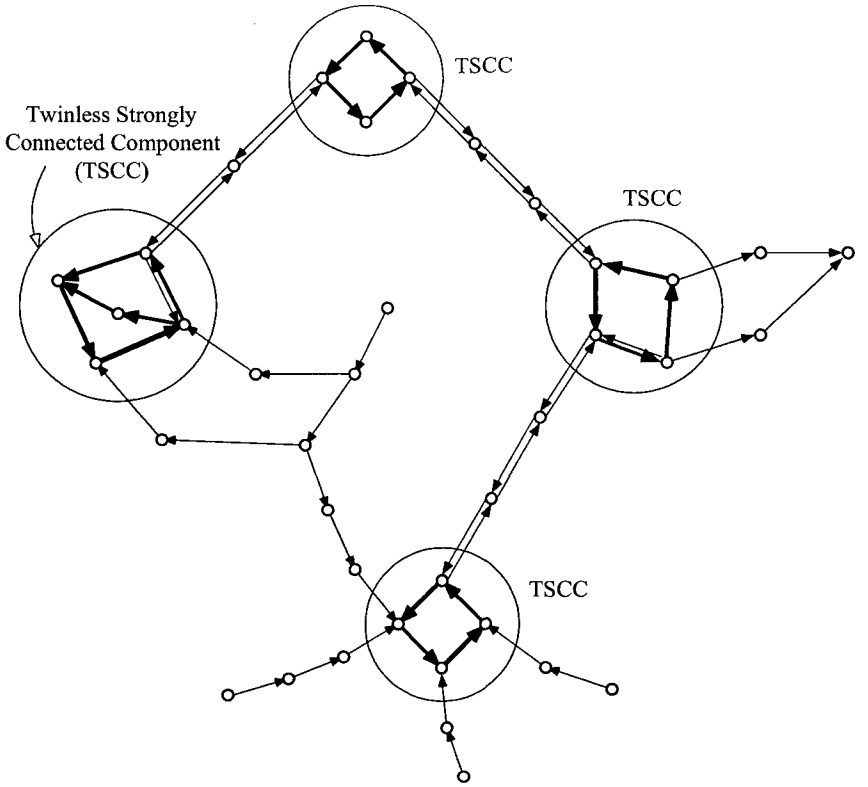


Fig. 1. Twinless Strongly connected components of a digraph. Bold arcs show twinless arcs that form a strongly connected component.

say a digraph is *Twinless Strongly Connected* if every pair of nodes is twinless reachable.

We now provide a slightly different, but equivalent, definition of twinless strongly connectedness. We say that a pair of bidirected arcs (i, j) and (j, i) are *twins*. Recall that a digraph is strongly connected if it contains a directed path between every pair of its nodes. Our alternate definition then is as follows. A digraph $D = (N, A)$ is *Twinless Strongly Connected* if for some subset A' of A, the digraph (N, A') is strongly connected and A' does not contain an arc together with its twin. A *Twinless Strongly connected component* (TSCC) of a digraph is the node set of a maximal twinless strongly connected subdigraph of D. Figure 1 gives an example of four TSCCs, that contain 3 or more nodes, in a digraph.

It should be apparent that every pair of nodes in a TSCC, as defined by the second definition, are twinless reachable. What may not readily apparent is the

converse. That is, if N^i is a TSCC under the first definition of a TSCC. Then, the subdigraph $D^i = (N^i, A^i)$, where $A^i = \{(x,y)|(x,y) \in A, x \in N^i, y \in N^i\}$, is twinless strongly connected as per the second definition. We will work with the second definition until we show, in the next section, that both definitions are indeed equivalent.

Additionally, when considering digraphs, it is clear that reachability is a transitive property. That is, if there is a directed path from node x to node y, and a directed path from node y to node z, then there is a directed path from node x to node z. It turns out that the twinless reachable property is also transitive, but this is not so obvious. Transitivity of the twinless reachable property means that, if a pair of nodes x and y are twinless reachable, and a pair of nodes y and z are twinless reachable, then the pair of nodes x and z are twinless reachable. Transitivity is necessary to define an equivalence relationship and we will show this property in the next section.

In this paper, we consider the following questions (analogous to those for SCCs) in connection with TSCCs. How do we recognize TSCCs of a digraph? Is it possible to recognize TSCCs of a digraph in linear time? We also consider the (unweighted) augmentation problem. That is, given a digraph $D = (N, A)$ find the minimum cardinality set of arcs A' to add to the digraph so that $D = (N, A \cup A')$ is twinless strongly connected (In a seminal paper Eswaran and Tarjan [4] introduced and solved the augmentation problem for strong connectivity). Our answer to these questions is affirmative. Specifically, we develop linear time algorithms to recognize all TSCCs of a digraph and to solve the augmentation problem.

The remainder of this paper is organized as follows. In Section 2 we first derive some structural properties of TSCCs. Specifically we show a correspondence between TSCCs in a strongly connected digraph and 2-edge connected components of an associated undirected graph. Using these structural properties, in Section 3 we describe a linear time algorithm for identifying the TSCCs of a digraph. In Section 4 we consider the augmentation problem and show how to solve the unweighted version of the augmentation problem in linear time. In Section 5 we describe some applications of this structure—one in telecommunications, and one in determining the structural rigidity of buildings. Finally, in Section 6 we discuss a connection between the notion of a twinless strongly connected digraph and strongly connected orientations of mixed graphs.

2 TSCCs of a Strongly Connected Digraph and 2-Edge Connected Components

We now derive some structural properties of TSCCs in a strongly connected digraph D. For ease of exposition, we introduce some additional notation. The *TSCC induced digraph of* D is the digraph $D^{\text{TSCC}} = (N^{\text{TSCC}}, A^{\text{TSCC}})$ obtained by contracting each TSCC in D to a single node. We replace parallel

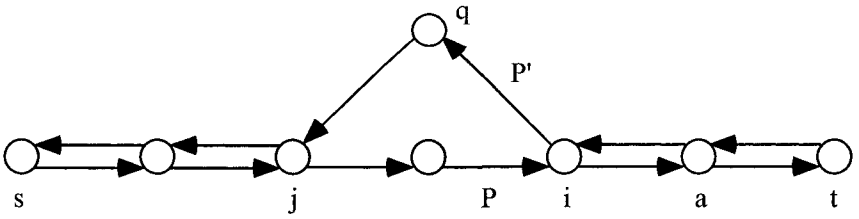


Fig. 2. Proof that the paths P and P' are twin paths.

arcs that the contraction creates by a single arc. Every node in the TSCC induced digraph D^{TSCC} corresponds to a TSCC in the original digraph D. Consequently, for any node $i \in N^{\text{TSCC}}$ we refer to the TSCC node i corresponds to in the original digraph, including the arcs and nodes in the TSCC, as TSCC(i).

For any digraph $D = (N, A)$, the *associated undirected graph* $G(D) = (N, E)$ is a graph with edges $E = \{\{i, j\} : (i, j) \in A \text{ and/or } (j, i) \in A\}$. If (i, j) belongs to A, we refer to the edge $\{i, j\}$ in E as an *image* of this arc. We say that two paths P and P' are *twin paths* if P is a path from node i to node j, and P' is a path from node j to node i that uses exactly the reversal (i.e., the twin) of each arc on the path P.

We first prove a useful property concerning the structure of directed paths between TSCCs in a strongly connected digraph.

Theorem 1 (Twin-arc). *Let $D = (N, A)$ be any strongly connected digraph and let D^{TSCC} be its TSCC induced digraph. The associated undirected graph of D^{TSCC} is a tree. Moreover, every edge in the associated tree is the image of a pair of twin arcs (and no other arcs) of D.*

Proof:
First, consider the TSCC induced subdigraph D^{TSCC} (note that since D is strongly connected, so is D^{TSCC}). We show that D^{TSCC} contains a twin path and no other path between any two nodes. As a result, the associated undirected graph of D^{TSCC} is a tree. We will establish this result by contradiction. Assume the digraph D^{TSCC} contains a path P from a node s to a node t and a path P' from node t to node s that are not twin paths. Let arc (i, q) be the first arc on P' that does not have a twin arc on the path P and let j be the first node following node i on the path P' that lies on the path P (see Figure 2). Then all nodes on P' between nodes i and j and all nodes on P between nodes j and i are twinless strongly connected and thus in the same TSCC. In other words, nodes i and j do not correspond to maximal twinless strongly connected subdigraph's of D (i.e., TSCCs). But, D^{TSCC} is obtained by contracting TSCCs in D and thus each node in D^{TSCC} is a TSCC. We now have a contradiction.

We now show that every pair of twin arcs in D^{TSCC} corresponds to a pair of twin arcs and no other arcs of D. As a result, every edge in the associated tree (i.e., $G(D^{\mathrm{TSCC}})$) is the image of a pair of twin arcs and no other arcs of D. Consider any two adjacent nodes in D^{TSCC}, say nodes a and t. Node t and node a correspond to TSCCs (possibly single nodes) in the expanded digraph (i.e., D). If the original (expanded) digraph contains two non-twin arcs (i, j) with $i \in \mathrm{TSCC}(a)$ and $j \in \mathrm{TSCC}(t)$ and (k, l) with $k \in \mathrm{TSCC}(t)$ and $l \in \mathrm{TSCC}(a)$, then the digraph obtained by the union of $\mathrm{TSCC}(a)$, $\mathrm{TSCC}(t)$ and the arcs (i, j) and (k, l) is twinless strongly connected, and we have a contradiction. Therefore, only a single arc (a, t) connects $\mathrm{TSCC}(a)$ to $\mathrm{TSCC}(t)$ and only a single arc, the twin of arc (a, t), joins $\mathrm{TSCC}(t)$ to $\mathrm{TSCC}(a)$. □

Since D^{TSCC} has the structure of a bidirected tree (that is, a tree with twin arcs in place of each edge [see Figure 3]) when D is strongly connected; we refer to D^{TSCC} as the *TSCC tree.*

Theorem 1 implies the following result concerning the relationship between a strongly connected digraph and its associated undirected graph.

Theorem 2. *The associated undirected graph $G(D)$ of a strongly connected digraph D is 2-edge connected if and only if D is a twinless strongly connected digraph.*

Proof:
If D is a twinless strongly connected digraph, then its associated undirected graph $G(D)$ must be 2-edge connected. Otherwise, if $G(D)$ is not 2-edge connected, deleting some edge $\{i, j\}$ from $G(D)$ disconnects the graph. In D, this edge corresponds to arc (i, j) or arc (j, i) or both arcs (i, j) and its twin (j, i). Eliminating these arcs destroys any directed path between nodes i and j. Consequently D is not twinless strongly connected; a contradiction.

To complete the proof, we show that if the associated undirected graph $G(D)$ is 2-edge connected, then D is a twinless strongly connected digraph. Suppose this is not true. Then $G(D)$ is 2-edge connected while D is not a twinless strongly connected digraph. Consider the TSCC tree of D. If D is not twinless strongly connected then its TSCC tree contains at least two nodes. If the TSCC tree contains 2 or more nodes, then its associated undirected graph (a tree) has at least one edge. Deleting an edge on this graph disconnects it. Since an edge on the associated undirected graph of the TSCC tree is an image of twin arcs and no other arcs in D, deleting the same edge in $G(D)$ disconnects $G(D)$. But then $G(D)$ is not 2-edge connected, resulting in a contradiction. Consequently, the TSCC tree is a single node and D is a twinless strongly connected digraph. □

Theorem 1 and Theorem 2 imply the following characterization of TSCCs in a strongly connected digraph.

Corollary 1 *The 2-edge-connected components of the associated undirected graph $G(D)$ of a strongly connected digraph D correspond in a one to one fashion with the TSCCs of D.*

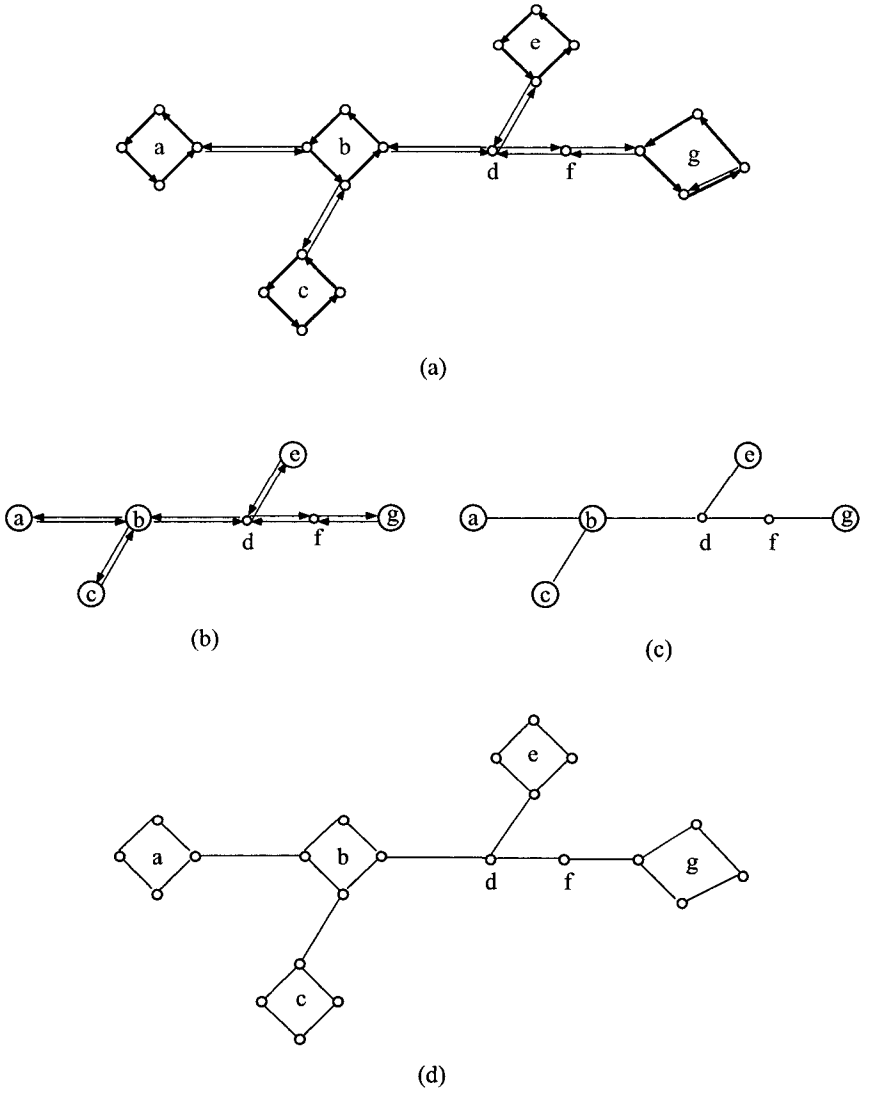


Fig. 3. Illustration of Theorem 1, Theorem 2, and Corollary 1. (a) Strongly connected digraph D, (b) D^{TSCC} the TSCC induced subdigraph of D (the TSCC tree), (c) Associated undirected graph of D^{TSCC}, (d) Associated undirected graph of D.

Notice that Corollary 1 assures us that the TSCCs of a digraph are uniquely defined. Also, from Theorem 1 it follows that the twinless reachable property is transitive.

Lemma 1 *Twinless reachability is a transitive property.*

Proof:
Suppose nodes a and b in a digraph are twinless reachable, and nodes b and c in the same digraph are twinless reachable. It immediately follows that nodes a, b, and c must all be in the same strongly connected component of the digraph. Consider the strongly connected component that contains nodes a, b, and c. From Theorem 1 it follows that nodes a and b must be in the same TSCC, and nodes b and c must be in the same TSCC. But that means nodes a and c are in the same TSCC. From the second definition of twinless strongly connectedness it follows that nodes a and c must be twinless reachable. □

Lemma 1 also shows that the twinless reachable condition defines an equivalence relationship. A binary relationship defines an equivalence relationship if it satisfies reflexivity, symmetry and transitivity. By definition twinless reachability satisfies reflexivity and symmetry, while Lemma 1 shows transitivity proving that it defines an equivalence relationship.

The proof of Lemma 1 also shows the equivalence of the two definitions.

Lemma 2 *The two definitions of a TSCC are equivalent.*

Proof:
It is readily apparent that all node pairs in a TSCC under the second definition are twinless reachable. The proof of Lemma 1 shows any pair of nodes that are twinless reachable must be in the same TSCC (as defined by the second definition of a TSCC). □

The previous lemmas also allow us to show that nodes on any directed path between two nodes in a TSCC are also in the TSCC. This will be useful to us when we consider augmentation problems.

Lemma 3 *Let D be any twinless strongly connected digraph, and P_{ij} be any directed path from node i to j with $i, j \in D$. Then $D_P = D \cup P_{ij}$ is a twinless strongly connected digraph.*

Proof:
Clearly D_P is strongly connected. Consider $G(D_P)$. From Theorem 2 $G(D)$ is 2-edge connected. Thus $G(D) \cup G(P_{ij})$ is also 2-edge connected. But $G(D) \cup G(P_{ij}) = G(D_P)$, showing $G(D_P)$ is 2-edge connected. Thus, by Theorem 2 D_P is also twinless strongly connected. □

3 Identifying Twinless Strongly Connected Components in Linear Time

With the characterization of the relationship between TSCCs in a strongly connected digraph and 2-edge connected components of the associated undirected graph it is now easy to develop a linear time algorithm (based on depth

first search) to identify all TSCCs. The first step consists of finding all strongly connected components of the directed graph. As noted in the outset of the paper this is easily done in linear time using depth first search [9]. A singleton node constituting a SCC of the digraph is also a TSCC of the digraph. If a SCC has cardinality 2, i.e., it consists of 2 nodes, then each node in the SCC is a TSCC. For each of the remaining SCCs (i.e., ones with cardinality greater than or equal to 3) we construct the strongly connected digraph (defined by the arcs between nodes of the SCC) and identify the TSCCs on the SCC.

Corollary 1 states that to identify the TSCCs of a strongly connected digraph, it is sufficient to identify all 2-edge-connected components of its associated undirected graph. Let D_S denote a strongly connected component of D. Consequently, we can convert D_S to its associated undirected graph G_S in $\mathcal{O}(|N| + |A|)$ time, and use the well-known method for identifying all 2-edge-connected components that is also based on depth first search (see exercise 23.2 in [3]).

4 The Augmentation Problem

In this section we consider the problem of augmenting a digraph so that it is twinless strongly connected. As mentioned in the introduction to this paper, Eswaran and Tarjan [4] introduced the augmentation problem. They showed how to minimally augment a digraph in linear time so that it is strongly connected. They also showed how to minimally augment an undirected graph, in linear time, so that it is 2-edge connected.

Our procedure to augment a digraph so that it is twinless strongly connected is roughly as follows. We first apply Eswaran and Tarjan's augmentation procedure to strongly connect the digraph. From Theorem 2, it follows that this strongly connected digraph is twinless strongly connected if and only if its associated undirected graph is 2-edge connected. Consequently, we can apply Eswaran and Tarjan's augmentation procedure (implicitly) to the associated undirected graph to determine the edges to add to make it 2-edge connected. In the corresponding digraph, we add an arc corresponding to each edge added, arbitrarily choosing a direction for the arc in the digraph. Theorem 2 assures us that this procedure gives a twinless strongly connected digraph. We will show that our procedure in fact works (i.e., adds the fewest number of arcs) if the digraph D is *carefully* modified by deleting certain carefully chosen arcs. As a result we present a linear time algorithm to solve the augmentation problem for twinless strong connectivity.

Since our procedure is based on Eswaran and Tarjan's augmentation algorithms we briefly review their procedures.

4.1 Augmenting for Strong Connectivity

Let $D = (N, A)$ be a directed graph, and define $D^{\text{SCC}} = (N^{\text{SCC}}, A^{\text{SCC}})$ to be the SCC induced digraph of D that is obtained by contracting each SCC in

D to a single node. We replace parallel arcs that the contraction creates by a single arc. It is well-known (and straightforward) that D^{SCC} is acyclic.

Eswaran and Tarjan show that it is sufficient to focus attention to the augmentation problem on the SCC induced digraph. To be specific let β be a mapping from N^{SCC} to N defined as follows. If $x \in N^{\mathrm{SCC}}$ then $\beta(x)$ defines any node in the strongly connected component of D corresponding to node x. They show that if A^{ASC} is a minimal set of arcs whose addition strongly connects D^{SCC}, then $\beta(A^{\mathrm{ASC}}) = \{(\beta(x), \beta(y)) | (x, y) \in A^{\mathrm{ASC}}\}$ is a minimal set of arcs whose addition that strongly connects D.

In the acyclic digraph D^{SCC}, a source is defined to be a node with outgoing but no incoming arcs, a sink is defined to be a node with incoming but no outgoing arcs, and an isolated node is defined to be a node with no incoming and no outgoing arcs. Let S, T and Q denote the sets of source nodes, sink nodes, and isolated nodes respectively in D^{SCC}, and assume without loss of generality $|S| \leq |T|$.

Eswaran and Tarjan's procedure finds an index r and an ordering $s(1), \ldots$ $\ldots, s(|S|)$ of the sources of D^{SCC} and $t(1), \ldots, t(|T|)$ of the sinks of D^{SCC} such that

1. there is a path from $s(i)$ to $t(i)$ for $1 \leq i \leq r$;
2. for each source $s(i)$, $r+1 \leq i \leq |S|$ there is a path from $s(i)$ to some $t(j)$, $1 \leq j \leq r$; and
3. for each sink $t(j)$, $r+1 \leq j \leq |T|$, there is a path from some $s(i)$, $1 \leq i \leq r$, to $t(j)$.

They show that a minimal augmentation of D^{SCC} is obtained from the arc set

$$A^{\mathrm{ASC}} = \{(t(i), s(i+1)) | 1 \leq i < r\} \cup \{(t(i), s(i)) | r+1 \leq i \leq |S|\} \cup \begin{cases} (t(r), s(1)) & \text{if } |Q| = 0 \text{ and } |S| = |T| \\ (t(r), t(|S|+1)) \cup \{(t(i), t(i+1)) | |S|+1 \leq i < |T|\} & \\ \cup (t(|T|), s(1)) & \text{if } |Q| = 0 \text{ and } |S| < |T| \\ (t(r), t(|S|+1)) \cup \{(t(i), t(i+1)) | |S|+1 \leq i < |T|\} & \\ \cup (t(|T|), q(1)) \cup \{(q(i), q(i+1)) | 1 \leq i < |Q|\} & \\ \cup (q(|Q|), s(1)) & \text{otherwise} \end{cases}$$

Notice that the augmenting set contains $|T| + |Q|$ arcs. Since there are $|S| + |Q|$ nodes with no incoming arcs, at least $|S| + |Q|$ arcs are needed to augment D^{SCC} so that it is strongly connected. Similarly, as there are $|T|+|Q|$ nodes with no outgoing arcs, at least $|T| + |Q|$ arcs are needed to augment D^{SCC} so that it is strongly connected. Thus $\max(|S|, |T|) + |Q|$ arcs is a lower bound on the number of arcs needed to augment D^{SCC} so that it is strongly connected.

We now show that the addition of the arcs in A^{ASC} strongly connects the digraph D^{SCC}. Actually we show a stronger result that the addition of these arcs makes D^{SCC} twinless strongly connected. Note however this does not mean adding $\beta(A^{\mathrm{ASC}})$ to D makes it twinless strongly connected.

Lemma 4 *When* $|N^{SCC}| > 2$ *the augmented digraph* $D^{ASC} = (N^{SCC}, A^{SCC} \cup A^{ASC})$ *is twinless strongly connected.*

Proof:
Observe that in D^{SCC} any path from a source to sink does not contain a source, sink, or isolated node as an intermediate node. Further D^{SCC} is acyclic and so does not contain twin arcs. Consider the pair $(t(i), s(i+1))$ for any $1 \leq i \leq r$. By design there is a path from $t(i)$ to $s(i)+1$ (i.e., the arc $(t(i), s(i+1))$). Additionally, the addition of arcs has created a directed path from $s(i+1)$ to $t(i)$ that does not use $(s(i+1), t(i))$. The path is defined by following the path from $s(i+1) \rightsquigarrow t(i+1) \rightarrow s(i+2) \rightsquigarrow t(i+2) \rightarrow \ldots \rightsquigarrow t(r) \rightarrow t(|S|+1) \rightarrow t(|S|+2) \rightarrow \ldots \rightarrow t|T| \rightarrow q(1) \rightarrow q(|Q| \rightarrow s(1) \rightsquigarrow t(1) \rightarrow s(2) \rightsquigarrow \ldots \rightarrow s(i) \rightsquigarrow t(i))$.[1] Consequently $s(i+1)$, $t(i)$, and all other nodes on the path from $s(i+1)$ to $t(i)$ are twinless reachable and thus in the same TSCC. Consequently, $s(1), \ldots, s(r), t(1), \ldots, t(r), t(|S|+1), \ldots, t(|T|), q(1), \ldots, q(|Q|)$ are in the same TSCC.

Now consider $s(i)$ and $t(i)$ for any $r+1 \leq i \leq |S|$. Augmentation has created a directed path from $t(i)$ to $s(i)$. Observe, there is a directed path from $s(i)$ to some $t(j), 1 \leq j \leq r$, and a path from some $s(k), 1 \leq k \leq r$, to $t(i)$. The augmentation also creates a directed path from $t(j)$ to $s(k)$ (by following the path described in the first part of the proof). Thus $s(i) \rightsquigarrow t(j) \rightsquigarrow s(k) \rightsquigarrow t(i)$ is a directed path from $s(i)$ to $t(i)$ that does not use arc $(s(i), t(i))$. Consequently $s(i)$ and $t(i)$, and all other nodes on the path from $s(i)$ to $t(i)$ are twinless reachable. This means that $s(1), \ldots, s(|S|), t(1), \ldots, t(|T|), q(1), \ldots, q(|Q|)$ are in the same TSCC.

Finally observe that every node in D^{SCC} that is not a source, sink, or an isolated node is on a path from a source to a sink. Thus, using Lemma 3, D^{ASC} is a twinless strongly connected digraph. □

4.2 Augmenting a Strongly Connected Digraph so that it is Twinless Strongly Connected

We now consider the following problem. Given a strongly connected digraph how do we minimally augment it so that it is twinless strongly connected. Before we begin we first make the following observation that immediately follows from the transitivity of the twinless reachable condition.

Property 1 *Let* γ *be a mapping from* N^{TSCC} *to* N *defined as follows. If* $x \in N^{TSCC}$ *then* $\gamma(x)$ *defines any node in the twinless strongly connected component of* D *corresponding to node* x. *If* A^{ATSC} *is a set of arcs whose addition twinless strongly connects* D^{TSCC}, *then* $\gamma(A^{ATSC}) = \{(\gamma(x), \gamma(y)) | (x, y) \in A^{ATSC}\}$ *is a set of arcs whose addition twinless strongly connects* D.

[1] We denote a path from node i to j by $i \rightsquigarrow j$, and an arc from node i to node j by $i \rightarrow j$.

Observe however that the converse is not true. Consequently, this property does not immediately show that it suffices to focus on D^{TSCC}.

Suppose D is a strongly connected digraph. Recall that D^{TSCC}, the TSCC induced digraph has the structure of a bidirected tree. Consider the set of leaf nodes L (also referred to as leaf TSCCs) of this TSCC tree and observe a leaf node on the TSCC tree has a pair of twin arcs directed into and out of it (referred to as twin leaf arcs). Consequently to make D twinless strongly connected we need to either add an arc directed into the leaf TSCC or out of the leaf TSCC that is not the twin of the twin leaf arcs. Since there are $|L|$ leaf TSCCs, we need at least $\lceil \frac{|L|}{2} \rceil$ arcs to make D twinless strongly connected.

The procedure to make D twinless strongly connected is as follows. Consider the associated undirected graph of D^{TSCC}. Recall since D is strongly connected, $G(D^{\mathrm{TSCC}})$ is a tree. Select one of the leaf nodes and perform DFS from this node. Number the leaf nodes of the TSCC tree in the order they are visited in the DFS procedure, and let $l(1), \ldots, l(|L|)$ denote this ordering. Augment D^{TSCC} with the arc set $A^{\mathrm{ATSC}} = \{(l(i), l(i + \lfloor |L|/2 \rfloor)) | 1 \leq i \leq \lceil |L|/2 \rceil\}$.

Lemma 5 *D^{ATSC} is Twinless Strongly Connected.*

Proof:
Observe that D^{ATSC} is strongly connected and thus by Theorem 2 to prove it is twinless strongly connected it suffices to show that $G(D^{\mathrm{ATSC}})$ is 2-edge connected. Consider $G(D^{\mathrm{TSCC}})$. The procedure described above is exactly Eswaran and Tarjan's procedure to augment a graph so that it is 2-edge connected. In other words Eswaran and Tarjan's procedure adds the edges[2] A^{ATSC} to $G(D^{\mathrm{TSCC}})$ to obtain a 2-edge connected graph $G(D^{\mathrm{ATSC}})$. □

Observe that the procedure adds $\lceil \frac{|L|}{2} \rceil$ arcs. Thus $\gamma(A^{\mathrm{ATSC}})$ minimally augments D so that it is twinless strongly connected.

4.3 Augmenting an Arbitrary Digraph so that it is Twinless Strongly Connected

We now describe how to put together the two procedures *carefully* so that the number of arcs added by the augmentation procedure is minimal. In particular, we will modify the digraph D by deleting certain arcs[3] so that when the two procedures are applied in sequence, the number of arcs added in the augmentation procedure is minimal.

First, we define some notation that we need to describe our procedure. Let $v \in N^{\mathrm{TSCC}}$. Then $\theta(v)$ is the node in D^{SCC} corresponding to the strong component in D^{TSCC} that contains v. For every $w \in N^{\mathrm{SCC}}$, $\psi(w)$ is the set of

[2] With a slight abuse of notation we use A^{ATSC} to denote the edges corresponding to the arcs in the set.

[3] We note that if $\hat{D}$ is a digraph obtained by deleting arcs in D, then an arc set that augments $\hat{D}$ so that it is twinless strongly connected also augments D to be twinless strongly connected.

nodes in the strong component of D^{TSCC} corresponding to node w. We will call a TSCC in D with exactly one incoming arc and one outgoing arc that are twin arcs a leaf TSCC, and refer to the pair of twin arcs directed into and out of a leaf TSCC as twin leaf arcs. It is fairly straightforward to see that in D^{TSCC}, a leaf TSCC is a node (i.e., TSCC) with exactly one incoming arc and one outgoing arc that are twin arcs.

Let L denote the set of leaf TSCCs in D^{TSCC}. As before, let S, T, and Q, denote the set of source nodes, sink nodes, and isolated nodes in D^{SCC}. We further classify the sources in D^{SCC}, based on whether or not they contain a leaf TSCC when expanded (to its constituent TSCCs) in D^{TSCC}. We denote the set of sources in D^{SCC} that do not contain a leaf TSCC as S^o. Similarly, we denote the set of sinks in D^{SCC} that do not contain a leaf TSCC as T^o. With respect to the set of isolated nodes Q in D^{SCC}, observe that an isolated node in D^{SCC} either corresponds to an isolated node in D^{TSCC}, or it corresponds to several TSCCs in D^{TSCC} that are strongly connected with each other. In the former case the isolated node does not contain any leaf TSCCs. In the latter case, from the fact that this strongly connected component is isolated (i.e., does not have any arcs directed into or out of it from nodes that are not in the strongly connected component) and the fact that the twinless strong component graph has the structure of a bidirected tree when the underlying digraph is strongly connected, it follows that there must be at least two leaf TSCCs contained in it. Let Q^o denote the set of isolated nodes in D^{SCC} that do not contain leaf TSCCs.

We now show that

$$\max\left\{|S|+|Q|, |T|+|Q|, \left\lceil \frac{|L|+|S^o|+|T^o|+2|Q^o|}{2} \right\rceil \right\}, \tag{1}$$

is a lower bound for the number of arcs needed to augment D so that it is strongly connected.

Claim 1 *Equation 1 describes a lower bound on the number of arcs needed to augment D so that it is strongly connected.*

Proof:
Clearly the number of arcs needed to make D strongly connected is a lower bound on the number of arcs needed to make D twinless strongly connected. Thus, $|S|+|Q|$ and $|T|+|Q|$ are lower bounds.

Observe that to make D twinless strongly connected, for each leaf TSCC, we will need to add an arc directed into or out of the leaf TSCC that is distinct from the twin leaf arcs. Notice further any source node in S^o corresponds to a strongly connected component with no leaf TSCC in D. This strongly connected component has no incoming arc and so requires at least one arc directed into it to make D strongly (or twinless strongly) connected. Similarly, any sink node in T^o corresponds to a strongly connected component with no leaf TSCC in D. This strongly connected component has no outgoing arc and

so requires at least one arc directed out of it to make D strongly (or twinless strongly) connected. Finally, observe that an isolated node in Q^o corresponds to a strongly connected component with no leaf TSCCs in D. This strongly connected component has no outgoing arcs or incoming arcs, and so requires at least one arc directed into it, and one arc directed out of it to make D strongly (or twinless) connected. Putting all of this together, we obtain that at least $\left\lceil \frac{|L|+|S^o|+|T^o|+2|Q^o|}{2} \right\rceil$ are required to make D twinless strongly connected. □

We now describe how to modify D so that when the two procedures described in Sections 4.1 and 4.2 are applied in sequence a minimal augmentation is obtained. To motivate the modifications needed consider the first step of the augmentation procedure (i.e., the application of Eswaran and Tarjan's augmentation procedure that we described in Section 4.1). Consider D^{SCC} and observe that the augmentation procedure adds exactly one incoming arc to the sources $s(1), \ldots, s(|S|)$ in N^{SCC}, adds exactly one outgoing arc from the sinks $t(1), \ldots, t(|S|)$, and adds one arc directed into and one arc directed out of the isolated nodes $q(1), \ldots, q(|Q|)$. If any $\psi(s(i))$ contains a leaf TSCC, then ensuring that the arc directed into $\psi(s(i))$ is directed into the leaf TSCC takes care of the leaf TSCCs requirement while simultaneously ensuring that an arc is added that is directed into the source $s(i)$. Similarly, if any $\psi(t(i))$ $(1 \leq i \leq |S|)$ contains a leaf TSCC, then ensuring that the arc directed out of $\psi(t(i))$ is directed out of the leaf TSCC takes care of the leaf TSCCs requirement while simultaneously ensuring that an arc is added that is directed out of the sink $t(i)$. For the isolated nodes, as noted earlier, $\psi(q(i))$ either contains two or more leaf TSCCs or is a singleton set. In the former case we can direct the arc into $\psi(q(i))$ into one of the leaf TSCCs in $\psi(q(i))$, and the arc out of $\psi(q(i))$ out of a different leaf TSCC in $\psi(q(i))$. In the latter case there is no choice in selecting the node in $\psi(q(i))$.

Finally consider the sink nodes $t(|S|+1), \ldots, t(|T|)$. For each of these sinks the augmentation procedure adds both arcs directed into the sink and out of the sink. If $\psi(t(i))$ contains two or more leaf TSCCs then we may proceed as the isolated nodes, selecting one leaf TSCC in $\psi(t(i))$ for the incoming arc, and another leaf TSCC in $\psi(t(i))$ for the outgoing arc. However, if $\psi(t(i))$ contains none or one TSCC then the augmentation procedure, if applied without adaptation, may add more arcs than necessary (as it adds an arc directed into this sink as well as directed out of the sink). Therein lies the problem (i.e., if $|S| = |T|$ we would not have had this problem).

To get around this problem we now describe how to modify the augmentation procedure. The modification we propose will delete arcs from the digraph D, to obtain a new digraph $\hat{D}$, so that the number of sources is increased. Specifically, we will increase the number of sources by taking a leaf TSCC and deleting its incoming arc. We will do this until the number of sources is equal to the number of sinks, or there are no more leaf TSCCs available for this purpose. We will show that when the two augmentation procedures are

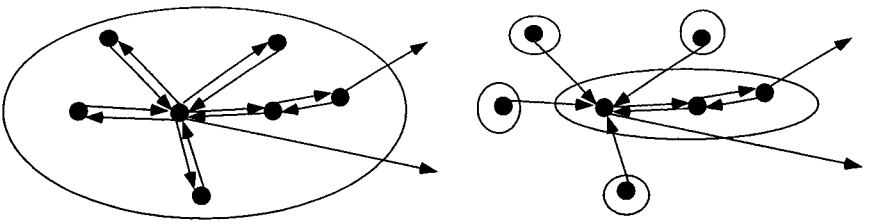


Fig. 4. Leaf TSCCs in a source may be converted into sources by deleting their incoming arcs.

applied in sequence to the digraph $\hat{D}$ the number of arcs added is equal to the lower bound in Equation 1. We now elaborate on how this may be done.

Consider a source $s(i)$ with $\psi(s(i))$ containing x leaf TSCCs. Then the number of sources in D^{SCC} may be increased by $y \leq x - 1$ by taking $y + 1$ leaf TSCCs in $\psi(s(i))$ and deleting their incoming arcs (see Figure 4 for an example). For a sink $t(i)$ with $\psi(t(i))$ containing x leaf TSCCs, the number of sources in D^{SCC} may be increased by $y \leq x-1$ by taking one of the leaf TSCCs and deleting its outgoing arc (creating a sink), and taking y of the remaining TSCCs and deleting their incoming arcs (creating sources). For an isolated node $q(i)$ with $\psi(q(i))$ containing x (≥ 2) leaf TSCCs, we may increase the number of sources by $y \leq x - 1$ and sinks by 1 by taking 1 leaf TSCC and deleting its outgoing arc (creating a sink) and taking y of the remaining leaf TSCCs and deleting their incoming arc. We refer to nodes that are neither, source nodes, sink nodes, or isolated nodes in D^{SCC} as *intermediate* nodes. Consider an intermediate node i in D^{SCC}. If $\psi(i)$ contains x leaf TSCCs then the number of sources may be increased by $y \leq x$ by deleting the incoming arc into y of the leaf TSCCs.

We are now ready to explain how to modify D^{TSCC} and apply the augmentation procedure. The algorithm TSCAUG is summarized in Figure 5. The first step is to identify the strongly connected components of D^{TSCC}, and each leaf TSCC in the strongly connected components. These may be done in linear time following the procedure described in Section 3. The next step is to classify each strongly connected component of D^{TSCC} as a source, sink, isolated, or intermediate strongly connected component. This may also be done in linear time (in fact since the procedure to find TSCCs requires identifying SCCs first it may be done as part of the procedure to identify TSCCs). Next, we consider the strongly connected component of D^{TSCC} one by one, while keeping track of the difference between the number of sinks and sources, to identify the arcs that are to be deleted to create $\hat{D}$. When considering a strongly connected component that is a source the procedure deletes an incoming arc from one leaf TSCC in the strongly connected component. If the number of sinks is greater than the number of sources, it also increases the number of sources by converting leaf TSCCs into sources so that the number of sources is equal

algorithm TSCAUG:

1: Construct D^{TSCC}. Identify each strongly connected component in D^{TSCC} and identify the leaf TSCCs in each strongly connected component of D^{TSCC}.
2. Classify each strongly connected component of D^{TSCC} as a source, sink, isolated, or intermediate strongly connected component.
3. Set k to be the difference between the number of sinks and sources. Consider each strongly connected component of D^{TSCC}.
    If it is a source containing leaf TSCCs
        Delete an incoming arc from one leaf TSCC.
        If $k > 0$, create upto k sources from the remaining leaf TSCCs by deleting their incoming arcs. Update k.
    If it is a sink containing leaf TSCCs
        Delete an outgoing arc from one leaf TSCC.
        If $k > 0$, create upto k sources from the remaining leaf TSCCs by deleting their incoming arcs. Update k.
    If it is isolated containing leaf TSCCs
        Delete an incoming arc from one leaf TSCC and delete an outgoing arc from another leaf TSCC.
        If $k > 0$, create upto k sources from the remaining leaf TSCCs by deleting their incoming arcs. Update k.
    If it is intermediate containing leaf TSCCs
        If $k > 0$, create upto k sources from the leaf TSCCs by deleting their incoming arcs. Update k.
4. Let A^{DEL} denote the arcs deleted from D^{TSCC} in Step 3. Set $\hat{D} = (N, A \backslash \gamma(A^{\mathrm{DEL}}))$. Or $\hat{D}^{\mathrm{TSCC}} = (N^{\mathrm{TSCC}}, A^{\mathrm{TSCC}} \backslash A^{\mathrm{DEL}})$.
5. Apply Eswaran and Tarjan's strong connectivity augmentation algorithm to $\hat{D}$ to obtain the set of arcs $\beta(\hat{A}^{\mathrm{ASC}})$. Set $\hat{D} = (N, \hat{A} \cup \beta(\hat{A}^{\mathrm{ASC}}))$.
6. Apply the algorithm described in Section 4.2 to $\hat{D}$ to obtain $\gamma(\hat{A}^{\mathrm{ATSC}})$. Set $\hat{D} = (N, \hat{A} \cup \gamma(\hat{A}^{\mathrm{ATSC}}))$.
7. Output the arcs $\beta(\hat{A}^{\mathrm{ASC}}) \cup \gamma(\hat{A}^{\mathrm{ATSC}})$.

Fig. 5. Algorithm to solve the twinless strong connectivity augmentation problem.

to the number of sinks, or no more leaf TSCCs remain in the strongly connected component. When considering a strongly connected component that is a sink it deletes an outgoing arc from a leaf TSCC in the strongly connected component, and creates additional sources using leaf TSCCs until the number of sources is equal to the number of sinks, or the leaf TSCCs in the strongly connected component are exhausted. When considering a strongly connected component that is isolated it creates an additional source and an additional sink, and also converts leaf TSCCs to sources as needed. Similarly when con-

sidering a strongly connected component that is intermediate it converts leaf TSCCs into sources if needed. Observe that this procedure considers each leaf TSCC in the digraph once and thus is an $\mathcal{O}(|N|)$ procedure. After obtaining $\hat{D}$ we first apply Eswaran and Tarjan's augmentation procedure to strongly connect $\hat{D}$. Recall that this procedure takes $\mathcal{O}(|A|)$ time. Next we apply the procedure described in Section 4.2 which also takes $\mathcal{O}(|A|)$ time. Since each of the steps in the algorithm take linear time, the procedure is an $\mathcal{O}(|N| + |A|)$ procedure.

We now show that the number of arcs added by the procedure is equal to the lower bound shown in Equation 1.

Claim 2 *The number of arcs added by algorithm TSCAUG is equal to the bound in Equation 1.*

Proof:
When the first step of the procedure is applied the procedure adds $|\hat{T}| + \hat{Q}|$ arcs. Observe that in $\hat{D}^{\text{SCC}}$ no source, sink, or isolated node contains leaf TSCCs. Thus no arc added in the first step is directed into or out of a leaf TSCC in $\hat{D}$. Consequently, the second step adds $\left\lceil \frac{|\hat{L}|}{2} \right\rceil$ arcs. Thus a total of $|\hat{T}| + |\hat{Q}| + \left\lceil \frac{|\hat{L}|}{2} \right\rceil$ arcs are added. Notice that,

$|\hat{T}| = |T| + (|Q| - |Q^o|)$,
$|\hat{Q}| = |Q^o|$,
$|\hat{L}| = \max\{0, |L| - (|S| - |S^o| + |T| - |T^o| + 2(|Q| - |Q^o|) + |T| - |S|)\}$.

So the procedure adds

$$(|T| + |Q| - |Q^o|) + |Q^o| + \left\lceil \frac{\max\{0, |L| - (|S| - |S^o| + |T| - |T^o| + 2(|Q| - |Q^o|) + |T| - |S|)\}}{2} \right\rceil$$

arcs. Or

$$\max\left\{|T| + |Q|, \left\lceil \frac{|L| + |S^o| + |T^o| + 2|Q^o|}{2} \right\rceil \right\}$$

arcs, which is the bound in Equation 1. □

5 Applications and Extensions

We now discuss two applications of the TSCC structure. The first application is in the design of survivable telecommunications networks. Magnanti and Raghavan [7] describe a dual-ascent (primal-dual) algorithm for the network design problem with low connectivity requirements (NDLC). In this problem, we are given an undirected graph $G = (N, E)$, with node set N and edge set E, a cost vector $\mathbf{c} \in \mathcal{R}_+^{|E|}$ on the edges E, and a connectivity requirement $r_i \in (0, 1, 2)$ for each node, and wish to design a minimum cost network

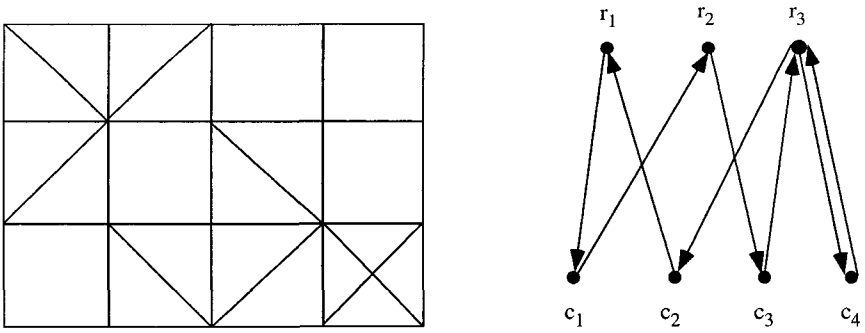


Fig. 6. A 3×4 grid with tension bracings and its corresponding bracing digraph.

that contains $r_{st} = \min(r_s, r_t)$ edge-disjoint paths between nodes s and t. To develop a strong mixed-integer programming formulation for the problem they direct the problem by replacing every edge $\{i, j\}$ by a pair of arcs (i, j) and (j, i), each with the same cost as edge $\{i, j\}$. They show that in the equivalent directed problem, we wish to design a minimum cost directed network that contains a directed path from every node with $r_i = 2$ to every node with $r_j = 1$ or 2. Further, any two nodes i and j with $r_i = r_j = 2$ must be twinless reachable. In their dual-ascent algorithm, it is necessary to identify twinless strongly connected components of a so-called auxiliary digraph (to identify ascent directions). The linear time algorithm described in Section 3 serves this purpose.

A second application, described in Baglivo and Graver [1], arises in determining whether the bracing of an architectural structure is rigid (i.e., will not deform). In this application, we are given an $n \times m$ square grid and wish to brace this structure by placing *tension bracings*. A tension bracing consists of a tension wire that can be compressed but not stretched. Baglivo and Graver show that the question of whether a square grid with tension bracings is rigid may be answered by looking at its associated *bracing digraph*. The bracing digraph is constructed as follows. Create a node for each row and each column of the $n \times m$ square grid (i.e., $n + m$ nodes). If a tension wire is attached to the lower left and upper right hand corners of the cell in the ith row and jth column of the square grid, then create a directed arc from the node representing column j to the node representing row i. On the other hand, if the tension wire is attached to the lower right and upper left hand corners of the cell in the ith row and jth column, then create a directed arc from the node representing row i to the node representing column j. Figure 6 gives an example of this construction. Baglivo and Graver show that a tension bracing of an $n \times m$ square grid is rigid if and only if the bracing digraph is strongly connected.

For aesthetic reasons an architect may want to ensure that no cell in the square grid has two diagonal tension bracings [8]. For an existing structure, a question that then arises is whether tension bracings can be removed from cells that contain two diagonal tension bracings while keeping the structure rigid. It is easy to see this corresponds to the question of whether some subdigraph of the bracing digraph is strongly connected and does not contain a pair of twin arcs. This is precisely the definition of twinless strong connectedness, and again the algorithm in Section 3 can be used to determine whether the bracing digraph is twinless strongly connected.

If the bracing digraph is not twinless strongly connected, then a natural question that arises is how it may be minimally augmented so that it is twinless strongly connected. This is an augmentation problem with a twist: the augmented digraph also needs to be bipartite (arcs need to go from nodes representing rows to columns or vice versa); and the solution to this problem remains an open question. Interestingly, the case when we wish to augment the bracing digraph (preserving bipartiteness) so that it is strongly connected was recently solved in [5].

6 On Strongly Connected Orientations of Mixed Graphs and Twinless Strongly Connected Digraphs

A referee pointed out the connection between the notion of a twinless strongly connected digraph and strongly connected orientations of mixed graphs. We discuss this connection here.

A mixed graph $M = (N, A, E)$ is a graph with node set N that contains both directed arcs A and undirected edges E. A natural question that arises in such graphs is whether the undirected edges can be directed (oriented) so that the resulting digraph is strongly connected. Boesch and Tindall [2] show that a mixed graph has a strongly connected orientation if and only if it is connected[4] and the underlying undirected graph[5] of the mixed graph is 2-edge connected.

The question of whether a digraph D is twinless strongly connected can be translated into a question of whether a mixed graph has a strongly connected orientation as follows. Replace each pair of twin arcs (i, j) and (j, i) in the digraph D by an undirected edge $\{i, j\}$ to obtain a mixed graph $M(D)$. It is easy to observe that the digraph D is twinless strongly connected if and only if $M(D)$ has a strongly connected orientation. Based on this connection it is possible to provide an alternate proof of some of the results in Section 2.

[4] In other words, it is possible to go from any node to any other node between using a sequence of undirected and directed arcs (directed arcs must be traversed in the direction they are oriented).

[5] The underlying undirected graph is obtained by replacing all directed arcs by undirected edges.

After learning about the connection between strong orientations of mixed graphs and twinless strongly connected digraphs, we did a web search and found Gusfield [6] considers the following mixed graph augmentation problem. Given a mixed graph $M = (N, A, E)$ find the minimum number of directed arcs to add to the mixed graph so that it has a strongly connected orientation. By transforming the digraph D to the mixed graph $M(D)$, as described in the previous paragraph, our problem may be solved as a mixed graph augmentation problem on $M(D)$. Alternatively, the mixed graph augmentation problem may be transformed to the problem of augmenting a digraph so that it is twinless strongly connected as follows. First, take each pair of twin arcs (i, j) and (j, i) in the mixed graph M and replace (j, i) by adding a new node k_{ij} to the mixed graph and two arcs (j, k_{ij}) and (k_{ij}, i).[6] Next replace each undirected edge $\{i, j\}$ by a pair of twin arcs (i, j) and (j, i) to obtain the digraph $D(M)$. It is easy to observe that the mixed graph M has a strong orientation if and only if $D(M)$ is twinless strongly connected. Thus the solution to the augmentation problem for the mixed graph M may be obtained by solving the twinless strongly connected augmentation problem on $D(M)$.

Gusfield's augmentation algorithm determines an optimal orientation of the undirected edges of the mixed graph (thus transforming it into a directed graph) after which Eswaran and Tarjan's augmentation algorithm for strong connectivity is applied. This procedure and especially the proofs of correctness are complex and quite intricate (as determining the optimal orientation of the undirected edges is quite challenging!). Our procedure in the context of directed graphs is much simpler, and the proofs of correctness are fairly straightforward. Our procedure nicely illustrates how the sequential application of Eswaran and Tarjan's two augmentation solve the problem of optimally augmenting a graph so it is twinless strongly connected, or alternatively shows how to optimally augment a mixed graph so that it admits a strong orientation.

References

1. J. A. Baglivo and J. E. Graver. *Incidence and Symmetry in Design and Architecture.* Cambridge University Press, Cambridge, UK, 1983.
2. F. Boesch and R. Tindall. Robbin's Theorem for mixed multigraphs. *The American Mathematical Monthly,* 87(9):716–719, 1980.

[6] Observe that the mixed graphs considered by Gusfield may contain both arcs (i, j) and (j, i), in which case nodes i and j are strongly connected. However, without this transformation, in our digraph i and j are not twinless strongly connected. Notice, that it suffices to focus on the mixed graph augmentation problem on this transformed mixed graph. Let $\xi(k_{ij}) = i$ or j for the new nodes added and $\xi(i) = i$ otherwise. The set of arcs A^{AMSC} added to the transformed mixed graph are mapped to arcs on the original mixed graphs by $\xi(A^{\text{AMSC}}) = \{(\xi(x), \xi(y)) | (x, y) \in A^{\text{AMSC}}\}$.

3. T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms.* MIT Press, 1990.
4. K. Eswaran and R. Tarjan. Augmentation problems. *SIAM Journal on Computing*, 5(4):653–665, 1976.
5. H. N. Gabow and T. Jordán. How to make a square grid framework with cables rigid. *SIAM Journal on Computing*, 30:649–680, 2000.
6. D. Gusfield. Optimal mixed graph augmentation. *SIAM Journal on Computing*, 16:599–612, 1987.
7. T. L. Magnanti and S. Raghavan. A dual-ascent algorithm for low-connectivity network design. Technical report, Smith School of Business, University of Maryland, College Park, 2006.
8. A. Recksi. Private Communication at Discrete Optimization 1999, Rutgers University, 1999.
9. R. Tarjan. Depth-first-search and linear graph algorithms. *SIAM Journal on Computing*, 1:146–160, 1972.

Part III

Modeling & Making Decisions

EOQ Rides Again!

Beryl E. Castello and Alan J. Goldman

Department of Applied Mathematics and Statistics
The Johns Hopkins University
Baltimore, MD 21218
castello@ams.jhu.edu, goldman@ams.jhu.edu

Summary. Despite the criticisms of Woolsey [25] and others, the classical "EOQ" model of Inventory Theory remains useful, both pedagogically and as a stepping-stone towards more realistic variants. Prominent among these are: (a) the one in which stock-outs are permitted but penalized, and (b) the one in which deliveries are gradual (i.e., continuous) rather than instantaneous. Here we also deal with a variant containing both features, (a) and (b). (In addition, a less common variant is considered.) In the spirit of Harris's original (1913) presentation, the use of calculus is avoided entirely. Instead, it is shown explicitly how all the variants can be solved by "reduction" to the basic model, which is in turn treated by using the simply-proved "base case" of the arithmetic-geometric mean inequality.

Key words: Economic order/production quantity; inventory management; shortage; algebra; EOQ; EPQ.

1 Introduction

The setting for an inventory cost analysis can be described as follows: the decision maker, on whose behalf the optimization is conducted, can be thought of as the manager of a firm which sells a product. The firm replenishes its stock by reordering more, from either an external wholesaler or a capability for manufacturing the goods internally. The decision maker must determine (1) how often to reorder, and (2) the size of the reorder quantity.

The usual objective of the analysis is to minimize the total cost per unit time associated with inventory. This is comprised of: (1) holding costs (the costs of maintaining or "carrying" inventory), (2) reordering or production costs (the "fixed" costs from individual replenishments of inventory), (3) procurement costs (the "variable" costs of purchasing each unit of inventory), and (4) shortage costs (the cost of subjecting customers to delays from a backlog of unfilled demand).

Inventory models have been studied extensively since Harris first presented the famous economic order quantity (EOQ) formula in 1913 ([17]; see [11] for

historical background). Until recently, almost all the literature concerning the EOQ and related inventory models[1] used differential calculus as the underlying solution technique, though frequently omitting the supplementary reasoning needed to establish optimality. (For example, see [8], [16], [3], [1], [21], and [10].) One early exception occurred in 1970, when Thierauf and Grosse [24] developed the optimal order quantity for the EOQ model using three different techniques: graphical analysis, algebraic analysis (which is later used to solve the EPQ model), and calculus. The authors based their algebraic method on the presumed validity of equating opposing costs (i.e., holding costs versus ordering costs).

In 1996, Grubbström [14] derived the EOQ formula without the use of calculus via an elegant "completing the square" argument. Grubbström and Erdem [15] then obtained the solution to an EOQ model with shortages without the use of derivatives, and proposed that their approach be used to introduce students without a background in calculus to inventory models. In 2001, Cárdenas-Barrón [4] revisited the work of Grubbström and Erdem and extended it to the economic production quantity (EPQ) model with shortages. In 2004, Ronald et al. [22] suggested that the method presented in the previous two articles might still be too sophisticated for some readers; they provided rationale (derivation rather than mere verification) for some of the identities used in [15] and [4]. A 2005 article by Chang et al. [5] attempted to simplify the algebra of the previous three articles, such that "the usual skill of completing the square" could more readily handle the objective functions for the EOQ and EPQ models with shortages.

In this work, we derive optimal solutions for four inventory models – EOQ, EPQ, EOQ with shortages, and EPQ with shortages – again without the use of derivatives. In each case, we begin with a step-by-step development of the objective function. To keep the methods entirely elementary, our approach uses the easily-proved "base case" of the arithmetic-geometric mean inequality to identify the optimal solution for the classical EOQ model. We then show that the three variants can be reduced to this basic EOQ model. The final section includes a similar treatment for a less common variant (see [12]) involving "deferred payment" for the delivered goods; that treatment can also be adapted to Chung and Huang's [7] extension of Goyal's variant from EOQ to EPQ.

[1] Literature-tracing here has been hindered considerably because different authors use different names and/or acronyms for the same scenario, and vice versa. For example, the economic production quantity (EPQ) model is sometimes called the "(economic) production lot size model." Also, the term "economic lot size" has been used for both the EOQ and EPQ models.

2 Notation

In our analyses, we make use of the following notation. (The positive quantities h, D, p, K, M, and s are data of the models.)

- Q the total order (production) quantity,
- h the per unit time inventory holding cost per unit,
- D the per unit time demand,
- p the purchase cost per unit,
- K the fixed purchasing cost per order,
- M the per unit time production,
- B the number of items backordered when production starts (or when an order is delivered),
- s the "penalty" cost per unit per unit time, due to shortages,
- I_M the maximum inventory level,
- T the cycle time (time between successive production starts / order deliveries),
- $I(t)$ the inventory level at time t,
- $W(t)$ the backlog level (waiting outages) at time t.

3 The Economic Order Quantity Model

Figure 1 illustrates the inventory level as a function of time. The cycle time in this model is the time between successive orders. Q units are ordered each time an order is placed, and shortages are not allowed.

The crux of this and the later models is the proper balancing of holding costs against reordering costs. If the order quantity is too high, then the average inventory level will be high, resulting in large holding costs. If the order quantity is too low, then the average inventory level will be low. Orders will need to be placed more frequently, resulting in larger reordering costs.

3.1 Assumptions

Development of the basic EOQ model involves the following assumptions:

1. The inventory consists of a single product.
2. Demand for the item occurs at a known and constant rate.
3. The item has a sufficiently long shelf-life so that no inventory is lost due to spoilage.
4. The lead time (the time between the placement of an order and its arrival) is zero or a known positive constant.
5. The maximum inventory level is not constrained by insufficiency of shelf space.

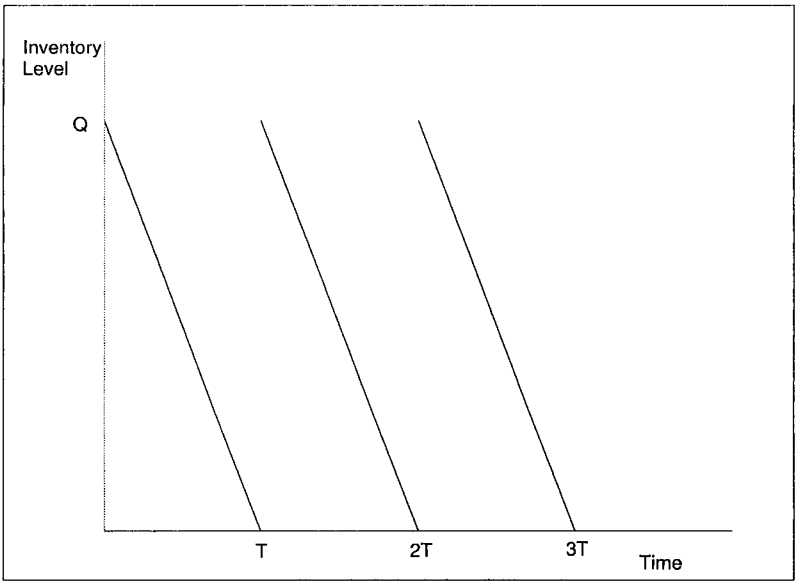


Fig. 1. Inventory Profile for the EOQ Model.

3.2 Model Formulation

Holding Costs. The holding cost during one cycle of length T is "h" times the area under the $I(t)$ curve. (This follows from the representation of the area as a definite integral, itself the limit of finite sums giving discrete approximations of the cycle's holding cost.) Here, that is the area of a triangle with base-length T and altitude $I_M = Q$. Hence, the holding cost during one cycle is $h \times \frac{1}{2}TQ$ and the total holding cost per unit time is $\frac{h}{2}Q$.

Ordering Costs. If the per unit time demand is D and the order quantity is Q, then the number of orders placed per unit time is D/Q. Thus, the per unit time ordering cost is $\frac{D}{Q}K$.

Purchase Costs. To meet demand, the firm must purchase D units per unit time, so the per unit time purchase cost is Dp.

Total Costs. The total per unit time inventory cost for the basic EOQ model is the sum of the per unit time holding costs, ordering costs and purchase costs:

$$f_{EOQ}(Q) = \frac{Q}{2}h + \frac{D}{Q}K + Dp. \tag{1}$$

3.3 Derivation of Optimal Order Quantity Using AM/GM/I

Let $a_1, a_2, \ldots, a_n$ be a finite set of nonnegative real numbers with arithmetic mean

$$am(a_1, a_2, \ldots, a_n) = \frac{a_1 + a_2 + \ldots + a_n}{n}$$

and geometric mean

$$gm(a_1, a_2, \ldots, a_n) = \sqrt[n]{a_1 a_2 \ldots a_n}.$$

The well-known *arithmetic-geometric mean inequality* (AM/GM/I) states that

$$gm(a_1, a_2, \ldots, a_n) \leq am(a_1, a_2, \ldots, a_n),$$

with equality if and only if $a_1 = a_2 = \ldots = a_n$.

To see how AM/GM/I can be used for optimization analyses, suppose the a_i's are variables and we want to minimize their SUM, or equivalently their arithmetic mean, subject to a constraint that fixes the value of their PRODUCT, or equivalently their geometric mean. Then the AM/GM/I gives a constant lower bound on the value of the minimand. Because that lower bound is actually attained (when all the a_i's are equal), it is in fact the desired minimum value, and the unique optimal solution has each a_i equal to the n-th root of the stipulated value of their product.

The most common proofs of the AM/GM/I are calculus-free (cf. the following proof and Footnote 2). Thus, the solution to an optimization problem such as the one described in the preceding paragraph may be obtained using elementary methods, i.e., without the use of calculus.

Proof of AM/GM/I (Special Cases). The inequality is clearly true when $n = 1$. Fortunately, we need only the case $n = 2$, which is easily shown to be true as follows[2]:

$$\begin{aligned}(\sqrt{x_1} - \sqrt{x_2})^2 \geq 0 &\Leftrightarrow x_1 - 2\sqrt{x_1 x_2} + x_2 \geq 0 \\ &\Leftrightarrow x_1 + x_2 \geq 2\sqrt{x_1 x_2} \\ &\Leftrightarrow \frac{1}{2}(x_1 + x_2) \geq \sqrt{x_1 x_2}.\end{aligned}$$

Equality holds at the start if and only if

$$\begin{aligned}(\sqrt{x_1} - \sqrt{x_2})^2 = 0 &\Leftrightarrow \sqrt{x_1} - \sqrt{x_2} = 0 \\ &\Leftrightarrow \sqrt{x_1} = \sqrt{x_2} \\ &\Leftrightarrow x_1 = x_2,\end{aligned}$$

hence it also holds at the end if and only if $x_1 = x_2$. □

[2] The general case can be proved using induction.

Our goal is to determine an optimal order quantity which minimizes Equation (1). But since the third summand Dp of Equation (1) is constant, minimizing Equation (1) is the same as minimizing the quantity

$$Q\left(\frac{h}{2}\right) + \frac{(DK)}{Q},$$

which is proportional to the arithmetic mean of the two nonnegative quantities

$$x_1 = Q\left(\frac{h}{2}\right) \quad \text{and} \quad x_2 = \frac{(DK)}{Q},$$

whose product, $DKh/2$, is constant (free of variable Q). Thus, by the AM/GM/I, their geometric mean, $\sqrt{x_1 x_2}$, provides a constant lower bound for their arithmetic mean. By the preceding paragraph, this constant lower bound is actually attained (so that the arithmetic mean, and thus Equation (1), is minimized) when (and only when!) $x_1 = x_2$, i.e., (uniquely) when

$$Q\frac{h}{2} = \frac{DK}{Q}$$

which happens when

$$Q_{EOQ} = \sqrt{\frac{2DK}{h}}.$$

For later use (e.g., in the proof of Lemma 10), we observe that the minimum value, exclusive of the constant summand Dp in $f_{EOQ}(Q_{EOQ})$, is given by

$$2x_1 = (Q_{EOQ})h = \sqrt{2DKh}.$$

4 The EOQ Model with Shortages Allowed

The next model we consider is similar to the basic EOQ model in that all goods are available for immediate sale when the order arrives. However, the order is not placed when the inventory level reaches zero, but rather we allow customers to be "wait-listed." In the literature this model is sometimes referred to as the *planned shortages model.* We will refer to it as EOQ-SH. Figure 2 illustrates the inventory level as a function of time when Q units are ordered when the number of orders waiting to be filled reaches a certain critical point, B, which must also be determined. These backorders are immediately filled from the newly-arrived order, leaving the maximum inventory level as $I_M = Q - B$. The cycle time in this model is the time between successive orders.

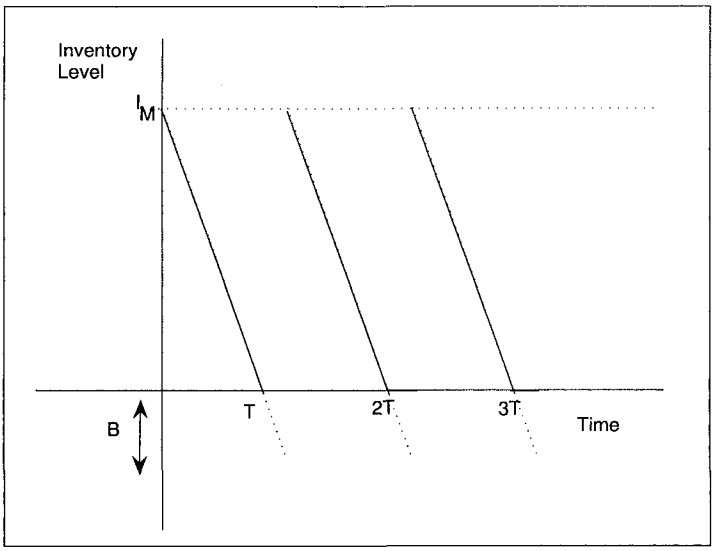


Fig. 2. Inventory Profile for the EOQ Model with Shortages Allowed.

4.1 Assumptions

In addition to the assumptions enumerated in Section 3.1 for the basic EOQ model, development of the EOQ-SH model involves assuming that:

1. No customers will be lost due to shortages.
2. There is a cost associated with allowing shortages. These, reflecting the datum "s" listed at the start of Section 2, are the variable costs in goodwill as customers await backordered items.
3. All backorders are immediately filled from the newly-arrived order.

4.2 Model Formulation

In this model, there are two decision variables: Q, the total order quantity, and B, the number of units on backorder at the time production starts.

It will prove useful to divide the cycle time T into two subintervals: $T_a = [0, t_1]$, the time during which there are no backorders and inventory is decreasing from its maximum level at a rate D; and $T_b = [t_1, T]$, the time during which there are backorders, the level of which is increasing at rate D.

Holding Costs. During one cycle of length T, holding costs accumulate only during T_a. The total holding cost for this time period is "h" times the area

under the $I(t)$ curve for $0 \leq t \leq t_1$. That is the area of a triangle with base-length t_1 and altitude $I_M = Q - B$. Hence, the holding cost during one cycle is $h \times \frac{1}{2} t_1 (Q - B)$ and the total per unit time holding cost is $(h/2)(t_1/T)(Q - B)$.

We note that the total demand during T_a is given by $Dt_1 = Q - B$, while the demand during the entire cycle is $DT = Q$. Division yields $t_1/T = (Q - B)/Q$. Hence, we can write the holding cost per unit time as

$$\frac{h}{2} \frac{(Q - B)^2}{Q}.$$

Shortage Costs. During one cycle of length T, shortage costs accumulate only during T_b. The total shortage cost for this time period is "s" times the area of a triangle with base-length $T - t_1$ and altitude B. Hence, the shortage cost during one cycle is $s \times \frac{1}{2}(T - t_1)B$, and the total per unit time holding cost is $(s/2)[(T - t_1)/T]B$.

Since $t_1/T = (Q - B)/Q = 1 - B/Q$, we have $(T - t_1)/T = B/Q$. Thus, the shortage cost per unit time may be written as

$$\frac{s}{2} \frac{B^2}{Q}.$$

Other Costs. The ordering cost and purchase cost for this model are the same as in the basic EOQ model.

Total Costs. The total inventory cost is given by the function

$$f_{EOQ-SH}(Q, B) = \frac{h}{2} \frac{(Q - B)^2}{Q} + \frac{s}{2} \frac{B^2}{Q} + \frac{D}{Q} K + Dp. \tag{2}$$

4.3 Derivation of Optimal Order Quantity and Maximum Backlog Amount

To determine an optimal order quantity, Q, and maximum backlog amount, B, we minimize the function Equation (2) subject to the constraint $0 < B \leq Q$. Although we are dealing with two variables and a "feasible region" that is not closed, our argument will avoid issues of existence and of sufficient conditions for an optimal solution.

Before presenting that argument in our specific case, we pause to outline its general logic. The aim is to minimize a function $f(Q, B)$ — here, the function given by Equation (2). Suppose that for each Q, we can find a value $B^*(Q)$ of B that minimizes $f(Q, B)$ for this fixed Q. That leads to a function $f(Q, B^*(Q)) = G(Q)$ of Q alone. And suppose, further, that we can then find a value Q^* of Q that minimizes $G(Q)$. To see that this process of sequential ("one variable at a time") optimization is correct, i.e., that $(Q^*, B^*(Q^*))$ is indeed an optimal solution to the original "simultaneous optimization" problem, just observe that for any allowable pair (Q, B),

$$f(Q,B) \geq f(Q,B^*(Q)) = G(Q) \geq G(Q^*) = f(Q^*,B^*(Q^*)),$$

where the first inequality follows from the definition of $B^*(Q)$, and the second one from the definition of Q^*. In what follows, Lemma 1 gives the determination of $B^*(Q)$, i.e., the first stage of "sequential", while Lemma 2 gives the second stage, the determination of Q^*.

Lemma 1. *For a fixed value of Q, Equation (2) is minimized, subject to the constraint $0 < B \leq Q$, when (and only when) $B = h/(h+s)Q$.*

Proof.[3] To clarify how Equation (2) depends on B when Q is fixed, it is natural to rewrite the terms of Equation (2) involving B as

$$\begin{aligned}\frac{h}{2}\frac{(Q-B)^2}{Q} + \frac{s}{2}\frac{B^2}{Q} &= \frac{h+s}{2Q}\left(B^2 - \frac{2hQ}{h+s}B + \frac{h}{h+s}Q^2\right) \\ &= \frac{h+s}{2Q}\left[\left(B - \frac{hQ}{h+s}\right)^2 - \left(\frac{hQ}{h+s}\right)^2 + \frac{h}{h+s}Q^2\right]. \quad (3)\end{aligned}$$

Thus, for fixed Q, Equation (2) is quadratic in B. To minimize $f_{EOQ-SH}(Q,B)$, we must choose B so as to minimize Equation (3), subject to the constraint $0 < B \leq Q$. Clearly Equation (3) is minimized (indeed, with the constraints met) when and only when

$$B = B_{EOQ-SH}(Q) = \frac{h}{h+s}Q.$$

□

Lemma 2. *When $B = B_{EOQ-SH}(Q)$, the function Equation (2) is minimized by choosing*

$$Q = \sqrt{2DK\frac{(h+s)}{hs}}.$$

Proof. When we substitute the expression for $B_{EOQ-SH}(Q)$ into Equation (2), we obtain

$$f_{EOQ-SH}(Q, B_{EOQ-SH}(Q)) = \frac{KD}{Q} + pD + \frac{Q}{2}\frac{hs}{(h+s)}. \quad (4)$$

But Equation (4) is the basic EOQ model Equation (1) with holding cost h replaced by $\hat{h} = (hs)/(h+s)$. Thus, by our previous optimization of Equation (1),

$$Q_{EOQ-SH} = \sqrt{\frac{2DK}{\hat{h}}} = \sqrt{2DK\frac{h+s}{hs}}.$$

[3] As in [15], [4], and [5], our derivation approach uses "square-completion."

We also note that

$$B_{EOQ-SH} = \frac{h}{h+s} Q_{EOQ-SH} = \sqrt{\frac{2DKh}{s(h+s)}}$$

and

$$I_M = Q_{EOQ-SH} - B_{EOQ-SH} = \sqrt{\frac{2DKs}{h(h+s)}}.$$

□

5 The Economic Production Quantity Model

Next we consider an EOQ variant in which items are made available for inventory (or sale) gradually as opposed to all at once. This "gradual delivery" might be a result of either (i) external"stretching-out" because the supplier's production rate or the transportation system's delivery rate is limited, or (ii) internal stretching-out because the retailer has limitations on the speed of handling receipt of a sizable order and getting it "on the shelves" accessible to customers.[4] (For brevity, and because of the predominant terminology noted next, we will refer to this process as "production" rather than "delivery or production".)

In the literature this model is most often referred to as the *economic production quantity* model. Sometimes the term *(economic) production lot size* model is also used. We will simply refer to it as *EPQ*. One of the earliest references to this inventory system may be found in [23]. Figure 3 illustrates the inventory level as a function of time when Q units are produced during each production run. The cycle time in this model is the time between successive production starts. Again, shortages are not allowed.

5.1 Assumptions

In addition to the assumptions enumerated in Section 3.1 for the basic EOQ model, development of the EPQ model involves the following assumptions:

1. The production rate is constant and strictly greater than the demand rate. (We observe that when the production rate is less than or equal to the demand rate the firm will have no need or opportunity to maintain inventory.)
2. Production is resumed only when $I(t) = 0$.

[4] When the firm manufactures the items it sells (rather than ordering stock from an independent vendor), production rate limitations would be an internal rather than an external issue.

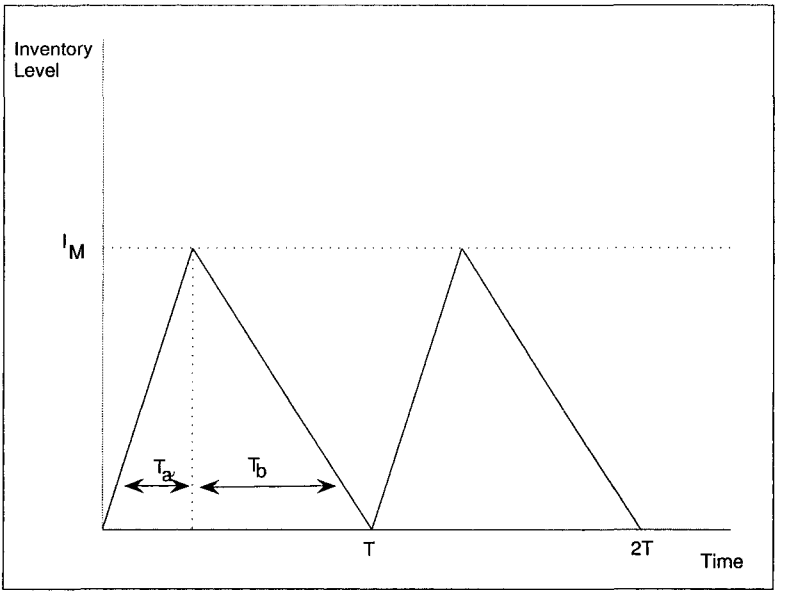


Fig. 3. Inventory Profile for the EPQ Model.

5.2 Model Formulation

Holding Costs. As before, the holding cost during one cycle of length T is "h" times the area under the $I(t)$ curve. That is the area of a triangle with base-length T and altitude I_M. Unlike the basic EOQ model, $I_M < Q$ because units are being demanded and sold while they are being produced.

The cycle time T consists of two time periods: $T_a = [0, t_1]$, the time during which the units are being both produced and sold (and inventory is increasing, at rate $M - D$); and $T_b = [t_1, T]$, the time when units are not being produced (and inventory is decreasing at rate D).

Since production of Q units takes place in time period T_a, then $Q = Mt_1$ and I_M is the level to which $I(t)$ rises during T_a, i.e.,

$$I_M = (M - D)(Q/M).$$

Thus the holding cost per cycle is

$$h \times \frac{1}{2}T[(M - D)(Q/M)].$$

Dividing by T (to obtain the average holding cost per unit time) gives the total per unit time holding cost of

$$\frac{Q}{2}(1 - D/M)h.$$

Other Costs. The ordering cost and purchase cost for this model are the same as in the basic EOQ model.

Total Costs. The total inventory cost is given by the function

$$f_{EPQ}(Q) = \frac{Q}{2}(1 - D/M)h + \frac{D}{Q}K + Dp. \tag{5}$$

5.3 Derivation of Optimal Order Quantity

If we define $\tilde{h} = (1 - D/M)h$ then the cost function Equation (5) may be rewritten as

$$f_{EPQ}(Q) = \frac{Q}{2}\tilde{h} + \frac{D}{Q}K + Dp,$$

which is Equation (1) with h replaced by $\tilde{h}$. Hence, by the previous optimization of Equation (1), the optimal order quantity for the EPQ model is given by

$$Q_{EPQ} = \sqrt{\frac{2DK}{\tilde{h}}} = \sqrt{\frac{2DK}{h(1 - \frac{D}{M})}}.$$

6 The EPQ Model with Shortages Allowed

Our final model is an EPQ (see Section 5) variant with shortages allowed. We will refer to this model as EPQ-SH. Figure 4 illustrates one possible inventory profile when, in each production "run", Q units are produced (made available) at a rate $M(> D)$, and such a "run" starts only when the number of units on backorder reaches a certain maximum level B. The cycle time in this model is the time between successive production starts. For simplicity, we assume that the start of an inventory cycle is characterized by the backorder amount at its maximum level (and, therefore, the inventory at level zero). The remaining assumptions for this variant of the EOQ model are the union of those for the EOQ-SH and the EPQ, discussed in Sections 4.1 and 5.1, respectively.

For this more complicated model, it is worthwhile to verify explicitly that its assumptions indeed produce repetitive behavior of the inventory system, cycle after cycle. This can be assured by checking that a cycle's "final state" coincides with its "initial state", i.e., that $I(T) = I(0)$ and $W(T) = W(0)$, or equivalently that $I(T) = 0$ and $W(T) = B$.

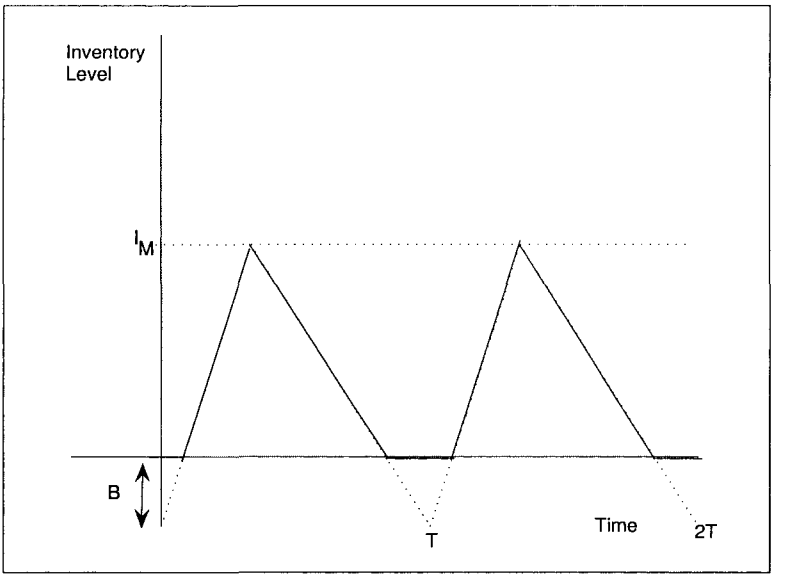


Fig. 4. An Inventory Profile for the EPQ Model with Shortages Allowed.

6.1 Model Formulation

In this model, there are two decision variables: Q, the total order quantity, and B, the number of units on backorder at the time production starts. Since T is defined as the time between successive production starts, we have $Q = DT$. In the following analysis, we examine the inventory and backlog levels between the starts of the first and second production runs. We assume that the first production run starts at $t = 0$ and the second at $t = T$.

It will prove useful to divide the time interval $[0, T]$ into as many as three subintervals: S_1, S_2, and S_3. In the first of these time intervals, new supply is arriving and is being used only to meet backorders, both from the previous cycle and arising during the current one, until none are left. Thus, shortage costs but no holding costs are accumulating. In the second subinterval, the remainder of the new supply arrives; holding costs but no shortage costs are accumulating. In the third subinterval, S_3, there is no production, and for all $t \in S_3$, $I(t) = 0$. Thus, as in S_1, shortage costs are being incurred, but no holding costs.

At the start of S_1, $W(0) = B$. These backorders are being filled at the rate M, i.e., the new supply is being used up as fast as it arrives. At the same time, new demand is coming in at a rate D and being left unfilled, creating new stockouts. So, during S_1, the waiting outages are decreasing at the net rate $M - D > 0$. Therefore, for $t \in S_1$,

$$W(t) = B - (M - D)t. \tag{6}$$

Whether or not subinterval S_2 exists, and the subsequent impact on S_3, depends upon which of the following two scenarios occurs:

- Case I: The new supply is not sufficient for the inventory level to rise above zero. (In this case, the waiting outages at the end of the production run will be nonnegative.)
- Case II: As in Figure 4, the new supply is sufficient for the backlog level to be reduced to zero. (In this case, there will be no waiting outages at the end of the production run.)

We analyze these two scenarios in turn, noting that the first of them is typically omitted in the literature.[5]

Case I: No Inventory Maintained

In this scenario, the production run ends at some time T_1 for which $Q = MT_1$. The "Case scenario" implies that for $T_1 < t < T$, all the supplies have been exhausted without the inventory level ever becoming positive, and so backlogs are increasing at the rate D. Thus, S_2 is of zero length and $S_3 = [T_1, T]$. During S_3, we have

$$W(t) = W(T_1) + D(t - T_1).$$

Lemma 3. *In Case I, the "repetitiveness conditions" hold:* $I(T) = 0$ *and* $W(T) = B$.

Proof. The inventory level is never positive, so in particular $I(T) = 0$. Next, since $W(T) = W(T_1) + D(T - T_1)$, we have

$$\begin{aligned} W(T) &= W(Q/M) + D[Q/D - Q/M] \\ &= B - (M - D)Q/M + Q - DQ/M \\ &= B. \end{aligned}$$

(Here Equation (6) has been used.) □

Holding Costs. In Case I, no holding costs arise since $I(t) = 0$ for all t.

Shortage Costs. The contribution of shortage costs to the total cost over one period is given by "s" times the area of the region under the "$W(t)$ curve." A typical curve for Case I is shown in Figure 5. We see that the region under the $W(t)$ curve can be described as consisting of two parts:

- a trapezoid over $[0, Q/M]$, whose left-hand boundary has height $W(0) = B$ and whose right-hand boundary has height $W(Q/M) = B-(M-D)Q/M$;

[5] [20] is one text that does include the full analysis.

- a trapezoid over $[Q/M, T]$, whose left-hand boundary has height $W(Q/M)$ and whose right-hand boundary has height $W(T) = B$.

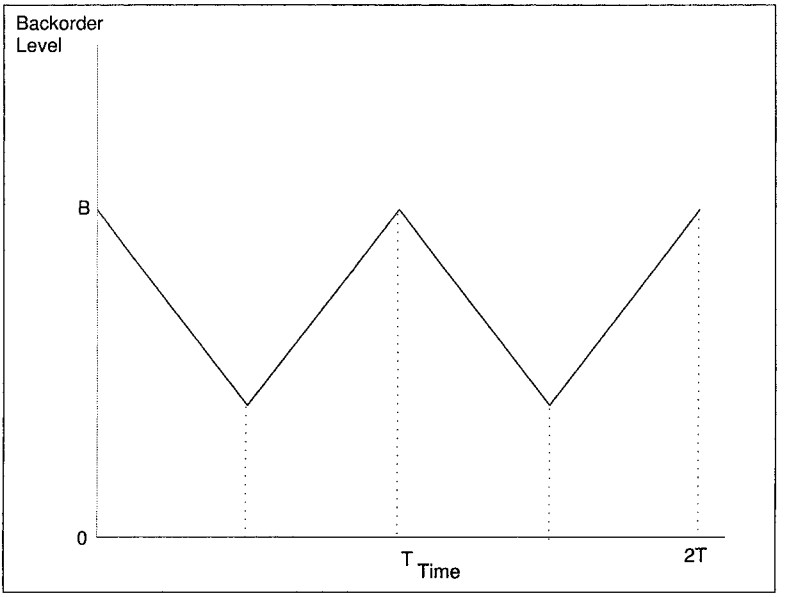


Fig. 5. Backorder Level Profile for Case I.

The area of each trapezoid may be calculated as

$$\frac{1}{2} \times \text{base length} \times (\text{left-hand height} + \text{right-hand height}).$$

For the first trapezoid, the area is $\frac{1}{2}(Q/M)[W(0) + W(Q/M)]$. For the second trapezoid, the area is $\frac{1}{2}(T - Q/M)[W(Q/M) + W(T)] = \frac{1}{2}(T - Q/M)[W(Q/M) + W(0)]$. The total area under the $W(t)$ curve is therefore (again Equation (6), with $t = Q/M$, has been used)

$$\begin{aligned}
&\frac{1}{2}\frac{Q}{M}\left[W(0) + W\left(\frac{Q}{M}\right)\right] + \frac{1}{2}\left(T - \frac{Q}{M}\right)\left[W\left(\frac{Q}{M}\right) + W(0)\right] \\
&= \frac{1}{2}T\left[W(0) + W\left(\frac{Q}{M}\right)\right] \\
&= \frac{1}{2}T\left[B + B - (M - D)\frac{Q}{M}\right]
\end{aligned}$$

$$= \frac{T}{2}Q\left[2\frac{B}{Q} - \left(1 - \frac{D}{M}\right)\right].$$

So the total shortage cost over a production cycle is

$$s\frac{T}{2}Q\left[2\frac{B}{Q} - \left(1 - \frac{D}{M}\right)\right].$$

Other Costs. The per cycle ordering cost and purchase cost for this model are the same as in the basic EOQ model.

Total Costs. The total inventory cost over a production cycle is

$$K + pQ + s\frac{T}{2}Q\left[2\frac{B}{Q} - \left(1 - \frac{D}{M}\right)\right].$$

The average cost per unit time is

$$f_1(Q,B) = \frac{KD}{Q} + pD + s\frac{Q}{2}\left[\frac{2B}{Q} - \left(1 - \frac{D}{M}\right)\right]. \tag{7}$$

If we let[6]

$$r = B/Q \quad \text{and} \quad q = D/M \tag{8}$$

then we can eliminate variable B in favor of r, and rewrite equation Equation (7) as

$$F_1(Q,r) = \frac{KD}{Q} + pD + s\frac{Q}{2}\left[2r - (1-q)\right]. \tag{9}$$

Case II: Inventory Maintained

Subinterval S_1.
In this scenario, let $S_1 = [0, T_1]$ so that T_1 is the time at which the backlog, given by Equation (6), is reduced to zero. Hence $0 = B-(M-D)T_1$ which gives $T_1 = B/(M-D)$. Algebraic manipulation and the re-use of the abbreviations in Equation (8) yields the equivalent expression

$$T_1 = \frac{rq}{1-q}T. \tag{10}$$

Only shortage costs accumulate during $S_1 = [0, T_1] = [0, B/(M-D)]$. During this time period, $W(t)$ changes at rate $M - D$ from $W(0) = B$ to $W(T_1) = 0$. The total shortage cost is given by s times the area under the $W(t)$ curve for

[6] Although independently conceived, the auxiliary variable r has some similarity to the "k" introduced above equation (10) in [22] for use in their analysis of EOQ-SH. In particular, that paper's k would be, in our notation, $B/(Q-B) = r/(1-r)$, and so is uniquely determined by r (and vice versa).

this time interval. That is the area of a right triangle with base length T_1 and height B. Thus, the shortage cost incurred during S_1 is

$$\frac{s}{2}\left(\frac{rq}{1-q}T\right)B = \left[\frac{sQTq}{2(1-q)}\right]r^2. \tag{11}$$

Subinterval S_2.
Only holding costs are incurred during S_2. It is useful to divide S_2 into two sub-subintervals, S_{21} and S_{22}. S_{21} lasts until the end of production, i.e., $S_{21} = [T_1, Q/M]$. S_{22} begins at $t = Q/M$ and ends at a time T_2 where $Q/M \leq T_2 \leq T$ and $I(T_2) = 0$.

During S_{21}, the arrival (at rate M) of the remainder of the new shipment is no longer needed to fill backlogs. That is, $W(t) = 0$ throughout this sub-subinterval. Instead production is being used to meet the new demand, which arrives at the rate D. Thus $I(t)$ is increasing from 0 at the rate $M - D$.

The total holding cost incurred during S_{21} is given by "h" times the area under the $I(t)$ curve for the interval $S_{21} = [B/(M-D), Q/M]$. This is the area of a right triangle with base-length

$$\begin{aligned}\frac{Q}{M} - \frac{B}{(M-D)} &= \frac{Q}{M}\left[1 - \frac{B}{Q(1-D/M)}\right] \\ &= \frac{Q}{M}\left[1 - \frac{r}{(1-q)}\right],\end{aligned}$$

and with altitude

$$\begin{aligned}I_M = I(Q/M) &= (M-D)\left[\frac{Q}{M} - \frac{B}{(M-D)}\right] \\ &= Q(1-q) - B \\ &= Q[(1-q) - r].\end{aligned} \tag{12}$$

For reasons that will become apparent shortly, we do NOT immediately write out an expression for the area of this triangle.

During S_{22}, there is no further production. The system begins S_{22} with $I_M = I(Q/M)$ units in inventory. During S_{22}, inventory diminishes at a rate D until it is exhausted at the end of S_{22}. For t in S_{22}, we have

$$I(t) = I_M - D(t - Q/M).$$

In particular, at the end of S_{22}, $0 = I_M - D(T_2 - Q/M)$ which implies that

$$T_2 = Q/M + (1/D)I_M.$$

Using Equation (8), Equation (12), and algebraic manipulation yields the equivalent expression

$$T_2 = Tq + T[(1-q) - r] = (1-r)T. \tag{13}$$

The total holding cost is given by "h" times the area under the $I(t)$ curve for the subinterval $S_2 = [T_1, T_2]$. That is the area of a triangle with base-length $T_2 - T_1$ and altitude I_M. The total holding cost is given by (using Equations (10), (12), and (13))

$$(h/2)(T_2 - T_1)I_M = \frac{hQT}{2}\left[(1-r) - \frac{rq}{(1-q)}\right][(1-q) - r],$$

which simplifies to

$$\frac{hQT}{2(1-q)}[(1-q) - r]^2. \tag{14}$$

Subinterval S_3.
During the third and final subinterval, $S_3 = [T_2, T]$, $W(t)$ increases from 0 at rate D, until the end, T, of the current inventory cycle; $I(t) = 0$ throughout S_3. (In particular, the repetitiveness condition $I(T) = 0$ is met) Thus, shortage costs (but no holding costs) accumulate, and for t in S_3, we have $W(t) = D(t - T_2)$. In particular, for the other repetitiveness condition (using Equation(13)),

$$\begin{aligned} W(T) &= D(T - T_2) = D(T - (1-r)T) \\ &= DrT \\ &= rQ = B. \end{aligned}$$

The total shortage cost from S_3 is "s" times the area under the $W(t)$ curve for the subinterval S_3. That is the area of a right triangle with altitude $W(T) = B$ and base-length quantifiable by Equation (13) as

$$T - T_2 = rT.$$

So the total shortage cost from S_3 is $(s/2)(Br)T$, which (by Equation (8) can be rewritten as

$$s\frac{Q}{2}r^2T. \tag{15}$$

Total Costs. Summing Equations (11), (14), and (15) gives the total holding and shortage costs during $[0, T]$ as

$$\frac{sQT}{2}\frac{r^2q}{(1-q)} + \frac{hQT}{2(1-q)}[(1-q) - r]^2 + \frac{sQT}{2}r^2$$

$$\begin{aligned} &= \frac{sQT}{2(1-q)}r^2 + \frac{hQT}{2(1-q)}[(1-q) - r]^2 \\ &= \frac{QT}{2(1-q)}\left\{sr^2 + h[(1-q) - r]^2\right\}, \end{aligned}$$

which, for fixed Q, is quadratic in r. The total cost during $[0, T]$ is obtained by adding $K + pQ$ to the last expression. Dividing the total cost during $[0, T]$

by $T = Q/D$ gives our "Case II" objective function (minimand), the average cost per unit time, as

$$F_2(Q, r) = KD/Q + pD + \frac{Q}{2(1-q)} g(r), \tag{16}$$

where

$$g(r) = sr^2 + h[(1-q) - r]^2. \tag{17}$$

6.2 Derivation of Optimal Order Quantity and Maximum Backlog Amount

For this model, as for EOQ-SH (see Section 4.3), we use sequential optimization to derive the optimal order quantity and maximum backlog amount. In what follows, Lemmas 5 and 6 give the determination of $r^*(Q)$ and Q^* for Case I; Lemmas 8 and 9 give the determination of $r^*(Q)$ and Q^* for Case II. Finally, Lemma 10 shows that Case II will always result in a lower inventory management cost.

Lemma 4. *If Case I occurs, then $B/Q + D/M \geq 1$, i.e., $r + q \geq 1$.*

Proof. If Case I occurs, then the waiting outages at the end of a production "run" will be nonnegative, i.e., $W(Q/M) \geq 0$. Since $W(t) = B - (M - D)t$ for all t in S_1, this condition reads

$$\begin{aligned} B - (M - D)Q/M \geq 0 &\Leftrightarrow B/Q \geq (M - D)/M \\ &\Leftrightarrow B/Q + D/M \geq 1. \end{aligned}$$

□

Lemma 5. *Choosing $r = r_1 = 1 - q$ (uniquely) minimizes Equation (9) for every fixed value of Q.*

Proof. For a fixed Q, Equation (9) is an increasing function of r, and so, it is minimized by choosing r (i.e., B) as low as possible while remaining in Case I, i.e., $r + q \geq 1$ (by Lemma 4). Thus, we must choose $r = 1 - q$, and note that this value is actually independent of Q. □

Lemma 6. *When $r = r_1$, the function Equation (9) is minimized (uniquely) by choosing Q to be*

$$Q_1 = \sqrt{\frac{2KD}{s(1 - D/M)}}.$$

Proof. When $r = r_1$, we have

$$F_1(Q, 1-q) = \frac{KD}{Q} + pD + s\frac{Q}{2}(1-q).$$

But this is the objective function for a basic-model EOQ problem Equation (1) with "adjusted holding cost" $h' = s(1-q)$. So by our earlier optimization of Equation (1), its minimum (over Q) occurs for

$$Q_1 = \sqrt{\frac{2KD}{h'}} = \sqrt{\frac{2KD}{s(1-q)}}. \tag{18}$$

□

Lemma 7. *If Case II occurs, then* $r + q < 1$.

Proof. If Case II occurs, then $Q/M > T_1 = B/(M-D)$. Thus

$$\begin{aligned} Q/M > B/(M-D) &\Leftrightarrow M/Q < (M-D)/B \\ &\Leftrightarrow B/Q < 1 - D/M \\ &\Leftrightarrow r < 1 - q. \end{aligned}$$

□

Lemma 8. *Choosing* r *to be*

$$r_2 = \frac{(1-q)h}{(h+s)}$$

(uniquely) minimizes Equation (16) for every fixed value of Q.

Proof. Since the coefficient of $g(r)$ given by (17) in $F(Q, r)$ is positive, for any fixed Q we want to choose r so as to minimize $g(r)$, subject to the "Case II" constraints $0 \le r < 1 - q$. We may rewrite $g(r)$ as

$$g(r) = (h+s)\left\{\left(r - \frac{h(1-q)}{h+s}\right)^2 + \Delta\right\},$$

where

$$\Delta = (1-q)^2 \frac{sh}{(h+s)^2}$$

is independent of r. Since $0 < r_2 = h(1-q)/(h+s) < 1 - q$, the constrained minimizer of the function $g(r)$ on the interval $0 \le r < 1 - q$ is indeed r_2. □

Lemma 9. *When* $r = r_2$, *the function Equation (16) is minimized (uniquely) by choosing* Q *to be*

$$Q_2 = \sqrt{\frac{2KD}{(1-q)}\frac{(h+s)}{hs}}.$$

Proof. Note that

$$(1-q) - r_2 = \frac{(1-q)s}{(s+h)},$$

a fact that simplifies the algebra for evaluating

$$g(r_2) = \frac{(1-q)^2 sh}{(h+s)},$$

which in turn gives

$$F_2(Q, r_2) = \frac{KD}{Q} + pD + \frac{Q}{2}\frac{(1-q)hs}{(h+s)}.$$

Now we need to choose $Q > 0$ to minimize this last expression. But it is the objective function for a basic-model EOQ problem Equation (1), with "holding cost" coefficient replaced by

$$h'' = \frac{(1-q)hs}{(h+s)}.$$

So by our optimization of Equation (1), its minimum occurs (uniquely) for

$$Q_2 = \sqrt{\frac{2KD}{h''}} = \sqrt{\frac{2KD}{(1-q)}\frac{(h+s)}{hs}}.$$

□

Lemma 10. *The least cost solution for the EPQ-SH is obtained by setting the order (production) quantity to* $Q = Q_2$.

Proof. By the remark ending Section 3, for Case I, we have, apart from the constant summand pD,

$$F_1(Q_1, r_1) = \sqrt{2KDh'} = \sqrt{2KDs(1-q)}.$$

The corresponding objective function value for Case II, again apart from the constant summand pD, is

$$F_2(Q_2, r_2) = \sqrt{2KDh''} = \sqrt{2KD(1-q)\frac{hs}{(h+s)}}.$$

We see that Case I is better if and only if

$$h' < h'' \Leftrightarrow s(1-q) < (1-q)hs/(h+s) \Leftrightarrow (h+s) < h,$$

which is impossible. □

In summary, for the EPQ-SH, the optimal order quantity, Q_{EPQ-SH}, is given by Q_2 and the optimal maximum backorder quantity is given by

$$\begin{aligned} B_{EPQ-SH} &= r_2 Q_2 \\ &= \left(\frac{(1-q)h}{(h+s)}\right)\sqrt{\frac{2KD}{(1-q)}\frac{(h+s)}{hs}} \\ &= \sqrt{\frac{2KDh(1-q)}{s(h+s)}}. \end{aligned}$$

7 Conclusions and Further Study

This paper has explicitly shown how the EOQ and EPQ inventory models, with and without shortages allowed, can be solved without the use of calculus. In particular, the EPQ without shortages and the EOQ and EPQ with shortages can all be solved by "reduction" to the basic EOQ model. The EOQ model is in turn treated by using the simply-proved "base case" of the arithmetic-geometric mean inequality.

It may be possible to apply the reduction approach to other deterministic inventory models. For example, Goyal ([12]; also see [13]) develops an interesting variant of the EOQ model in which the supplier permits a fixed delay in settling the amount owed.[7,8] Goyal represents the total variable cost per time unit in terms of the time, T, between successive orders (as opposed to the order quantity that we have used in the variants presented previously), making use of the following quantities[9]:

- S: the cost of placing one order;
- I_c: the interest charges per dollar investment in stocks per year;
- I_d: the interest which can be earned per dollar in a year; and
- t: the permissible delay in settling accounts.

By using the relationship $T = Q/D$, we can restate Goyal's [12] formulas for the average cost per unit time in terms of the variable Q.[10] Goyal considers two cases:

- **Case I: the cycle time is greater than or equal to the permissible delay in settling payments**

[7] Goyal [12] does not specify a solution approach.

[8] Upendra Dave [9] critiques certain aspects of the model developed in [12]. Goyal's response to the critique may be found in [13].

[9] The symbols h, D, and p have the same meanings in [12] as in the current document.

[10] Equations (1) and (2) of [12] contain a typographical error: $I_c - I_c$ instead of $I_c - I_d$.

To reduce this case to the familiar EOQ model given in Equation (1), define

$$h_{G1} = h + pI_c \quad \text{and} \quad K_{G1} = S + \frac{1}{2}Dpt^2(I_c - I_d).$$

The economic order quantity in this case is, by our optimization of Equation (1),

$$Q_{G1} = \sqrt{\frac{2DK_{G1}}{h_{G1}}} = \sqrt{\frac{2D(S + \frac{1}{2}Dpt^2(I_c - I_d))}{h + pI_c}},$$

in agreement with equation (3) of [12].

- **Case II: the cycle time is less than the permissible delay in settling payments**
  To reduce this case to the familiar EOQ model given in Equation (1), define

$$h_{G2} = h + pI_d \quad \text{and} \quad K_{G2} = S.$$

The economic order quantity in this case is, by our optimization of Equation (1),

$$Q_{G2} = \sqrt{\frac{2DK_{G2}}{h_{G2}}} = \sqrt{\frac{2DS}{h + pI_d}},$$

in agreement with equation (6) of [12].

The decision-maker's choice between the two cases is governed by the same "which curve to be on?" logic used in analyzing situations with volume discounts where the price structure has only a single breakpoint.

Chung and Huang [7] generalize [12], extending its idea to the EPQ setting and using calculus to obtain the optimal solution. Our techniques readily adapt to this case as well. A natural extension of these models that could be investigated is the application of "reduction to EOQ" to EOQ/EPQ models with both shortages and permissible delays in payment.

In the body of this paper, we have made explicit a number of assumptions customary in the development of such models. As suggested by the preceding discussion involving [12], further study should reveal elementary solution techniques under the removal of one or more of these restrictions. We outline some additional possibilities below.

- **Defective Items.** Although most traditional inventory models assume that all of the stock is of perfect quality, Huang [18] and Chou, et al. [6] study the EOQ and EPQ models with backlogging and defective items. Chou et al. note that Huang follows an algebraic approach similar to that presented in [15] and [4], whereas in [6], the goal is to simplify the calculations given in [18].
- **Time-independent shortage costs.** In their EOQ model with shortages, Lawrence and Pasternack [19] include a *"time-independent" shortage cost*, as well as the time-dependent shortage cost used in our models.

Lawrence and Pasternack describe this cost as the "administrative one" associated with recording a backorder and contacting a customer when the order arrives. This cost differs from the one captured by the datum "s" because it is a *fixed* cost per backorder, independent of how long a customer waits for the goods to arrive. It would be useful to determine whether solution via algebraic methods is still possible with the introduction of this additional term.

Acknowledgments

The second author acknowledges the influence of his late Johns Hopkins University colleague, Eliezer Naddor, on this paper's topic. Its title was prompted by the classic Western novel [2] of Max Brand (nom de plume of Frederick Schiller Faust, also the inventor of "Dr. Kildare").

References

1. R.L. Ackoff and M.C. Sasieni. *Fundamentals of Operations Research.* John Wiley & Sons, Inc., New York, 1968.
2. M. Brand. *Destry Rides Again.* Houghton-Mifflin, 1930.
3. J. Buchan and E. Koenigsberg. *Scientific Inventory Management.* Prentice-Hall, Englewood Cliffs, N.J., 1963.
4. L. E. Cárdenas-Barrón. The economic production quantity (EPQ) with shortage derived algebraically. *International Journal of Production Economics*, 70(3):289–292, 2001.
5. S.J. Chang, J.P. Chuang, and H.-J. Chen. Short comments on technical note — the EOQ and EPQ models with shortages derived without derivatives. *International Journal of Production Economics*, 92(2):242–243, 2005.
6. C.-L. Chou, Y.-F. Huang, and H.-F. Huang. A note on the deterministic inventory models with shortage and defective items derived without derivatives. *Journal of Applied Sciences*, 6(2):325–327, 2006.
7. K.-J. Chung and Y.-F. Huang. The optimal cycle time for EPQ inventory model under permissible delay in payments. *International Journal of Production Economics*, 84(3):307–318, 2003.
8. C.W. Churchman, R.L. Ackoff, and E.L. Arnoff. *Introduction to Operations Research.* John Wiley & Sons, Inc., New York, 1957.
9. U. Dave. On "Economic Order Quantity under conditions of permissible delay in payments" by Goyal. *Journal of the Operational Research Society*, 36(11):1069, 1985.
10. E. Elsayed and T.O. Boucher. *Analysis and Control of Production Systems.* Prentice-Hall, Englewood Cliffs, N.J., 1985.
11. D. Erlenkotter. Ford Whitman Harris and the Economic Order Quantity model. *Operations Research*, 38(6):937–946, 1990.
12. S. Goyal. Economic Order Quantity under conditions of permissible delay in payments. *Journal of the Operational Research Society*, 36(4):335–338, 1985.

13. S. Goyal. On "Economic Order Quantity under conditions of permissible delay in payments" by Goyal: Reply. *Journal of the Operational Research Society*, 36(11):1069–1070, 1985.
14. R. W. Grubbström. Material requirements planning and manufacturing resource planning. in *International Encyclopedia of Business and Management*, M. Warner, editor, Routledge, 3400–3420, 1996.
15. R. W. Grubbström and A. Erdem. The EOQ with backlogging derived without derivatives. *International Journal of Production Economics*, 59(1–3):529–530, 1999.
16. F. Hanssmann. *Operations Research in Production and Inventory Control.* John Wiley & Sons, Inc., New York, 1962.
17. F. W. Harris. How many parts to make at once. *Operations Research*, 38(6):947–950, 1990. Reprinted from *Factory, The Magazine of Management*, Volume 10, Number 2, 1913.
18. Y.-F. Huang. The deterministic inventory models with shortage and defective items derived without derivatives. *Journal of Statistics and Management Systems*, 6(2):171–180, 2003.
19. J.A. Lawrence Jr. and B.A. Pasternack. *Applied Management Science: Modeling, Spreadsheet Analysis, and Communication for Decision Making*, Second Edition. John Wiley & Sons, Inc., New York, 2002.
20. E. Naddor. *Inventory Systems.* Robert E. Krieger Publishing Company, Malabar, FL, 1982.
21. J.L. Riggs. *Production Systems: Planning, Analysis, and Control.* John Wiley & Sons, Inc., New York, 1970.
22. R. Ronald, G.K. Yang, and P. Chu. Technical note: The EOQ and EPQ models with shortages derived without derivatives. *International Journal of Production Economics*, 92(2):197–200, 2004.
23. E. Taft. The most economical production lot. *Iron Age*, 101:1410–1412, 1918.
24. R.J. Thierauf and R.A. Grosse. *Decision Making through Operations Research.* John Wiley & Sons, Inc., New York, 1970.
25. G. Woolsey. A requiem for the EOQ. *Production and Inventory Management*, 298(3):68–72, 1988.

The Federal Express Local Sort Facility Employee Scheduling Problem

Lawrence Bodin[1], Zhanging Zhao[2], Michael Ball[3], Atul Bhatt[4], Guruprasad Pundoor[4] and Joe Seviek[4]

[1] Robert H. Smith School of Business
University of Maryland
College Park, MD 20742
lbodin@rhmsith.umd.edu

[2] Supply Chain Transformation Group
Intel Corporation
Chandler, AZ 85226
Zhengingzhao98@gmail.com

[3] Robert H. Smith School of Business & Institute for Systems Research
University of Maryland
College Park, MD 20742
mball@rhmsith.umd.edu

[4] Federal Express Corporation
Memphis TE, 38125
akbhall@fedex.com, gpundoor@fedex.com, jhseviek@fedex.com

Summary. In this paper, we consider the problem of developing daily and weekly schedules for the employees at a Federal Express local sort facility. We define the key data components in this analysis, describe the work rules, present a mathematical programming formulation of the daily scheduling problem as an extension of a set partitioning problem, and present a simplification of the weekly scheduling problem. Although this paper describes a work in progress, this study has defined the key issues that will have to be considered when implementing the final system. This system has the potential to be used at hundreds of FedEx local sort facilities if the final testing of the final system is successful.

Key words: Personnel scheduling; integer programming; set partitioning.

1 Introduction

The problem described in this paper concerns the development of daily and weekly work schedules for the couriers, Customer Service Agents (CSA) and handlers at a sort facility for the Express Division of Federal Express Corporation (FedEx). This sort facility carries out the local delivery and pickup of packages for FedEx Express.

The operations of a FedEx local sort facility can be outlined as follows. Packages arrive at the sort facility primarily from airports in the early morning. These packages are then sorted at the sort facility and loaded onto local delivery vehicles. Couriers then deliver the packages that required delivery before 10:30 am. After these deliveries are completed, the courier delivers other packages that are scheduled for delivery later in the day and picks up packages that are to be delivered in the next couple of days at some other location. The couriers then return to the sort facility.

The packages that were picked up by the couriers are sorted, loaded onto vehicles, and most of these packages are delivered to airports for transshipping at FedEx's central sort facilities (or *hubs*). Additional checks such as international certification have to be performed on some of the packages at the local sort facility before they are delivered to the airports. In principal, therefore, one can regard the overall FedEx Express operations as a combination of deliveries and pickups at the local sort facilities and the resorting of the packages at the hubs. The effective and efficient operation at the local sort facilities is critical in ensuring that the overall FedEx Express delivery and pickup services run smoothly.

The personnel that work at a local sort facility are either regular employees or casual employees. *Regular employees* at the local sort facility sort the packages, run the routes, maintain the sort facility and act as the customer service agents. The regular employees that are going to work on a given week are known in advance and there is no guarantee that the same set of regular employees will be available two weeks in a row. Thus, the procedure described in this paper may have to be run weekly since the regular employees available on any day in two consecutive weeks may be different. Regular employees are called *employees* in the remainder of this paper.

Casual employees are additional employees that FedEx Express can employ as needed. Both casual employees and regular employees can be used to work some extra days; this creates a situation called the extended workweek for an employee (see Section 6.2). Casual employees and employees that have their workweek extended are not considered in the daily scheduling problem, since the need for casual employees and extended workweek employees is determined after the daily scheduling procedure described in this paper has been run.

Currently, FedEx has over 700 local sort facilities that carry out FedEx Express's local delivery and pickup operations. These facilities can have anywhere from under 20 to over 300 employees handling the sort facility's daily operations. Nationwide, tens of thousands of FedEx employees are involved in these local sort operations daily. Currently, schedulers using the FedEx's Compass System form the daily and weekly schedules for the employees at the sort facility without any optimization. It is envisioned that the optimization system that evolves from this study will become a module in the Compass System.

In this paper, we will describe the procedures we developed for forming the daily and weekly schedules for the employees at a FedEx Express sort facility. This project is a work in progress. We wanted to present this preliminary description of our analysis of this project in this volume for Saul Gass, because we believe that this analysis encompasses many of the modeling and analysis

principles that Saul espouses in his speeches, seminars, and papers. To get an understanding of the problem, we talked to FedEx employees at the corporate level in Payroll, Human Resources, and Industrial Engineering as well as employees at local sort facilities. Saul often noted that to properly model and analyze a real situation, the analyst must see the problem, feel the problem and discuss the problem with the persons directly involved with the problem.

After giving an overview of related work in Section 2, we give the problem's key data components in Section 3 and precisely define the problem that is being solved in Section 4. In Section 5, we present the core approach for solving the daily scheduling problem. In Sections 6 and 7, we describe various conditions that identify the key characteristics of FedEx's daily and weekly scheduling problem. We present some final remarks in Section 8.

2 Related Work

The problem addressed in this paper is a type of personnel scheduling problem. There is, of course, a long history within the Operations Research community of the study of scheduling in general and personnel scheduling in particular. One class of early models addressed problems in which there were requirements for many workers of the same type over a time horizon. For example, in each 15 minute interval there might be a requirement that a certain minimum number of workers be on duty. A simple version of this problem can be formulated as a network flow problem (see [6] for a description of this model and its application to telephone operator scheduling). More complex versions require the use of heuristics or integer programming. A popular approach involves the use of generalized set partitioning models, where columns correspond to possible work schedules and each row a time period that must be covered; here, unlike the (pure) set partitioning model, each row must be covered multiple times (e.g., [1], [4]).

There is a large body of work on crew scheduling problems for which a large set of tasks must be covered, where each task has a fixed start and end time (see [5] and [3]). These problems most typically arise when crews must carry out tasks that represent work units derived from a service with an associated timetable. Two well-studied application areas are transit crew scheduling and air crew scheduling. Set partitioning and generalized set partitioning approaches are widely used to address these problems.

The problem addressed in this paper is closest to this latter class of problems. However, in our case, some tasks have flexibility relative to the time at which they can be carried out. Nonetheless, the possible schedules are tightly constrained, in that many tasks have narrow time windows in which they can be executed, and there can be precedence constraints relative to the order in which tasks are executed.

In addition, the problem described herein can have constraints requiring that the same employee carry out certain combinations of tasks and constraints that have specified employee to task assignment restrictions. This problem probably falls most naturally within the general class of resource constrained scheduling problems [2]. When such problems are "tightly constrained," set partitioning approaches

generally are preferred. As described later in the paper, we recommend a set partitioning approach; our reasons for doing so are quite consistent with the motivation associated with the prior work described above for closely related problems. As with many such problems, the particular details of our application lead to interesting challenges, both in column generation and in incorporating side constraints into the models.

3 Key Data Components

The two key data components in this study are the Task Specifications and the Employee Qualification Conditions. These data components are now described.

3.1 Task Specification

The operations at a sort facility are broken down into a set of tasks such as the Pre-Trip Tasks, Belt Loader Tasks, Fine Sort Tasks, Route Tasks, Post-Trip Tasks and Courier Clerical Tasks. These tasks partition the operations of a sort facility into a set of definable work elements. The following can be specified about each of these tasks:

- Task duration: The task duration states the total time in minutes required to process this task and must be defined for every task.
- Precedence relationships between the tasks. These precedence relations state conditions on the tasks such as the following:
  i. The start time of one task must not exceed the start time or the end time of another task.
  ii. The end time of one task must not exceed the start time or the end time of another task.
- Specification of a definitive start time or end time of a task: For example, this specification could be Task 5 must start at exactly 7:30 am.
- Co-assignment responsibilities. The co-assignment responsibilities specify that a defined subset of tasks must be handled by the same FedEx employee. For example, the person servicing route 246 must also do the fine sort task for route 246 and the courier clerical task for route 246. These co-assignment responsibilities are very important and drive the process given below.
- Qualification conditions. The qualification conditions for a task define the knowledge that an employee must possess in order to service the task. These conditions may also define specific employee types such as a courier must service the task. The qualification conditions serve as a link between the Task Specifications and the Employee Qualification Conditions.

Finally, all tasks are defined to begin and end at the sort facility.

3.2 Employee Qualification Conditions

The employees are classified as three work types: couriers, Customer Service Agent (CSA), and handlers. The most common shift types for these employees are the following:

- Full time 5 day a week 8 hours/day employees.
- Full time 4 day a week 10 hours/day employees.
- Part Time Employees.

There are many requirements on an employee's daily and weekly work schedule in order to ensure that these schedules are feasible. Some of these requirements are described in Sections 5 and 6.

Each employee is qualified to perform certain tasks at the sort facility. Although no employee is qualified to perform every type of task at the sort facility, a few employees, called *Swing Employees* or *Swings*, are qualified to perform many of the tasks regarding the operations of the sort facility. An employee has to be qualified to carry out a task in order for the employee to be assigned to service the task. The qualification conditions serve as a link between the Task Specifications and the Employee Qualification Conditions.

Most couriers can only be assigned to service a few routes, because most couriers only know the geography for a few routes in the sort facility. Swing employees are qualified to deliver and pick up packages on many routes, as well as being qualified to carry out many of the tasks at the sort facilities. CSA and handlers stay in the sort facility and do not drive a route.

4 The Problems Being Solved

The two scheduling problems being solved are the *Daily Scheduling Problem* and the *Weekly Scheduling Problem*. The employees that are scheduled to work each day of the week are known. Further, FedEx pays each employee scheduled to work on a given day, regardless of whether the employee is assigned any tasks or not.

It is important to note that each employee is paid for the number of pay hours in the employee's daily and weekly schedule. The pay that the employee receives includes the daily and weekly overtime in the employee's schedule (whichever is greater) and guaranteed minimums on the length of a workday and length of a workweek. For example, if an employee is scheduled to work 7 hours on Monday, then the employee is paid for the 7 hours that the employee works. The Daily Scheduling Problem and Weekly Scheduling Problem are now described.

4.1 Daily Scheduling Problem

In the daily scheduling problem, each task will eventually be assigned to a qualified employee or to a causal employee (if necessary), and each employee is assigned a collection of tasks that makes up his workday. As noted above, an employee is paid for a day that he is scheduled to work, even if there are no tasks assigned to the employee on that day. The shift types and penalties for the daily work schedules for the employees are displayed in Table 1.

Table 1. Daily Pay Hour Restrictions for FedEx Express Sort Facility Employees

Employee Type	**Minimum Pay**	**Regular Pay**	**Overtime Pay (1.5* Regular Pay)**	**Overtime Pay (2* Regular Pay)**
Full Time 5 days/week	4 hours	4-8 hours	8-12 hours	$\geq$ 12 hours
Full Time 4 days/week	4 hours	4-10 hours	10-14 hours	$\geq$ 14 hours
Part Time	2 hours	2-8 hours	8-12 hours	---

The columns in Table 1 are defined as follows:

- **Employee Type** is the type of employee being considered in the row.
- **Minimum Pay** is the minimum number of hours and minutes each employee must be paid if he is assigned to work less than these hours.
- **Regular Pay** is the range of hours in a daily work schedule where the employee is paid regular time only. There is no overtime pay associated with this daily work schedule.
- **Overtime Pay (1.5* Regular Pay)** is the range of hours in a daily work schedule that the employee is paid 1.5 times the regular pay for the overtime hours that the employee works.
- **Overtime Pay (2* Regular Pay)** is the range of hours in a daily work schedule that the employee is 2* regular time. This employee is also paid 1.5 time regular time for the hours worked in the column named Overtime Pay (1.5*Regular Pay)

Thus, for a five day a week full time employee, the employee is given regular pay for a shift up to 8 hours. A five day a week full time employee whose shift duration is between 8 and 12 hours receives 1.5 * regular pay for all hours between 8 and 12 hours and regular pay for the first 8 hours. It the shift duration is over 12 hours, then the employee is paid 8 hours of regular time for the first 8 hours, 6 hours of pay (including 2 hours of overtime for hours 8-12 in the daily schedule) and 2* regular pay for any time over 12 hours in duration.

The only penalty costs in the daily scheduling problem are the daily overtime, minimum daily pay hours and slack time between tasks. The slack time between

two tasks is considered in the column generation as part of the duration of the daily work schedule. As an example, if each of two tasks on a daily work schedule is 10 minutes in duration and one task ends at 7am and the other begins at 7:05 am, then it will take 25 minutes to complete these two tasks (6:50 am to 7:15 am) and the employee is paid for 25 minutes of work.

We analyzed some daily schedules from a FedEx sort facility. These daily schedules were extracted from FedEx's Compass System. We used an approximate set of work rules, since we are still reconciling some of the work rules at a sort facility. In this case, we found that there was about 6-7% penalty cost in the daily schedules.

Based on this analysis, we decided to use a variant of a set partitioning problem for solving this problem. We wanted an approach that is guaranteed to find a feasible solution and maximize the savings. We did not want to risk using a heuristic approach that could get 'locked into' an inferior solution. This integer program is presented in Section 5.

4.2 Weekly Scheduling Problem

The weekly scheduling problem collects the daily schedules for the employees and forms the weekly schedules. The characteristics of a weekly schedule for a full time employee and a part time employee are given below.

- **Full Time Employee**
  a. Minimum duration of the weekly schedule is 35 hours. The employee is paid up to 35 hours if his weekly schedule is less than 35 hours in duration.
  b. Maximum duration of the weekly schedule is 40 hours before paying weekly overtime.
  c. Weekly overtime is paid at the rate of 1.5*Regular Pay for any week where the number of paid hours exceeds 40 hours.
- **Part Time Employee**
  a. Minimum duration of weekly schedule is 17.5 hours. The employee is paid up to 17.5 hours if his weekly schedule is less than 17.5 hours in duration.
  b. Maximum duration of weekly schedule is 40 hours before paying weekly overtime.
  c. Weekly overtime is paid at the rate of 1.5*Regular Pay for any week where the number of paid hours exceeds 40 hours.

4.3 Overtime Paid to an Employee over a Work Week

Overtime paid to an employee over a workweek is equal to the following:

Overtime = Max[weekly overtime over the workweek, total daily overtime over the workweek].

It turns out that we can simplify the computation of overtime by the critical observation given in Section 4.4.

The only two penalty costs in the weekly scheduling problem are the following: i) the minimum workweek durations - 35 hours in a week for a full time employee and 17.5 hours in a week for a part time employee, and ii) overtime attributable to any time an employee works over 40 hours in a week.

Assume the following conditions hold:

- Full time employees that work a 4 day workweek get paid overtime if their workday is longer than 10 hours.
- Full time employees that work a 5 day workweek or part time employees get paid overtime if their workday is longer than 8 hours.

Then, we can conclude that it is impossible for an employee to earn more overtime attributable to the weekly constraint of 40 hours than the total daily overtime over the week that he can earn by working more than 8 hours on the individual days in the week. This critical observation, quantified in Section 3.4, significantly simplifies the weekly scheduling problem.

4.4 A Critical Observation

Define the following terms.

- D_{ij} = maximum number of hours that employee i can work on day j in the workweek before earning daily overtime (i.e., maximum regular pay hours on day j), j = 1,2,…,6 where j = 1 is Monday, j = 2 is Tuesday, etc. For FedEx, $D_{ij} = 10$ for a 4 day a week employee and $D_{ij} = 8$ hours for all employees other than a 4 day a week employee. Employee i may be scheduled to work less than D_{ij} hours in a day and is scheduled to be paid for the hours that he is scheduled to work. As an example, if a 5 day a week employee is scheduled to work 7 hours on Monday, then this employee is paid for 7 hours of work on Monday.
- $B_{ij} = 1$ if employee i is scheduled to work on day j and $B_{ij} = 0$ otherwise.

- $D_i = \sum_{j=1}^{6} \left(B_{ij} \cdot D_{ij} \right)$ equals the total number of hours over the workweek that employee i can work from day j = 1,2,...,6. This term computes the maximum number of hours that an employee can work and not be paid weekly overtime.
- W_i = maximum number of hours over the workweek that employee i can work over the week before earning weekly overtime. For FedEx, $W_i = 40$ for all employees.

We then make the following general observation that is critical for the design of the daily and weekly scheduling algorithms:

> **For any employee *i*, if $W_i \geq D_i$, then the total overtime paid to employee *i* that can be attributed to weekly overtime is no greater than the sum of the daily overtime that the employee earns for each day of the week that the employee works.**

If $D_i > W_i$, then it is possible for the employee to earn more overtime hours attributable to the weekly constraint than the daily constraints. These observations are illustrated in the following two examples.

Example 1:

For employee *i*, let $W_i = 40$, $B_{ij} = 1$, j=1, 2, 3, 4, 5, $B_{i6} = 0$ and $D_{ij} = 8$ for j = 1, 2, 3, 4, 5, 6. Then, $D_i = \sum_{j=1}^{6} \left(B_{ij} \cdot D_{ij} \right) = 40$ is the maximum number of hours the employee can work without being paid weekly overtime. Further, assume that the employee earns daily overtime after working 8 hours in a day.

Assume that the daily and weekly schedules for the week are formed, and the employee is scheduled to work 7 hours on Monday and 9 hours on Tuesday through Friday. The employee earns 4 hours of daily overtime (1 hour of overtime on Tuesday through Friday) or 2 hours of additional overtime pay if overtime is paid at time and a half. The employee is scheduled to work 43 = 7 + 9 + 9 + 9 + 9 hours during the week and earn 3 hours of weekly overtime (or 1.5 hours of overtime pay if overtime is paid at time and a half). Thus, the employee earns 2 hours of overtime pay, since he earns 2 hours of daily overtime over the week, 1.5 hours of weekly overtime, and the overtime paid to the employee is the maximum of the daily overtime earned by the employee over the week and the weekly overtime earned by the employee.

Example 2:

If $D_i > W_i$, then the critical observation is not guaranteed to hold. Let $D_i = 40$. Assume that employee i works 7.5 hours a day from Monday through Friday. Then, his total workweek is 37.5 hours, and there is no daily overtime. However, if $W_i = 36$, then there is 1.5 hours of weekly overtime.

4.5 Algorithm for the Weekly Scheduling Problem

The implications of the result in Section 4.4 are as follows:

- We do not have to worry about minimizing weekly overtime in developing the weekly schedule for the employees.
- We only have to ensure that the employees have as little weekly minimum pay penalties as possible in developing these weekly schedules.

Based on this result, the algorithm that we suggest for the weekly scheduling problem is as follows.

- Solve the daily scheduling problem for each day of the week.
- Develop a heuristic procedure that swaps tasks between daily schedules on the same day in such a way that the total daily and weekly penalty costs are reduced.

We have tested this procedure on a data set from a small sort facility, and everything appeared to work very well. We realize that we will have to make some changes to this procedure as the development of the entire system proceeds. To accomplish this, we plan to test this procedure on additional data bases from medium and large local sort facilities.

5 Problem Formulation of the Daily Scheduling Problem

We propose using a variant of a *set partitioning problem* (VSPP) for solving the daily scheduling problem. The set partition problem is run for each day of the week as a separate optimization problem. In this formulation, the employees who are available to work each day of the week are known in advance.

As formulated in (1) to (4) below, the VSPP is always feasible. An optimal solution to the VSPP for any day of the week finds the following:

- A feasible daily work schedule for each employee scheduled to work that day of the week. It is possible that an employee is not assigned any tasks to work that day. The employee is paid regardless of whether he has tasks assigned to him or not.

- Any task left unassigned can be assigned to a relief employee or assigned to a regular employee under the notion of an extended workweek (see Section 6.1 for a description of an extended workweek for an employee).
- The total cost of the work schedules in a solution is minimal according to the objective function defined in Section 5.1.

5.1 Mathematical Programming Formulation of the VSPP

The mathematical programming formulation of the set partitioning problem representing the daily scheduling problem is the following.

Variables

- $X_{ij} = 1$ if employee *i* is assigned to work the *j*'th schedule enumerated for employee *I* in the final solution for the given day, and $X_{ij} = 0$ otherwise.
- $E_i = 1$ if employee *i* is not assigned to service at least one task on the given day, and $E_i = 0$ otherwise.
- $T_k = 1$ if task *k* is not assigned to an employee in the solution to the problem, and $T_k = 0$ if task k is assigned to an employee in the solution to the problem.

Parameters

- $a_{ikj} = 1$ if employee *i* can handle task *k* on the *j*th work schedule enumerated for employee *i*, and $a_{ikj} = 0$ otherwise.
- c_{ij} is the daily cost of employee *i* servicing the *j*th work schedule enumerated for employee *i*. This cost includes the normal hourly cost plus penalties that account for the daily schedule not satisfying the minimum duration constraint and the daily schedule having daily overtime embedded within it.
- f_i is a constant and represents the minimum daily cost for employee *i*. f_i can be larger than the minimum daily cost for an employee to ensure that as many tasks as possible are assigned to employees.
- s_k is the cost for bringing in an additional employee (relief employee or extended workweek employee) to service task *k*. As with f_i, s_k can be larger than the cost for bringing in an additional employee to ensure that as many tasks as possible are assigned to the employees scheduled to work that day.

Formulation of the VSPP

The objective function of the VSPP is the following:

$$\text{Minimize } Z = \sum_i \sum_j c_{ij} \cdot X_{ij} + \sum_i f_i \cdot E_i + \sum_k s_k \cdot T_k \qquad (1)$$

The sums in (1) are taken over all employees *i*, work schedules *j* enumerated for employee *i*, and tasks *k*.

The VSPP has the following constraints

$$\sum_i \sum_j a_{ikj} \cdot X_{ij} + T_k = 1 \text{ for all tasks } k \qquad (2)$$

The sum in (2) is taken over all employees *i* and work schedules *j* enumerated for employee *i*.

$$\sum_j X_{ij} + E_i = 1 \text{ for all employee } i \qquad (3)$$

The sum in (3) is taken over all work schedules *j* for employee *i*.

All decision variables are Boolean variables, and the range of these decision variables are over all employees i, all feasible schedules j for employee i, and all tasks k.

$$X_{ij} = \{0,1\},\ E_i = \{0,1\},\ T_k = \{0,1\} \qquad (4)$$

The equations (1) to (4) represent the following.

- (1) is the objective function of the set partitioning problem. The objective function attempts to minimize the total cost of the daily schedule. Penalty terms are added into the objective function to account for the conditions that a task is not assigned to any of the employees or there is no task assigned to an employee.
- (2) ensures that every task on the day being analyzed is assigned to an employee or is unassigned if it is to be handled by a relief employee or an employee under an extended workweek.
- (3) ensures that every employee available for work on the day being analyzed is assigned a work schedule or is paid a specified amount if he receives no tasks to carry out on the given day.
- (4) ensures that every variable in the solution is either 0 or 1. If $X_{ij} = 1$ in the final solution, then the *j*'th work schedule enumerated for employee *i* is used in the daily schedule for the day. If $E_i = 1$, then employee *i* (who is scheduled to

work on the day) receives no tasks to carry out that day. If $T_k = 1$, then task k is not assigned to any employee who is scheduled to work on the day.

5.2 Comments on the VSPP

- The VSPP always finds a feasible solution.
- The VSPP may require an excessive amount of time to generate the columns and an excessive amount of time to find the optimal solution.
- In most cases, the fixed costs in a daily scheduling problem comprise well over 90% of the total cost of a daily schedule, and the optimal feasible solution will have some extra cost in it. This approach guarantees that the best solution or a solution that is close to the best solution is found.
- Simpler heuristic solution may have great difficulty finding very good solutions.
- The objective function for each daily schedule is not linear with respect to the total duration of each daily schedule. Penalty costs and overtime costs, as well as other factors such as lunch breaks (see Section 5.1), have to be considered, and they can be difficult to handle with heuristic approaches.
- Our experimentation with a problem involving about 25 employees and 120 tasks and an approximate set of work rules (the work rules continue to be refined) formed an integer program that had about 1400 columns. This problem was solved using CPLEX, and the linear program terminated optimal. A savings of about 2% over the existing solution was realized.

6 Daily Work Rules and Generating the Columns of the VSPP

The implementation of the VSPP requires an understanding of the FedEx Express work rules. These work rules allow checks to be performed to ensure that a generated daily schedule is feasible and that the cost of the schedule is correctly determined. We have not as yet completed the design of the column generation for the daily work schedules. The generation of these daily schedules is complicated by the requirement that the final system has to operate at all of the FedEx Express local sort facilities, and the daily and weekly work rules can differ at each of these facilities. That is to say, there are *local practices and state regulations* at the sort facilities that have to be considered when carrying out the column generation. We are still trying to get clarification on some of these issues. Some of the more important constraints and considerations in this generation are described in Section 6.1.

Further, we want to generate the columns of the VSPP in such a way that no redundant columns are generated. An overview of our proposed process for carrying out this column generation is described in Section 6.2.

6.1 Issues in Generating the Daily Work Schedules

As noted previously, an employee generally gets paid for the time worked subject to the daily minimums and overtime considerations given in Table 1. Some of the major issues that have to be considered in generating daily work schedules are as follows.

- **Split Shifts.** A daily work schedule is called a split shift if there is at least a two hour of idle time between any two consecutive tasks on the daily work schedule. Split shifts have different pay rules than regular shifts, and these pay rules may be subject to local practices. One shift is made up of the tasks on the daily work schedule serviced before a break of at least two hours, and the other shift is made up of the tasks that follow this break. Generally, a split shift is comprised of two shifts. Not all employees are eligible to work split shifts, and employees who are eligible to work a split shift are specified in the data.
- **Stretch and Flex.** Each employee is given a fixed amount of time at the beginning of each shift each day for a task called *Stretch and Flex*. The Stretch and Flex is used for the employee to 'warm up' and for management to hold a quick meeting with the employees.
- **Lunch Breaks.** The accounting for lunch breaks in a daily schedule is tricky and is subject to local practices, since it appears that different local sort facilities can have their own lunch break rules. The duration of the lunch break can be a function of the duration of the work schedule being generated. The determination of the lunch break is only used for reporting purposes if the lunch break is not paid for. The duration of the lunch break is specified in each daily schedule generated, so that the employee knows the duration of the lunch break in his daily schedule.
- **Determining the Order of Servicing a Set of Tasks on a Daily Work Schedule.** A Backtrack – Forward Track procedure has been developed to determine the order of servicing a set of tasks that can possibly make up a feasible daily work schedule and the cost of this proposed daily work schedule. A side benefit of this approach is that it determines if a set of tasks can feasibly be serviced on a daily work schedule. Although this procedure is not fully defined in this paper, details of this approach are given in Section 6.2.

6.2 Overview of the Approach for Generating All Feasible Daily Work Schedules

We break the set of tasks that are specified to be carried out on a given day into four sets called Set A, Set B, Set C and Set D.

- Each element in Set A contains all co-assigned tasks that have to be serviced by the same employee and have to be serviced one after the other. Generally, these co-assigned tasks represent a route and the tasks that have to be carriied out by the same employee immediately preceding and following the route.

Each element in Set A is called a SuperTask. A SuperTask becomes a row in the VSPP. This row replaces the rows representing the co-assigned tasks making up the SuperTask. The use of Supertasks reduces the number of rows and columns of the VSPP without reducing the accuracy of the solution.

- Set B contains all co-assigned tasks in the data that are not part of the SuperTasks in Set A. These co-assigned tasks do not have to be serviced consecutively along with the appropriate SuperTask from Set A. In other words, if a daily work schedule consists of tasks from Set B and a Supertask from Set A, then it is possible to insert tasks from Set C and Set D into the daily work schedule between the Task from Set B and the Supertask from Set A.
- Set C contains all tasks that are not co-assigned to any other task. Each of these tasks has a time window that constrains the task's start time and end time.
- Set D contains all of the remaining tasks; that is to say, Set D contains all tasks that are not co-assigned to any other task, and each task in set D does not have a time window that constrains the task's start time and end time.

Furthermore, set W is defined to be the set of available employees. An **element** is defined to be a member of any of these five sets.

The algorithm begins to build a daily work schedule for an employee by selecting an element corresponding to an employee in Set W. It then selects an element from Set A if such an element exists. If this SuperTask can feasibly be serviced by the employee from set W, a (SuperTask, employee) pairing is formed. (If this pairing is not feasible, the algorithm then goes to the next element in A and repeats the process.)

Given that (SuperTask, employee) pairing is feasible, the algorithm attempts to build up the daily work schedule with elements from set B, then set C and then set D. Whenever an element is added to a daily work schedule, the cost of the daily work schedule is computed. When the daily work schedule is complete, it is added to the set of feasible daily work schedules for the day under consideration. Some workers (CSA, handlers) do not service any routes so their daily work schedules only consist of elements from sets C and D. Daily work schedules for the CSA and handlers are generated by this process.

Any time a feasible element becomes a candidate to be added to the proposed daily work schedule, the Backtrack – Forward Track procedure is used to determine if the proposed daily work schedule is feasible and, if so, what are the various costs and resulting feasible time windows associated with each task in the schedule. In this procedure, the best position of the new task in the proposed daily work schedule is determined. This procedure is still undergoing development, and we believe further analysis is necessary before the final version of this procedure is implemented.

7 Weekly Scheduling Issues

As noted previously, the weekly scheduling problem is relatively simple to analyze. However, there are a couple of issues with the weekly scheduling problem that have to be mentioned.

7.1 Extended Workweek

Earlier, we mentioned that the possibility that certain tasks may not be assigned to an employee when solving the VSPP. We noted earlier in this paper that these tasks may be handled by relief employees or employees on an extended workweek. An extended workweek for an employee occurs when an employee not scheduled to work on a given day is called in to work a shift. In this way, a 5 day a week employee may work 6 or 7 days and a 4 day a week employee may work 5-7 days during the week. Generally, the employee working the extended workweek is paid time and a half for the first extra day of work and double time for the other extra days of work.

7.2 Days Off

Employees take days off such as vacation days, have authorized leave, have to attend special meetings, training, etc. These issues have to be considered in determining the cost of a weekly schedule for the employee. We are currently investigating these issues.

8 Final Remarks

In this paper, we have attempted to document many of the critical aspects of the FedEx Express local sort facility personnel scheduling problem. We have described the problem, proposed an algorithm for solving the problem and discussed many of the issues confronting FedEx in attempting to develop a system that will be useful in all of their local sort facilities. This is a very important problem for Federal Express. The solution gives FedEx the opportunity to achieve a considerable amount of savings and get control on the important issue of scheduling the employees at their local sort facilities.

Acknowledgments

We wish to thank Amy Langston, John Gunckel, and Richard Wootten of Federal Express in Memphis TE for their support throughout this project. In addition, many Federal Express employees in Payroll, Human Resources, Industrial Engineering and at the local sort facilities have been involved in discussions on this project. We thank all of them. We also wish to thank the referee for his many

insightful comments. Partial support for this project has come from a grant to the University of Maryland from Federal Express Corporation.

References

1. M. Brusco, L. Jacobs, R. Bongiorno, D. Lyons and B. Tang. Improving Personnel Scheduling at Airlines, *Operations Research,* 43: 741 – 751, 1995.
2. J. Desrosiers, Y. Dumas, M. Solomon and F. Soumis. Time Constrained Routing and Scheduling, in *Handbook of Operations Research and Management Science: Network Models*, M. Ball, T. Magnanti, C. Monma and G. Nemhauser, eds., Elsevier, Amsterdam, 673-762, 1995.
3. M. Gamache, F. Soumis, G. Marquis and J. Desrosiers. A Column Generation Approach for Large-Scale Aircrew Rostering Problems, *Operations Research*, 47: 247-263, 1999.
4. A. Mason, D. Ryan and D. Panton. Integrated Heuristic, Simulation and Optimization Approaches to Staff Scheduling, *Operations Research*, 46: 161-175, 1998.
5. D. Ryan. The Solution of Massive Generalized Set Partitioning Problems in Crew Rostering, *J. Operational Research Soc*, 43: 459-467, 1992.
6. M. Segal. The Operator Scheduling Problem: A Network Flow Approach, *Operations Research*, 22: 808-823, 1974.

Sensitivity Analysis in Monte Carlo Simulation of Stochastic Activity Networks

Michael C. Fu

Robert H. Smith School of Business & Institute for Systems Research
Department of Decision and Information Technologies
University of Maryland
College Park, MD 20742
mfu@rhsmith.umd.edu

Summary. Stochastic activity networks (SANs) such as those arising in Project Evaluation Review Technique (PERT) and Critical Path Method (CPM) are an important classical set of models in operations research. We focus on sensitivity analysis for stochastic activity networks when Monte Carlo simulation is employed. After a brief aside reminiscing on Saul's influence on the author's career and on the simulation community, we review previous research for sensitivity analysis of simulated SANs, give a brief summary overview of the main approaches in stochastic gradient estimation, derive estimators using techniques not previously applied to this setting, and address some new problems not previously considered. We conclude with some thoughts on future research directions.

Key words: Stochastic activity networks; PERT; CPM; project management; Monte Carlo simulation; sensitivity analysis; derivative estimation; perturbation analysis; likelihood ratio/score function method; weak derivatives.

1 Introduction

In the vast toolkit of operations research (OR), two of the most useful methods/models without a doubt are simulation and networks. Numerous surveys of practitioners consistently place these in the top 10 in terms of applicability to real-world problems and solutions. Furthermore, there is a large industry of supporting software for both domains. Thus, almost all degree-granting programs in operations research offer standalone courses covering these two topics. At the University of Maryland's Robert H. Smith School of Business, these two courses at the Ph.D. level have course numbers BMGT831 and BMGT835. They form half of the core of the methodological base that all OR Ph.D. students take, along with BMGT834 Stochastic Models and BMGT830 Linear programming, which Saul Gass taught for much of his career at Maryland,

using his own textbook first published in 1958, and translated into numerous other languages, including Polish, Russian, and Spanish. Reflecting back on my own academic career, the first presentation I ever gave at a technical conference was at the 1988 ORSA/TIMS Spring National Meeting in Washington, D.C., in which Saul was the General Chair. When the INFORMS Spring Meeting returned to Washington, D.C. in 1996, Saul was again on the advisory board and delivered the plenary address, and this time, through my connection with him, I served on the Program Committee, in charge of contributed papers.

Saul's contributions to linear programming and his involvement in ORSA and then INFORMS are of course well known, but what may not be as well known are his contributions to the simulation community and simulation research. Saul was heavily involved with the Winter Simulation Conference — the premier annual meeting of stochastic discrete-event simulation researchers, practitioners, and software vendors — during what could be called the formative and critical years of the conference. Saul served as the ORSA representative on the Board of Directors in the early 1980s. During these years, "he contributed much insight into the operation of conferences and added prestige" [21], which clearly helped launch these meetings on the path to success. In addition, Saul has also contributed to simulation research in model evaluation and validation through "(i) the development of a general methodology for model evaluation and validation [15, 18], and (ii) the development of specific validation techniques that used quantitative approaches [16, 17]" [21]. Arjang Assad's article in this volume details numerous additional instances of Saul's involvement with simulation during his early career.

The networks studied in this paper are a well-known class of models in operations research called *stochastic activity networks* (SANs), which include the popular Project Evaluation Review Technique (PERT) and Critical Path Method (CPM). A nice introduction to PERT can be found in the encyclopedia entry written by Arjang Assad and Bruce Golden [2], two long-time colleagues of Saul hired by him in the 1970s while he was chairman of the Management Science & Statistics department at the University of Maryland, and co-authors of other chapters in this volume. Such models are commonly used in resource-constrained project scheduling; see [6] for a review and classification of existing literature and models in this field. In SANs, the most important measures of performance are the completion time and the arc (activity) criticalities. Often just as important as the performance measures themselves are the sensitivities of the performance measures with respect to parameters of the network. These issues are described in the state-of-the-art review [8]. When the networks become complex enough, Monte Carlo simulation is frequently used to estimate performance. However, simulation can be very expensive and time consuming, so finding ways to reduce the computational burden are important. Efficient estimation schemes for estimating arc criticalities and sensitivity curves via simulation are derived in [4, 5].

In this paper, we consider the problem of estimating performance measure sensitivities in the Monte Carlo simulation setting. This problem has been studied by various researchers. In particular, infinitesimal perturbation analysis (IPA) and the likelihood ratio (LR) method have been used to derive efficient estimators for local sensitivities — see [3] for IPA used with conditional Monte Carlo, and [1] for the LR method used with control variates; whereas [7] use design of experiments methodology and regression to fit performance curves. The use of conditioning in [3] for variance reduction differs from its use in deriving estimators for performance measures to which IPA cannot be applied. Here, we present some alternative estimators: weak derivative (WD) estimators, and smoothed perturbation analysis (SPA) estimators. These can serve as alternatives to IPA and LR estimators. We also consider estimation of performance measure sensitivities that have not been considered in the literature. The rest of the paper is organized as follows. In Section 2, we present the problem setting and briefly review the PA, LR, and WD approaches to derivative estimation. The derivative estimators are presented in Section 3. Some concluding remarks, including extensions and further research, are made in Section 4.

2 Problem Setting

We consider a directed acyclic graph, defined by a set of nodes $\mathcal{N}$ of integers $1, \ldots, |\mathcal{N}|$, and a set of directed arcs $\mathcal{A} \subset \{(i,j) : i,j \in \mathcal{N}; i < |\mathcal{N}|, j > 1\}$, where (i,j) represents an arc from node i to node j, and, without loss of generality, we take node 1 as the source and node $|\mathcal{N}|$ as the sink (destination). For our purposes, we also map the set of directed arcs to the integers $\{1, \ldots, |\mathcal{A}|\}$ by the lexicographically ordering on elements of $\mathcal{A}$. Both representations of the arcs will be used interchangeably, whichever is the most convenient in the context. Let $\mathcal{P}$ denote the set of paths from source to sink. The input random variables are the individual activity times given by X_i, with cumulative distribution function (c.d.f.) $F_i, i = 1, \ldots, |\mathcal{A}|$, and corresponding probability density function (p.d.f.) or probability mass function (p.m.f.) f_i. Assume all of the activity times are *independent*. However, it should be clear that duration of paths in $\mathcal{P}$ will not in general be independent, such as in the following example, where all three of the path durations are dependent, since X_6 must be included in any path.

Example: 5-node network with $\mathcal{A} = \{(1,2),(1,3),(2,3),(2,4),(3,4),(4,5)\}$ mapped as shown in Figure 1; $\mathcal{P} = \{(1,4,6),(1,3,5,6),(2,5,6)\}$.

Let $P^* \in \mathcal{P}$ denote the set of activities on the *optimal* (critical) path corresponding to the project duration (e.g., shortest or longest path, depending on the problem), where P^* itself is a random variable. In this paper, we consider the total project duration, which can be written as

$$Y = \sum_{j \in P^*} X_j. \tag{1}$$

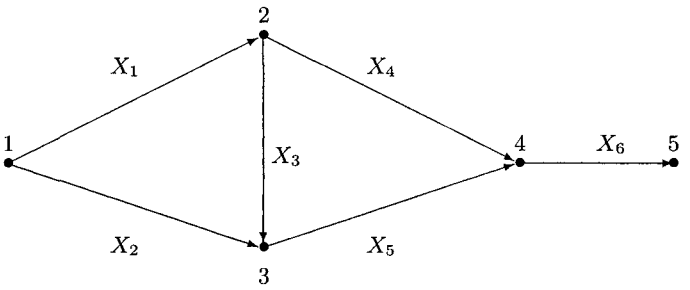


Fig. 1. Stochastic Activity Network.

Another important performance measure is arc criticality, which is the probability that a given activity is on the optimal (or critical) path, i.e., $P(i \in P^*)$ for activity i. As the focus of this paper is sensitivity analysis, we wish to estimate derivatives of performance measures involving Y with respect to a parameter θ. We consider three cases:

1. $dE[Y]/d\theta$, where θ appears in the activity time distributions (i.e., in some p.d.f. f_i);
2. $dP(Y > y)/d\theta$ for some given $y \geq 0$, where again θ appears in the activity time distributions;
3. $dP(Y > \theta)/d\theta$, where θ occurs directly in the tail distribution performance measure (so this is essentially the negative of the density function evaluated at the point θ).

The first case has been addressed previously in [3] using IPA and in [1] using the LR method. We will review these methods briefly, and also present new estimators based on the use of weak derivatives (WD). Neither the second nor the third case has been considered in the literature, but both the LR method and WD approaches can be extended to the second case in a straightforward manner, whereas the IPA estimator would fail for that form of performance measure, requiring the use of smoothed perturbation analysis (SPA) to obtain an unbiased estimator. The third case essentially provides an estimate of the density of Y if taken over all possible values of θ. This form of performance measure presents some additional challenges not seen in previous work.

Before deriving the estimators, we give a brief overview of IPA, SPA, the LR method, and the WD approach. Details can be found in [12]. For illustrative purposes, we will just consider the first case above, where the performance measure is an expectation:

$$J(\theta) = E[Y(\theta)] = E[Y(X_1, \ldots, X_T)]. \tag{2}$$

Y is the (univariate) output performance measure, $\{X_i\}$ are the input random variables, and T is the number of input random variables. In the SAN setting, $T = |\mathcal{A}|$, and Y is given by (1). Stochastic simulation can be viewed as a means of carrying out the so-called "law of the unconscious statistician" (cf. p.7 in [20]; this term was removed in the 1996 second edition):

$$E[Y(\mathbf{X})] = \int y dF_Y(y) = \int Y(\mathbf{x}) dF_{\mathbf{X}}(\mathbf{x}). \tag{3}$$

Coming into the simulation are input random variables $\{X_i\}$, for which we know the distribution $F_{\mathbf{X}}$; coming out of the simulation is an output random variable Y, for which we would like to know the distribution F_Y; and what we have is a way to generate samples of the output random variables as a function of the input random variables via the simulation model. If we knew the distribution of Y, there would generally be no need for simulation.

For simplicity, we assume for the remainder of the discussion in this section that the parameter θ is scalar, because the vector case can be handled by taking each component individually. In the right-hand side of (3), the parameter appearing directly in the sample performance $Y(\cdot;\theta)$ corresponds to the view of perturbation analysis (PA), whereas its appearance in the distribution $F_{\mathbf{X}}(\cdot;\theta)$ leads to the likelihood ratio (LR) method or weak derivative (WD) approach.

Let f denote the joint density of all of the input random variables. Differentiating (3), and assuming interchangeability of integration and differentiation:

$$\frac{dE[Y(X)]}{d\theta} = \begin{cases} \displaystyle\int_{-\infty}^{\infty} Y(x) \frac{df(x;\theta)}{d\theta} dx, & (4) \\ \displaystyle\int_0^1 \frac{dY(X(\theta;u))}{d\theta} du, & (5) \end{cases}$$

where x and u and the integrals are all T-dimensional. For notational simplicity, these T-dimensional multiple integrals are written as a single integral throughout, and we also assume that one random number u produces one random variate x (e.g., using the inverse transform method). In (4), the parameter appears in the distribution directly, whereas in (5), the underlying uncertainty is considered the uniform random numbers.

For expositional ease, we begin by assuming that the parameter only appears in X_1, which is generated independently of the other input random variables. So for the case of (5), we use the chain rule to write

$$\begin{aligned} \frac{dE[Y(X)]}{d\theta} &= \int_0^1 \frac{dY(X_1(\theta;u_1), X_2, \ldots)}{d\theta} du \\ &= \int_0^1 \frac{\partial Y}{\partial X_1} \frac{dX_1(\theta;u_1)}{d\theta} du. \end{aligned} \tag{6}$$

In other words, the estimator takes the form

$$\frac{\partial Y(X)}{\partial X_1}\frac{dX_1}{d\theta}, \tag{7}$$

where the parameter appears in the transformation from random number to random variate, and the derivative is expressed as the product of a sample path derivative and derivative of a random variable. This approach is called infinitesimal perturbation analysis (IPA), and the main requirement for it to yield an unbiased estimator is that the sample performance be almost surely continuous, which is not satisfied for certain forms of performance measure (e.g., probabilities) and will be violated for some stochastic systems.

Assume that X_1 has marginal p.d.f. $f_1(\cdot;\theta)$, and that the joint density for the remaining input random variables $(X_2,\ldots)$ is given by f_{-1}, which has no (functional) dependence on θ. Then the assumed independence gives $f = f_1 f_{-1}$, and the expression (4) involving differentiation of a density (measure) can be further manipulated using the product rule of differentiation to yield the following:

$$\frac{dE[Y(X)]}{d\theta} = \int_{-\infty}^{\infty} Y(x)\frac{\partial f_1(x_1;\theta)}{\partial\theta} f_{-1}(x_2,\ldots)dx \tag{8}$$

$$= \int_{-\infty}^{\infty} Y(x)\frac{\partial \ln f_1(x_1;\theta)}{\partial\theta} f(x)dx. \tag{9}$$

In other words, the estimator takes the form

$$Y(X)\frac{\partial \ln f_1(X_1;\theta)}{\partial\theta}. \tag{10}$$

On the other hand, if instead of expressing the right-hand side of (8) as (9), the density derivative is written as

$$\frac{\partial f_1(x_1;\theta)}{\partial\theta} = c(\theta)\left(f_1^{(2)}(x_1;\theta) - f_1^{(1)}(x_1;\theta)\right),$$

it leads to the following relationship:

$$\frac{dE[Y(X)]}{d\theta} = \int_{-\infty}^{\infty} Y(x)\frac{\partial f(x;\theta)}{\partial\theta}dx,$$

$$= c(\theta)\left(\int_{-\infty}^{\infty} Y(x) f_1^{(2)}(x_1;\theta) f_{-1}(x_2,\ldots)dx\right.$$

$$\left. - \int_{-\infty}^{\infty} Y(x) f_1^{(1)}(x_1;\theta) f_{-1}(x_2,\ldots)dx\right).$$

The triple $\left(c(\theta), f_1^{(1)}, f_1^{(2)}\right)$ constitutes a weak derivative (WD) for f_1, which is not unique if it exists. The corresponding WD estimator is of the form

$$c(\theta)\left(Y(X_1^{(2)}, X_2, \ldots) - Y(X_1^{(1)}, X_2, \ldots)\right), \tag{11}$$

where $X_1^{(1)} \sim f_1^{(1)}$ and $X_1^{(2)} \sim f_1^{(2)}$. In other words, the estimator takes the difference of the sample performance at two different values of the input random variable X_1. The term "weak" derivative comes about from the possibility that $\frac{df_1(\cdot;\theta)}{d\theta}$ in (8) may not be proper, and yet its *integral* may be well-defined, e.g., it might involve delta-functions (impulses), corresponding to mass functions of discrete distributions.

If in the expression (5) the interchange of expectation and differentiation does not hold (e.g., if Y is an indicator function), then as long as there is more than one input random variable, appropriate conditioning will allow the interchange as follows:

$$\frac{dE[Y(X)]}{d\theta} = \int_0^1 \frac{dE[Y(X(\theta;u))|Z(u)]}{d\theta} du, \tag{12}$$

where $Z \subset \{X_1, \ldots, X_T\}$. This conditioning is known as smoothed perturbation analysis (SPA), because it is intended to "smooth" out a discontinuous function. It leads to an estimator of the following form:

$$\frac{\partial E[Y(X)|Z]}{\partial X_1} \frac{dX_1}{d\theta}. \tag{13}$$

Note that taking Z as the entire set leads back to (7), the IPA estimator. The chief difficulty in applying the methodology is determining the appropriate Z such that $E[Y(X)|Z]$ is both smooth, and its derivative can be easily estimated. Further details can be found in [14, 11].

3 Derivations of the Estimators

We begin with $dE[Y]/d\theta$, where θ is some parameter occurring in the distribution of X_1 only, as considered at the end of the last section. Then, the IPA estimator can be obtained by straightforward differentiation of the expression for Y given by (1), noting that θ only affects Y through X_1:

$$\frac{dY}{d\theta} = \frac{dX_1}{d\theta} \mathbf{1}\{1 \in P^*\},$$

where $\mathbf{1}\{\cdot\}$ denotes the indicator function. The LR/SF estimator is given by (10), and the WD estimator is given by (11).

If we allow the parameter to possibly appear in all of the distributions, then the IPA estimator is found by applying the chain rule:

$$\frac{dY}{d\theta} = \sum_{i \in P^*} \frac{dX_i}{d\theta},$$

whereas the LR/SF and WD estimators are derived by applying the product rule of differentiation to the underlying input distribution, applying the independence that allows the joint distribution to be expressed as a product of marginals. In particular, the LR/SF estimator is given by

$$Y(X)\left(\sum_{i=1}^{T}\frac{\partial \ln f_i(X_i;\theta)}{\partial\theta}\right),$$

where f_i is the p.d.f. for X_i. The IPA and LR estimators differ from the respective ones in [3] and [1], in that those both use variance reduction techniques to improve the estimators further.

The WD estimator is of the form

$$\sum_{i=1}^{T} c_i(\theta)\left(Y(X_1,\ldots,X_i^{(2)},\ldots,X_T)-Y(X_1,\ldots,X_i^{(1)},\ldots,X_T)\right),$$

where $X_i^{(j)} \sim f_i^{(j)}, j = 1,2; i = 1,\ldots,T$, and $(c_i(\theta), f_i^{(2)}, f_i^{(1)})$ is a weak derivative for f_i.

Examples: Consider two common distributions: the exponential with mean θ and the normal (Gaussian) with mean θ and standard deviation σ. In both cases, let θ_i be the corresponding parameter in X_i. Then we have the following estimators:

(a) exponential distribution, X_i with mean θ_i

$$\text{IPA:} \quad \tfrac{X_i}{\theta_i}\mathbf{1}\{i\in\mathcal{P}^*\},$$

$$\text{LR:} \quad Y(X)\tfrac{1}{\theta_i}\left(\tfrac{X_i}{\theta_i}-1\right),$$

$$\text{WD:} \quad \tfrac{1}{\theta_i}\left(Y(X_1,\ldots,X_i^{(2)},\ldots)-Y(X)\right),$$

where $X_i^{(2)}$ has the following Erlang distribution (p.d.f.):

$$\frac{1}{\theta_i^2}xe^{-x/\theta_i}\mathbf{1}\{x>0\}.$$

(b) normal distribution, X_i with mean θ_i and standard deviation σ_i

$$\text{IPA:} \quad \mathbf{1}\{i\in\mathcal{P}^*\},$$

$$\text{LR:} \quad Y(X)\tfrac{X_i-\theta_i}{\sigma_i^2},$$

$$\text{WD:} \quad \tfrac{1}{\sqrt{2\pi}\sigma_i}\left(Y(X_1,\ldots,X_i^{(2)},\ldots)-Y(X_1,\ldots,X_i^{(1)},\ldots)\right),$$

with $X_i^{(1)} = \theta_i - X'$ and $X_i^{(2)} = \theta_i + X'$, where X' has the following Weibull distribution (p.d.f.):

$$8\sigma_i^4 xe^{-(2\sigma_i^2 x)^2}\mathbf{1}\{x>0\}.$$

If instead of an expectation, we were interested in estimating the tail distribution, e.g., $P(Y > y)$ for some fixed y, the WD and LR/SF estimators would simply replace Y with the indicator function $\mathbf{1}\{Y > y\}$. However, IPA does not apply, since the indicator function is inherently discontinuous, so an extension of IPA such as SPA is required. On the other hand, if the performance measure were $P(Y > \theta)$, then since the parameter does not appear in the distribution of the input random variables, WD and LR/SF estimators cannot be derived without first carrying out an appropriate change of variables, which we will shortly demonstrate.

To derive an estimator via conditioning for the derivative of $P(Y > y)$, we first define the following:

$$\begin{aligned}
\mathcal{P}_j &= \{P \in \mathcal{P} \mid j \in P\} \;=\; \text{set of paths containing arc } j,\\
|P| &= \text{length of path } P,\\
|P|_{-j} &= \text{length of path } P \text{ with } X_j = 0.
\end{aligned}$$

The idea will be to condition on all activity times except a set that includes activity times dependent on the parameter.

In order to proceed, we need to specify the form of Y. We will take it to be the longest path. Other forms, such as shortest path, are handled analogously. Again, assuming that θ occurs in the density of X_1 and taking the conditioning quantities to be everything except X_1, i.e., $Z = \{X_2, \ldots, X_T\}$, we have

$$\begin{aligned}
L_Z(\theta) = P_Z(Y > y) &\equiv E[\mathbf{1}\{Y > y\}|X_2, \ldots, X_T]\\
&= \begin{cases} 1 & \text{if } \max_{P\in\mathcal{P}} |P|_{-1} > y;\\ P_Z(\max_{P\in\mathcal{P}_1} |P| > y) & \text{otherwise;} \end{cases}
\end{aligned}$$

where P_Z denotes the conditional (on Z) probability. Since

$$\begin{aligned}
P_Z(\max_{P\in\mathcal{P}_1} |P| > y) &= P_Z(X_1 + \max_{P\in\mathcal{P}_1} |P|_{-1} > y)\\
&= P_Z(X_1 > y - \max_{P\in\mathcal{P}_1} |P|_{-1}) = \bar{F}_1(y - \max_{P\in\mathcal{P}_1} |P|_{-1}; \theta),
\end{aligned}$$

where $\bar{F} \equiv 1 - F$ denotes the complementary c.d.f., we have

$$L_Z(\theta) = \bar{F}_1(y - \max_{P\in\mathcal{P}_1} |P|_{-1}; \theta) \cdot \mathbf{1}\{\max_{P\in\mathcal{P}} |P|_{-1} \le y\} + \mathbf{1}\{\max_{P\in\mathcal{P}} |P|_{-1} > y\}.$$

Differentiating, we get the estimator:

$$\frac{dL_Z}{d\theta} = \frac{\partial \bar{F}_1(y - \max_{P\in\mathcal{P}_1} |P|_{-1}; \theta)}{\partial\theta} \cdot \mathbf{1}\{\max_{P\in\mathcal{P}} |P|_{-1} \le y\}, \tag{14}$$

which applies for both continuous and discrete distributions, as the following example illustrates.

Example: For the 5-node example, $\mathcal{P} = \{146, 1356, 256\}$, $\mathcal{P}_1 = \{146, 1356\}$, $|146|_{-1} = X_4 + X_6, |1356|_{-1} = X_3 + X_5 + X_6, |256|_{-1} = X_2 + X_5 + X_6$. If X_1 is exponentially distributed with mean θ, $\partial \bar{F}_1(x;\theta)/\partial\theta = e^{-x/\theta}(x/\theta^2)$, and the estimator is given by

$$\exp\left(-\frac{y - \max(X_3 + X_5, X_4) - X_6}{\theta}\right) \frac{y - \max(X_3 + X_5, X_4) - X_6}{\theta^2}$$
$$\cdot \mathbf{1}\{\max(X_3 + X_5, X_4, X_2 + X_5) + X_6 \leq y\},$$

whereas if X_1 is Bernoulli — i.e., is equal to x_{low} with probability θ and equal to $x_{high} > x_{low}$ otherwise — then $\partial \bar{F}_1(x;\theta)/\partial\theta = \mathbf{1}\{x_{low} \leq x < x_{high}\}$, and the estimator is given by

$$\mathbf{1}\{x_{low} \leq y - \max(X_3 + X_5, X_4) - X_6 < x_{high}\}$$
$$\cdot \mathbf{1}\{\max(X_3 + X_5, X_4, X_2 + X_5) + X_6 \leq y\}.$$

Clearly, the estimator (14) was derived without loss of generality, so for the parameter θ being in the distribution of X_i, we have the estimator

$$\frac{\partial \bar{F}_i(y - \max_{P \in \mathcal{P}_i} |P|_{-j}; \theta)}{\partial\theta} \cdot \mathbf{1}\{\max_{P \in \mathcal{P}} |P|_{-j} \leq y\}. \tag{15}$$

For the shortest path problem, simply replace "max" with "min" throughout in the estimator (15).

To derive an LR/SF or WD derivative estimator for the performance measure $P(Y > \theta)$, there are two main ways to do the change of variables: subtraction or division, i.e.,

$$P(Y - \theta > 0), \quad P(Y/\theta > 1).$$

Note that this requires translating the operation on the output performance measure back to a *change of variables* on the input random variables, so this clearly requires some additional knowledge of the system under consideration. In this particular case, it turns out that two properties make it amenable to a change of variables: (i) additive performance measure; (ii) clear characterization of paths that need to be considered. The simplest change of variables is to take

$$\tilde{X}_i = X_i/\theta \ \forall i \in \mathcal{A},$$

so that θ now appears as a scale parameter in *each* distribution f_i. If $\tilde{Y}$ represents the performance measure after the change of variables, then we have

$$P(Y > \theta) = P(\tilde{Y} > 1),$$

and this can be handled as previously discussed.

Another change of variables that will work in this case is to subtract the quantity θ from an appropriate set of arc lengths. In particular, the easiest sets are the nodes leading out of the source or the nodes leading into the sink:

$$\tilde{X}_i = X_i - \theta,$$

for arcs i corresponding to directed arcs $(1, j) \in \mathcal{A}$ or $(j, |\mathcal{N}|) \in \mathcal{A}$. In the 5-node example of Figure 1, this would be either {1,2} or {6}. Note that minimal cut sets will not necessarily do the trick. For example, in the 5-node example, {1,5} is a cut set, but both members are contained on the path (1,3,5,6), so subtracting θ from these two arc lengths would lead to possibly erroneous results.

Again, if $\tilde{Y}$ represents the performance measure after the change of variables, then we have

$$P(Y > \theta) = P(\tilde{Y} > 0).$$

Now the parameter θ appears in the distribution, specifically as a location parameter, but only in a relatively small *subset* of the $\{f_i\}$. Since this transformation results in the parameter appearing in fewer number of input random variables, it may be preferred, because for both the LR/SF and WD estimators, the amount of effort is proportional to the number of times the parameter appears in the distributions. The extra work for a large network can be particularly burdensome for the WD estimator. However, for the LR estimator, this type of location parameter is problematic, since it changes the support of the input random variable, making it inapplicable.

Lastly, we apply PA to the problem of estimating $dP(Y > \theta)/d\theta$. We note that this estimation problem only makes sense in the *continuous* distribution case, where it is essentially an estimation of (the negative of) the p.d.f., since in the discrete case, the corresponding derivative is 0; thus, assume henceforth that each X_i has a p.d.f. f_i. Again, this type of performance measure cannot be handled by IPA, so we use SPA. The idea is to condition on a special set of activity times such that both the set itself and its complement have a non-zero probability of having a corresponding activity on the critical path.

Recall the following network definitions. A *cut set* is a set of arcs such that their removal from the network leaves no path from source to sink. A *minimal cut set* is a cut set such that the removal of any arc in the set no longer leaves a cut set. In the 5-node example, the minimal cut sets are {1,2}, {1,5}, {2,3,4}, {4,5}, {6}. By definition, a minimal cut set will have an activity on the critical path. The following observation is key:

Lemma. Let $\mathcal{C}$ be a minimal cut set for the network, and let $Z = \{X_i : i \notin \mathcal{C}\}$. If there exists an $i \notin \mathcal{C}$ such that $P(i \in P^*) > 0$, then $P_Z(Y > \theta)$ is a.s. continuous with respect to θ.

Thus, if one can find a minimal cut set that satisfies the condition in the lemma, one can in principle derive an unbiased derivative estimator for $P(Y > \theta)$ by conditioning on the complement set of activity times and then taking

the sample path derivative. Note, however, that finding such a minimal cut set may be a computationally formidable task for large networks. Furthermore, as we shall see, in order to take the sample path derivative in a convenient form, it is desirable that the activities in the cut set actually partition the path space. We illustrate these ideas in an extended example using the 5-node network.

Example: For the 5-node example, we consider all of the minimal cut sets. (i) Using minimal cut set $\mathcal{C} = \{6\}$, we take $Z = \{X_1, X_2, X_3, X_4, X_5\}$, so we have

$$P_Z(Y > \theta) = P_Z(\max_{P \in \mathcal{P}_6} X_6 + |P|_{-6} > \theta) = \bar{F}_6(\theta - \max_{P \in \mathcal{P}_6} |P|_{-6}).$$

Differentiating, the final estimator is given by

$$\frac{dP_Z(Y > \theta)}{d\theta} = -f_6(\theta - \max_{P \in \mathcal{P}_6} |P|_{-6}).$$

Note that the form of the estimator only involves the p.d.f.s of those arcs in the cut set. If X_6 follows an exponential distribution, the specific estimator is given by

$$-\frac{1}{E[X_6]} \exp\left(\frac{\max(X_1 + X_3 + X_5, X_1 + X_4, X_2 + X_5) - \theta}{E[X_6]}\right)$$
$$\cdot \mathbf{1}\{\max(X_1 + X_3 + X_5, X_1 + X_4, X_2 + X_5) \leq \theta\}.$$

(ii) Using minimal cut set $\mathcal{C} = \{1, 2\}$, we take $Z = \{X_3, X_4, X_5, X_6\}$, so we have

$$\begin{aligned} P_Z(Y > \theta) &= 1 - P_Z(\max_{P \in \mathcal{P}} |P| \leq \theta) \;=\; 1 - P_Z(\max_{P \in \mathcal{P}_1 \cup \mathcal{P}_2} |P| \leq \theta) \\ &= 1 - P_Z(\max_{P \in \mathcal{P}_1} |P| \leq \theta, \max_{P \in \mathcal{P}_2} |P| \leq \theta) \\ &= 1 - P_Z(X_1 + \max_{P \in \mathcal{P}_1} |P|_{-1} \leq \theta, X_2 + \max_{P \in \mathcal{P}_2} |P|_{-2} \leq \theta) \\ &= 1 - F_1(\theta - \max_{P \in \mathcal{P}_1} |P|_{-1}) F_2(\theta - \max_{P \in \mathcal{P}_2} |P|_{-2}), \end{aligned}$$

where we have used the fact that $\mathcal{P}_1$ and $\mathcal{P}_2$ partition the path space, and X_1 and X_2 are independent. Differentiating, the final estimator is given by

$$\begin{aligned} \frac{dP_Z(Y > \theta)}{d\theta} = &-f_1(\theta - \max_{P \in \mathcal{P}_1} |P|_{-1}) F_2(\theta - \max_{P \in \mathcal{P}_2} |P|_{-2}) \\ &- F_1(\theta - \max_{P \in \mathcal{P}_1} |P|_{-1}) f_2(\theta - \max_{P \in \mathcal{P}_2} |P|_{-2}). \end{aligned}$$

Using minimal cut set $\mathcal{C} = \{4, 5\}$ will yield the analogous estimator

$$\frac{dP_Z(Y>\theta)}{d\theta} = -f_4(\theta - \max_{P\in\mathcal{P}_4}|P|_{-4})F_5(\theta - \max_{P\in\mathcal{P}_5}|P|_{-5}) - F_4(\theta - \max_{P\in\mathcal{P}_4}|P|_{-4})f_5(\theta - \max_{P\in\mathcal{P}_5}|P|_{-5}).$$

(iii) Using minimal cut set $\mathcal{C} = \{2,3,4\}$, we take $Z = \{X_1, X_5, X_6\}$, and similar analysis yields

$$P_Z(Y>\theta) = 1 - F_2(\theta - \max_{P\in\mathcal{P}_2}|P|_{-2})F_3(\theta - \max_{P\in\mathcal{P}_3}|P|_{-3})F_4(\theta - \max_{P\in\mathcal{P}_4}|P|_{-4}),$$

again since $\mathcal{P}_2, \mathcal{P}_3$ and $\mathcal{P}_4$ partition the path space, and X_2, X_3, and X_4 are mutually independent. Differentiating, the final estimator is given by

$$\begin{aligned}\frac{dP_Z(Y>\theta)}{d\theta} = &-f_2(\theta - \max_{P\in\mathcal{P}_2}|P|_{-2})F_3(\theta - \max_{P\in\mathcal{P}_3}|P|_{-3})F_4(\theta - \max_{P\in\mathcal{P}_4}|P|_{-4})\\ &- F_2(\theta - \max_{P\in\mathcal{P}_2}|P|_{-2})f_3(\theta - \max_{P\in\mathcal{P}_3}|P|_{-3})F_4(\theta - \max_{P\in\mathcal{P}_4}|P|_{-4})\\ &- F_2(\theta - \max_{P\in\mathcal{P}_2}|P|_{-2})F_3(\theta - \max_{P\in\mathcal{P}_3}|P|_{-3})f_4(\theta - \max_{P\in\mathcal{P}_4}|P|_{-4}).\end{aligned}$$

(iv) Using minimal cut set $\mathcal{C} = \{1,5\}$, we take $Z = \{X_2, X_3, X_4, X_6\}$. Note, however, that $\mathcal{P}_1$ and $\mathcal{P}_5$ do not partition the path space, since path (1,3,5,6) is in both sets. We shall now see how this leads to difficulties:

$$\begin{aligned}P_Z(Y>\theta) &= 1 - P_Z(\max_{P\in\mathcal{P}}|P| \le \theta) \;=\; 1 - P_Z(\max_{P\in\mathcal{P}_1\cup\mathcal{P}_5}|P| \le \theta)\\ &= 1 - P_Z(X_1 + |146|_{-1} \le \theta, X_5 + |256|_{-5} \le \theta,\\ &\qquad X_1 + X_5 + |1356|_{-1,-5} \le \theta),\end{aligned}$$

which cannot be factored nicely as in the previous cases.

In general, if the $\{\mathcal{P}_i\}$ do partition the path space, i.e., $\mathcal{P} = \bigcup_{i\in\mathcal{C}}\mathcal{P}_i$ and $\mathcal{P}_i \cap \mathcal{P}_j = \phi$ for $i \neq j$, then

$$\begin{aligned}P_Z(Y>\theta) &= 1 - P_Z(\max_{P\in\mathcal{P}}|P| \le \theta) \;=\; 1 - P_Z(\max_{P\in\cup_i\mathcal{P}_i}|P| \le \theta)\\ &= 1 - P_Z(\bigcap_{i\in\mathcal{C}}\{\max_{P\in\mathcal{P}_i}|P| \le \theta\})\\ &= 1 - P_Z(\bigcap_{i\in\mathcal{C}}\{X_i + \max_{P\in\mathcal{P}_i}|P|_{-i} \le \theta\})\\ &= 1 - \prod_{i\in\mathcal{C}}P_Z(X_i + \max_{P\in\mathcal{P}_i}|P|_{-i} \le \theta)\\ &= 1 - \prod_{i\in\mathcal{C}}F_i(\theta - \max_{P\in\mathcal{P}_i}|P|_{-i}),\end{aligned}$$

which upon differentiation yields the estimator

$$-\sum_{i\in\mathcal{C}} f_i(\theta - \max_{P\in\mathcal{P}_i} |P|_{-i}) \prod_{j\neq i} F_j(\theta - \max_{P\in\mathcal{P}_j} |P|_{-j}),$$

where the product is equal to 1 if it is empty as in case (i) of the 5-node example just considered.

4 Directions for Future Research

In the general PERT/CPM stochastic activity network framework, we have derived some new unbiased Monte Carlo simulation-based derivative estimators for a setting where there are existing estimators available, and also for new settings. This gives the simulation analyst a wider array of options in carrying out performance evaluation. Depending on the types of networks and performance measures considered, different estimators may be more effective and useful. Clearly, comparisons of variance properties of the various estimators, whether through theoretical analysis and/or empirical/numerical experimental testing, would be beneficial to the potential user of these estimators. Preliminary experiments in [19] indicate that the SPA estimators are quite promising, as they generally show lower variance than the WD and LR estimators. However, in the 5-node Bernoulli example, it turns out that the SPA and one version of the WD estimator coincide. In the cases where they are applicable, the IPA estimators provide the lowest variance, which turns out to be essentially the same as a finite difference estimator with common random numbers, a procedure that can be easily implemented in many settings, but still requires additional simulations for every parameter of interest in the network.

The extension to arc criticalities is an important one. Another interesting extension to consider is the case where the input random variables (individual activity times $\{X_i\}$) are not necessarily independent. Using these derivative estimates in simulation optimization is another fruitful area of research; see, e.g., [9, 10, 12, 13]. Investigating quasi-Monte Carlo estimators and also estimators that are not sample means but quantiles or order statistics are also useful topics for future research.

Acknowledgments

This work was supported in part by the National Science Foundation under Grant DMI-0323220, and by the Air Force Office of Scientific Research under Grant FA95500410210. An earlier version was presented at the International Conference on Automatic Control and Systems Engineering, Cairo, Egypt, in December 2005 [11]. In addition to numerous editorial enhancements, substantive changes in this version include correcting the SPA estimator for $dP(Y > y)/d\theta$ in the 5-node example, adding the discrete (Bernoulli) case for the same example, and greatly expanding the scope of the analysis and adding the 5-node extended example for the $dP(Y > \theta)/d\theta$ SPA estimator.

References

1. V.G. Adlakha and H. Arsham. A simulation technique for estimation in perturbed stochastic activity networks. *Simulation*, 58:258–267, 1992.
2. A.A. Assad and B.L. Golden. PERT. in *Encyclopedia of Statistical Sciences*, Volume 6:691–697, 1992.
3. R.A. Bowman. Stochastic gradient-based time-cost tradeoffs in PERT network using simulation. *Annals of Operations Research*, 53:533–551, 1994.
4. R.A. Bowman. Efficient estimation of arc criticalities in stochastic activity networks. *Management Science*, 41:58–67, 1995.
5. R.A. Bowman. Sensitivity curves for effective project management. *Naval Research Logistics*, 50:481–497, 2003.
6. P. Brucker, A. Drexl, R. Mohring, K. Neumann, and E. Pesch. Resource-constrained project scheduling: Notation, classification, models, and methods. *European Journal of Operational Research*, 112:3–41, 1999.
7. J.G. Cho and B.J. Yum. Functional estimation of activity criticality indices and sensitivity analysis of expected project completion time. *Journal of the Operational Research Society*, 55:850–859, 2004.
8. S.E. Elmaghraby. On criticality and sensitivity in activity networks. *European Journal of Operational Research*, 127:220–238, 2000.
9. M.C. Fu. Optimization via simulation: A review. *Annals of Operations Research*, 53:199–248, 1994.
10. M.C. Fu. Optimization for simulation: Theory vs. practice (Feature Article). *INFORMS Journal on Computing*, 14:192–215, 2002.
11. M.C. Fu. Sensitivity analysis for stochastic activity networks. *Proceedings of the International Conference on Automatic Control and Systems Engineering* (on CD-ROM), 2005.
12. M.C. Fu. Stochastic Gradient Estimation, Chapter 19 in *Handbooks in Operations Research and Management Science: Simulation*, S.G. Henderson and B.L. Nelson, eds., Elsevier, 2006.
13. M.C. Fu, F.W. Glover, and J. April. Simulation optimization: A review, new developments, and applications. *Proceedings of the 2005 Winter Simulation Conference*, 83–95, 2005.
14. M.C. Fu and J.Q. Hu. *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*. Kluwer Academic, Boston, MA, 1997.
15. S.I. Gass. Decision-aiding models: validation, assessment, and related issues for policy analysis. *Operations Research*, 31:601–663, 1983.
16. S.I. Gass. Model accreditation: A rationale and process for determining a numerical rating. *European Journal of Operational Research*, 66:250–258, 1993.
17. S.I. Gass and L. Joel. Concepts of model confidence. *Computers and Operations Research*, 8:341–346, 1987.
18. S.I. Gass and B.W. Thompson. Guidelines for model evaluation: An abridged version of the U. S. general accounting office exposure draft. *Operations Research*, 28:431–479, 1980.
19. C. Gröer and K. Ryals. Sensitivity analysis in simulation of stochastic activity networks: A computational study; project report for BMGT835, Spring 2006. also submitted for presentation and publication, Tenth INFORMS Computing Society Conference on International Impacts of Computing, Optimization, and Decision Technologies.

20. S.M. Ross. *Stochastic Processes.* John Wiley & Sons, New York, NY, 1983.
21. R.G. Sargent. Personal communication, January 2006.

The EM Algorithm, Its Randomized Implementation and Global Optimization: Some Challenges and Opportunities for Operations Research

Wolfgang Jank

Robert H. Smith School of Business
Department of Decision and Information Technologies
University of Maryland
College Park, MD 20742
wjank@rhsmith.umd.edu

Summary. The EM algorithm is a very powerful optimization method and has become popular in many fields. Unfortunately, EM is only a local optimization method and can get stuck in sub-optimal solutions. While more and more contemporary data/model combinations yield multiple local optima, there have been only very few attempts at making EM suitable for global optimization. In this paper we review the basic EM algorithm, its properties and challenges, and we focus in particular on its randomized implementation. The randomized EM implementation promises to solve some of the contemporary data/model challenges, and it is particularly well-suited for a wedding with global optimization ideas, since most global optimization paradigms are also based on the principles of randomization. We review some of the challenges of the randomized EM implementation and present a new algorithm that combines the principles of EM with that of the Genetic Algorithm. While this new algorithm shows some promising results for clustering of an online auction database of functional objects, the primary goal of this work is to bridge a gap between the field of statistics, which is home to extensive research on the EM algorithm, and the field of operations research, in which work on global optimization thrives, and to stimulate new ideas for joint research between the two.

Key words: Monte Carlo EM; stochastic optimization; mixture model; clustering; global optimization; online auctions; functional objects.

1 Introduction

In this paper we want to shed new light on the Expectation-Maximization (EM) algorithm. The EM algorithm is a highly successful tool especially in statistics, but it has also received much attention outside the field. Applications include hierarchical models, neural networks, clustering and text mining.

The basic algorithm has been modified many times to overcome today's challenges such as complex data-models and huge databases. Yet, to this day, there has been barely any attempt at making EM suitable for solving global optimization problems.

The EM algorithm has experienced much success which can be partly attributed to its unique properties. One of these properties is that, in contrast to many other optimization methods, it guarantees an increase in the likelihood function in every iteration of the algorithm. Another property is that, since it operates on the log scale, it allows for significant analytical and numerical simplifications, especially for models in the exponential family. However, one of the biggest shortcomings of EM is that it is only a local optimization procedure and can consequently get stuck in sub-optimal solutions. One possible reason why this problem has not experienced much attention is that most applications of EM, especially within the statistics literature, have been simple enough to not experience significant impediments due to this shortcoming. However, increasing model- and data-complexity make this shortcoming more and more of a concern. It is well-known, for instance, that the mixture-model can have many local solutions, depending on the number of mixtures and the size and the dimension of the data, and that, as a consequence, the EM algorithm can produce solutions far from the true solution. EM is therefore more and more likely to yield unsatisfactory results in real-world applications.

Methods to find the global (and true) optimum are very prominent in the literature surrounding operations research, but surprisingly only very few of these methods have found their way into the classical statistics literature. This is very likely one of the reasons why EM is (still) associated with only local optimization qualities. Another reason may be a simple disconnect in language between the fields of operations research and statistics. The goal of this work is to bridge this gap and to stimulate joint research effort between the two fields.

In particular, we aim at bridging this gap by pointing to *randomized* EM implementations. Randomized EM implementations have been proposed to overcome complicated model-structures, and, more recently, also to overcome computational inefficiency due to huge databases. Randomized EM implementations have also been shown to be able to overcome local traps; however the ability to do so is based purely on chance, and there are no features in place to more systematically steer-free of such traps. Yet, one of the key features of randomized EM implementations is that, in contrast to the traditional EM algorithm, they are similar in nature to global optimization methods. Global optimization methods, like randomized EM implementations, exploit the principles of randomness to overcome local solutions, yet, unlike randomized EM, they do so more systematically. Thus, the randomized version of EM appears to be the natural ground for a wedding with the principles of global optimization.

We proceed as follows. In Section 2 we introduce the EM algorithm together with its basic properties and motivate the need for global optimization

qualities. We proceed by explaining randomized variants of EM and their usefulness. We devote Section 3 to challenges associated with randomized EM implementations to make the OR researcher familiar with implementation difficulties and possible research opportunities. In Section 4 we propose a possible combination between the EM algorithm and the principles of global optimization, but we are quick to point out that this may not be the only avenue, and our goal is simply to spark new research ideas. We proceed by applying our global optimization version of EM to a challenging curve-clustering problem of a large online auction database to illustrate some of its properties. We conclude with final remarks in Section 5.

2 The EM Algorithm and Its Randomized Implementation

2.1 Deterministic EM

The EM algorithm [12] is an iterative procedure for finding the maximum of likelihood functions in incomplete data problems. Let x denote the observed (or incomplete) data, and let z denote the unobserved (or missing) data. Then, in the notation of the EM algorithm, the pair (x, z) is referred to as the complete data. Let $f(x, z; \theta)$ denote the joint distribution of the complete data, dependent on a parameter vector θ. The goal is to find $\hat{\theta}$, the maximizer of the *marginal* likelihood[1] $L(\theta) \equiv L(\theta; x) = \int f(x, z; \theta) dz$. EM accomplishes that goal in iterative fashion. In each iteration, the EM algorithm performs an expectation and a maximization step. Let $\theta^{(t-1)}$ denote the current parameter value. Then, in the tth iteration of the algorithm, the E-step computes the conditional expectation of the complete data log-likelihood, conditional on the observed data and the current parameter value,

$$Q(\theta|\theta^{(t-1)}) = E\left[\log f(x, z; \theta)|\, x; \theta^{(t-1)}\right]. \tag{1}$$

This conditional expectation is often referred to as the "Q-function," since it plays a central role in the EM algorithm. The t^{th} EM update $\theta^{(t)}$ maximizes the Q-function. That is, $\theta^{(t)}$ satisfies

$$Q(\theta^{(t)}|\theta^{(t-1)}) \geq Q(\theta|\theta^{(t-1)}) \tag{2}$$

for all θ in the parameter space. This is the M-step. Given an initial value $\theta^{(0)}$, the EM algorithm produces a sequence $\{\theta^{(0)}, \theta^{(1)}, \theta^{(2)}, \ldots\}$ that, under mild regularity conditions [6, 53], converges to $\hat{\theta}$.

[1] Notice that $\hat{\theta}$ may only be a local maximum if several locally optimal solutions exist.

2.2 Example: EM for Model-Based Clustering

The EM algorithm is a very popular tool especially in clustering applications. Let $x_1, \ldots, x_n$ be a set of n independent p-dimensional data vectors arising from a mixture of a finite number of g groups in some unknown proportions $\pi_1, \ldots, \pi_g$. Specifically, we assume that the mixture density of the jth data point x_j $(j = 1, \ldots, n)$ can be written as

$$f(x_j; \theta) = \sum_{i=1}^{g} \pi_i f_i(x_j; \psi_i), \tag{3}$$

where the sum of the mixture proportions $\pi_i \geq 0$ $(i = 1, \ldots, g)$ equals one and the group-conditional densities $f_i(x_j; \psi_i)$ depend on an unknown parameter vector ψ_i. Let $\theta = (\pi_1, \ldots, \pi_g, \psi_1, \ldots, \psi_g)$ be the vector of all unknown parameters. Then, the corresponding log-likelihood is given by

$$\log L(\theta; x) = \sum_{j=1}^{n} \log\{\sum_{i=1}^{g} \pi_i f_i(x_j; \psi_i)\}. \tag{4}$$

The clustering goal is to calculate the parameter value $\hat{\theta}$ that maximizes the likelihood (4).

One can maximize the log-likelihood in (4) assuming that some of the information is unobserved. Specifically, we assume that the x_j's arise from one of the g groups. Let $z_1, \ldots, z_n$ denote the corresponding g-dimensional group-indicator vectors. That is, let the ith element of z_j equal one if and only if x_j originates from the ith group and zero otherwise. Notice that the group-indicator vectors z_j are unobserved. Let us write $x = (x_1, \ldots, x_n)$ for the observed (or incomplete) data and similarly $z = (z_1, \ldots, z_n)$ for the unobserved (or missing) data. Then the complete data is (x, z) and the complete data log-likelihood of θ can be written as

$$\log L_c(\theta; x, z) = \sum_{i=1}^{g} \sum_{j=1}^{n} z_{ij} \{\log \pi_i + \log f_i(x_j; \psi_i)\}, \tag{5}$$

where z_{ij} denotes the ith component of z_j.

The EM algorithm is an ideal method to solve this incomplete data problem, and it allows for significant simplifications. If we assume a normal mixture model, then the E-step and M-step are available in closed form, which makes the method straightforward to implement. That is, let

$$f_i(x_j; \psi_i) = \phi(x_j; \mu_i, \Sigma_i), \;\; \psi_i = (\mu_i, \Sigma_i), \tag{6}$$

where $\phi(\cdot; \mu, \Sigma)$ denotes the p-dimensional normal density with mean μ and covariance matrix Σ. Then, in the E-step we calculate the conditional expectation of the z_{ij}'s via [36]

$$\tau_{ij}^{(t-1)} = E(z_{ij}|x_j;\theta^{(t-1)}) = \frac{\pi_i^{(t-1)}\phi(x_j;\mu_i^{(t-1)},\Sigma_i^{(t-1)})}{\sum_{l=1}^{g}\pi_l^{(t-1)}\phi(x_j;\mu_l^{(t-1)},\Sigma_l^{(t-1)})} \tag{7}$$

for all $i = 1,\ldots,g$ and $j = 1,\ldots,n$. The normal case allows significant computational advantages by working with the corresponding sufficient statistics,

$$T_{i1}^{(t)} = \sum_{j=1}^{n}\tau_{ij}^{(t-1)} \tag{8}$$

$$T_{i2}^{(t)} = \sum_{j=1}^{n}\tau_{ij}^{(t-1)}x_j \tag{9}$$

$$T_{i3}^{(t)} = \sum_{j=1}^{n}\tau_{ij}^{(t-1)}x_j x_j^T. \tag{10}$$

In the M-step, we update the parameter estimates using only these sufficient statistics

$$\pi_i^{(t)} = T_{i1}^{(t)}/n \tag{11}$$

$$\mu_i^{(t)} = T_{i2}^{(t)}/T_{i1}^{(t)} \tag{12}$$

$$\Sigma_i^{(t)} = \{T_{i3}^{(t)} - T_{i1}^{(t)^{-1}}T_{i2}^{(t)}T_{i2}^{(t)^T}\}/T_{i1}^{(t)}. \tag{13}$$

The E-step and M-step are repeated until convergence. Convergence is often assessed by monitoring the improvements in the parameter estimates and/or the improvements in the log-likelihood function.

Notice that the normal mixture model presents one of the simplest cases to which the EM algorithm is applied. The conditional expectation in (7), and hence the E-step, is available in closed-form. Moreover, the M-step, by nature of (11), (12) and (13), also has a closed-form solution. This is not the case for most other models. For instance, simply replacing the normal mixture density (6) by another density, especially one outside the exponential family, yields a much more involved version of the EM algorithm. Other examples include the very popular hierarchical models, or mixed models, which result in E-steps with intractable integrals, which are typically also of very high dimension [34, 3]. In that case, one has to resort to approximation at the E-step and quite often also at the M-step. This leads to a randomized EM implementation which we will describe in more detail below.

2.3 Properties of EM

EM is popular within statistics but also in many other areas. One reason for this popularity are its widely-appreciated properties [48].

Arguably, one of the most outstanding properties of the EM algorithm is that, unlike many other optimization methods, it guarantees an improvement of the likelihood function in every update of the algorithm. Specifically, it can be shown [12] that any parameter value θ^* that satisfies

$$Q(\theta^*|\theta^{(t-1)}) > Q(\theta^{(t-1)}|\theta^{(t-1)}) \tag{14}$$

results in an increase of the likelihood function; that is, $L(\theta^*) > L(\theta^{(t-1)})$. In fact, one can show that the log-likelihood can be written as $\log L(\theta) = Q(\theta|\theta^{(t-1)}) - H(\theta|\theta^{(t-1)})$ where $H(\theta|\theta^{(t-1)}) = E\left[\log f(z|x;\theta)|\, x;\theta^{(t-1)}\right]$ is the conditional expectation of the conditional density of the missing data, $f(z|x;\theta)$. The likelihood-ascent property is then a simple consequence of Jensen's inequality applied to (14).

The likelihood-ascent property implies that the output of EM always *increases* the likelihood function, which is in contrast to, say, Newton-Raphson, which, upon convergence, requires additional verification that the final value is not a minimizing point. In addition, this property also alleviates the maximization step. Recall that the M-step, at least in principle, requires a full maximization of the Q-function in (1). This can be hard or even impossible depending on the complexity of the model. The likelihood-ascent property alleviates this problem in the sense that any parameter value θ^* that satisfies (14), and not necessarily only the full maximizer of (1), will contribute to the overall progress of the algorithm. That version of EM is often referred to as a *Generalized* EM (GEM) algorithm [12]. A related version is the *Expectation-Conditional-Maximization* (ECM) algorithm which also relieves complicated M-steps by breaking them up into simpler, conditional maximization steps [38].

But the EM algorithm is also known and appreciated for other convenient properties. EM is popular, because it operates on the log-scale and therefore allows significant analytical simplification, especially for models in the exponential family (such as the normal mixture model described in Section 2.2). Another by-product of the log-scale is that the method has a wider domain of attraction than, say, Newton-Raphson, and as a consequence enjoys much greater numerical stability especially with respect to the choice of the starting values. Notice though that, in contrast to Newton-Raphson, EM is typically applied only to problems where the goal is to maximize a likelihood function.

EM is also a very adaptive method that can be tailored towards many modern real-world problems. Recent modifications of EM have been used to handle huge databases in an efficient manner [39, 49], and also to update parameter estimates and predictions in real time [40, 41].

And finally, while EM has received a lot of interest from other areas, it has been – and still is – very popular in statistics. One reason for that may be the seminal paper by [12] in the statistics literature, which resulted in an early exposure to the principles of EM within the statistics community. Over the years, statisticians may have grown very familiar with the method and very comfortable with its properties, which may explain some of its popularity.

Another reason may lie in the principles of missing data, data augmentation and imputation which the method exemplifies. Missing data have always been at the heart of statistics research, and the EM algorithm embodies modern solutions to that problem. At any rate, the EM algorithm is extremely popular, and there is an opportunity to wed the method with other algorithms, particularly those developed in the OR literature, in order to tackle new and contemporary optimization problems.

2.4 Applications of EM

The EM algorithm has found an array of different applications. One of the more common applications within the statistics literature is for the fitting of linear mixed models or generalized linear mixed models [34, 35]. Another very common application is for the estimation of mixture models [36]. Other applications range from mixtures of experts [26], neural networks [1], signal processing [13], text mining [42], graphical models [28] and many, many more.

2.5 Challenges of EM

One of the biggest challenges for the EM algorithm is that it only guarantees convergence to a *local* solution. The EM algorithm is a greedy method in the sense that it is attracted to the locally optimal solution closest to its starting value. This can be a problem when several locally optimal solutions exist. This problem frequently occurs in the mixture model (3).

Consider Figure 1. The top panel of Figure 1 shows 40 observations, $x_1, \ldots, x_{40}$, simulated according to a mixture of two univariate normal distributions, $x_i \overset{iid}{\sim} [p_1 N(\mu_1, \sigma_1^2) + p_2 N(\mu_2, \sigma_2^2)]$, with $p_1 = p_2 = 0.5$, $\mu_1 = -1$, $\mu_2 = 2$, $\sigma_1^2 = 0.001$ and $\sigma_2^2 = 0.5$. Notice that this is a special case of the normal mixture model in (3) with $p = 1$ and $g = 2$. Notice also that the first mixture component has almost all its mass centered around its mean $\mu_1 = -1$. This results in a log-likelihood for μ_1 depicted in the bottom panel of Figure 1. We can see that, as expected, the global optimum of this log-likelihood is achieved at $\mu_1 = -1$. However, we can also see at least five local optima, located around the values $\mu_1 = 1, 1.5, 2, 2.5$ and 3. Clearly, depending on where we start EM, it may be trapped very far away from the global (and true) parameter value.

There has been extensive work on solving the above optimization problem. One very promising solution is via the cross-entropy (CE) method (see e.g. `http://iew3.technion.ac.il/CE/`). In fact, [5] compare CE with the EM algorithm for solving mixture problems such as in (3) and find a superior global-search performance for CE. However, the authors fail to point out some of the shortcomings of CE. In fact, CE is based on simulating candidate solutions θ's from a suitable distribution and then picking, among all candidates, those with the best performance. Simulating candidates efficiently

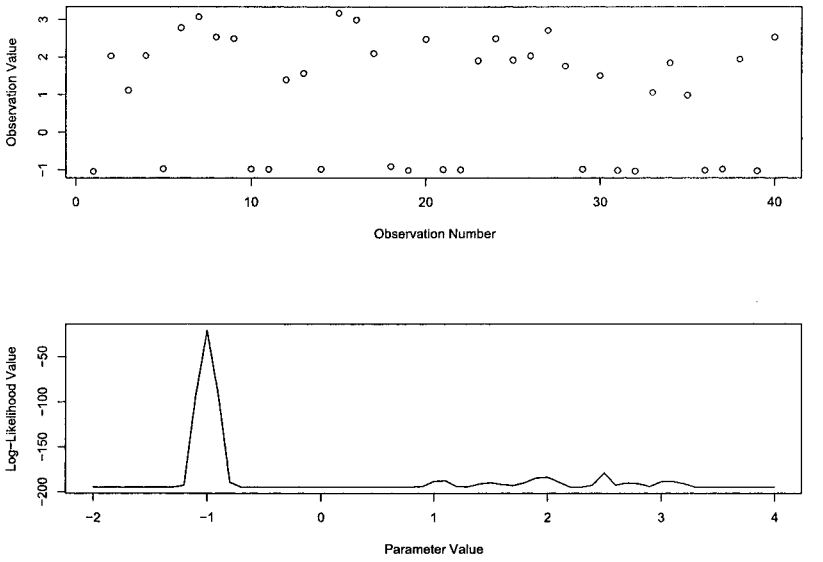


Fig. 1. Log-likelihood function for a simple two-component mixture problem. The top panel shows the simulated data. The bottom panel shows the log-likelihood function for μ_1, the mean of the first likelihood component, holding all other parameters constant at their true values.

though can be problematic in the mixture model. The reason is that, for instance in the case of the normal mixture, the parameter of interest is given by $\theta = (\pi_1, \ldots, \pi_g, \mu_1, \ldots, \mu_g, \Sigma_1, \ldots, \Sigma_g)$. That is, each candidate contains, among other things, the covariance matrices Σ_i of the multivariate normal distribution which, by definition, have to be positive definite. However, simulating positive definite matrices, especially in higher dimensions, is not at all obvious and can be computationally very challenging [17]. Now consider again the EM algorithm. EM overcomes this problem statistically by relying entirely on the sufficient statistics (13) for estimating Σ. But these sufficient statistics yield, per construction, positive definite matrices by default. This suggests that a wedding between the ideas of global optimization and the principles of EM could be advantageous on many different fronts!

2.6 Randomized EM Implementations

The EM algorithm is a *deterministic* method. What we mean by that is that it converges to the same stationary point if initiated repeatedly from the same

starting value. This is in contrast to *randomized* versions of EM. Randomized EM versions distinguish themselves from their deterministic origin in that repeat applications from the same starting value will *not* necessarily lead to the same answer. Randomized EM versions have become popular with the availability of more and more powerful computing. Randomized EM versions overcome many of the computational limitations that EM encounters in complex models. In what follows, we will describe the most basic randomized EM implementation. Our point of view is strongly influenced by the Monte Carlo EM algorithm, so we are quick to point out that this may not be the only viewpoint. However, it does not matter too much which point of view one assumes, because, in the end, all randomized EM versions are related to one another.

Many contemporary models result in a complicated E-step. This complication can be due to analytical intractability [34] or due to computational intensity [41] of the Q-function in (1). One remedy against this problem is by approximating the Q-function appropriately. While approximation can be done in several ways, by far the most popular approach is via simulation. This leads to the concept of Monte Carlo and the Monte Carlo EM algorithm.

The Monte Carlo EM (MCEM) algorithm, in its most basic form, has been around for over 10 years [52]. MCEM simply approximates the expectation in (1) by the Monte Carlo average

$$\tilde{Q}_{m_t}(\theta|\theta^{(t-1)}) = \frac{1}{m_t}\sum_{k=1}^{m_t} \log f(x, z_k; \theta), \tag{15}$$

where $z_1, \ldots, z_{m_t}$ are simulated from the conditional distribution of the missing data, $f(z|x;\theta^{(t-1)})$. Then, by the law of large numbers, $\tilde{Q}_{m_t}$ will be a reasonable approximation to Q if m_t is large enough. The MCEM algorithm proceeds in the same way as its deterministic counterpart, simply replacing Q by $\tilde{Q}_{m_t}$. We refer to this algorithm as the basic randomized EM implementation. Other randomized EM implementations include a stochastic-approximation version of the Q-function [11] or versions with $m_t = 1$ for all t [9]. We will get back to those version later in this manuscript.

Notice that our definition in (15) is general and applies particulary to those situations where the E-step has no closed-form solution. In those instances where the conditional expectation (1) can be calculated analytically, such as for the normal mixture in (7), the randomized version of EM simplifies significantly, since then simulation from the potentially very complicated conditional distribution $f(z|x;\theta^{(t-1)})$ can be traded-in for simple random sampling. We describe this case next.

2.7 Example: Randomized EM for Model-Based Clustering

One can derive a randomized EM version for the normal mixture model in (3) readily. Notice that for this model, the Q-function in (1) is given by the conditional expectation over $\log L_c(\theta; x, z)$ in (5), which is available in closed-form

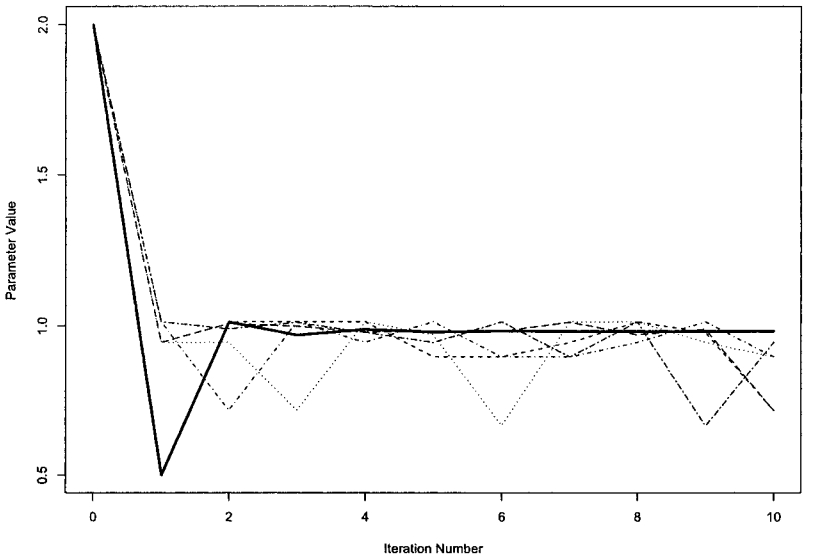


Fig. 2. Parameter-path of EM and randomized EM for the example in Section 2.5. We estimate only μ_1 and hold all other parameters constant at their true values. Each method is started at $\mu_1 = 2$ and run for 10 iterations. The thick solid line shows the parameter path of EM. The broken lines show the parameter path of 5 separate runs of randomized EM in (16) using a constant sample size $m_t = 20$ per iteration.

and given by the sum $\sum_{ij} \tau_{ij}\{\log \pi_i + \log f_i(x_j;\psi_i)\}$. Thus, we can approximate $Q(\theta|\theta^{(t-1)})$ simply by sub-sampling the entire database. Let $(x_1,\ldots,x_n)$ denote the full database and let $(x_1,\ldots,x_{m_t}) \subset (x_1,\ldots,x_n)$ be a randomly chosen sample of size m_t $(m_t < n)$. We can then approximate the Q-function in (1) by

$$\tilde{Q}_{m_t}(\theta|\theta^{(t-1)}) = \sum_{j=1}^{m_t}\sum_{i=1}^{g} \tau_{ij}^{(t-1)}\{\log \pi_i + \log f_i(x_j;\psi_i)\}. \tag{16}$$

Notice that as $m_t \to n$, $\tilde{Q}_{m_t} \to Q$. Thus, if we use $\tilde{Q}_{m_t}$ instead of Q, we sacrifice accuracy (by using only an approximation to the Q-function) for computational efficiency (by using only a small subset $x_{t_1},\ldots,x_{t_m}$ instead of the entire database). [8] propose a novel approach based on the likelihood-ascent property for finding a good balance between accuracy and computational efficiency in each EM iteration [25].

Notice that in order to implement a randomized version of EM, some care is necessary. One has to carefully choose the way in which one simulates[2], how much one simulates (i.e. the simulation sample size m_t), for how many iterations one runs the method, how one diagnoses convergence, and more. Consider Figure 2 which shows the parameter path of EM and randomized EM for the example in Section 2.5. The thick solid line corresponds to EM and the broken lines correspond to 5 separate runs of randomized EM. Notice that we only estimate μ_1 and hold all other parameters constant at their true values. Recall from Figure 1 that the global optimum for μ_1 is near -1.

We notice first that all algorithms converge to a value near 1, which is far from the global optimum. But notice also that while the EM algorithm stabilizes at 1 after only a few iterations, all 5 randomized EM implementations continue to show random variation. This is due to the error inflicted by only using a Monte Carlo approximation in (16) rather than the true Q-function. This error presents new challenges for implementing the method. In particular, one has to decide how large a value of m_t one desires in every iteration. Larger values of m_t yield a more accurate approximation $\tilde{Q}_{m_t}$; however, they also result in a higher computational burden especially when the database is large. The EM algorithm is also known to make good progress with smaller values of m_t (and thus a less accurate $\tilde{Q}_{m_t}$) in the early iterations. Either way, the random fluctuations of the randomized EM updates also present new challenges in terms of monitoring the method and diagnosing its convergence. We will give an overview over the different challenges, some existing solutions, and associated research opportunities in Section 3.

2.8 Advantages of Randomized EM

There are several advantages associated with a randomized implementation of the EM algorithm. Most commonly, randomized EM version are used to overcome intractable, typically high-dimensional integrals at the E-step [34, 2]. More recently, randomized implementations, if implemented intelligently, have also been found to be able to speed-up convergence of the EM algorithm [8]. But randomized EM versions can do more.

Randomized variants of EM may in fact be able to overcome local solutions [9, 11]. The basic idea is to inflict noise in the form of random perturbations into the deterministic updating scheme of EM. The hope is that a random perturbation will "push" the method away from a local trap and lead to a better solution. The Monte Carlo EM algorithm is stochastic in nature since it selects the samples $z_1, \ldots, z_{m_t}$ in random fashion, and therefore two runs of Monte Carlo EM from the same starting values, as seen for instance in Figure 2, will lead to different parameter paths. Thus there is a chance that

[2] As pointed out earlier, for the mixture model simulation reduces to simple random sampling; however, in more complex models this is typically more complicated and one has to simulate from $f(z|x;\theta^{(t-1)})$.

Monte Carlo EM can overcome local traps. However, it is also clear that the ability to overcome these traps is entirely due to chance, since the basic form of the algorithm has no tools in place that do so more systematically. This is in contrast to, say, the Genetic Algorithm which selects, among a group of random candidate solutions, those which show better promise of success in the next generation.

In this paper we present a version of the Monte Carlo EM algorithm with more systematic, built-in features useful to steer-free of local traps. This version is based on the ideas of evolutionary computation and borrows concepts from the Genetic Algorithm. We first give an overview over challenges when implementing a randomized EM version and discuss associated research opportunities in Section 3. We then present our global optimization version of EM in Section 4.

3 Challenges and Opportunities of Randomized EM

In this section we discuss problems and challenges associated with the randomized implementation of EM. These challenges can be classified into five major categories: simulation, approximation, maximization, iteration and convergence.

3.1 Simulation

The first step in implementing a randomized EM version involves simulating from the conditional distribution $f(z|x;\theta)$ of the missing data. Three basic questions have to be answered in this step. First: What type of simulation do we want to use? Three different simulation paradigms are available to us: i.i.d. simulation via rejection sampling [3], independent simulation via importance sampling [3], or generating dependent samples via MCMC (Markov Chain Monte Carlo) [34, 8]. All of these three paradigms have their own set of advantages and disadvantages. The second question that has to be answered is: How can we use this sampler in the most efficient way? (That is, with the least amount of simulation effort.) While much of the EM literature focuses on finding the correct sampler, there exist approaches (typically outside the literature on EM) that promise a more efficient use of the simulated data using variance-reduction methods like Quasi-Monte Carlo. And the last question to be addressed is: How much simulation do we need? This question can also be asked in a slightly different way: Do we want to keep the amount of simulation constant in every EM iteration? And if not, how do we increase the simulations throughout the EM iterations in an automated way? Naturally, if we do not keep the amount of simulation constant, then we would like to increase them in a way that results in a) the least amount of simulated data at the end of the day (i.e. the shortest computing times), while guaranteeing b) sufficient accuracy of the results.

Type of Simulation

The challenge is to find a sampler that draws from $f(z|x;\theta)$, or at least from a distribution very close to it. We will refer to $f(z|x;\theta)$ as the target distribution. The problem is easier if z is of smaller dimension or, alternatively, breaks-down into several low-dimensional components. On the other hand, finding an appropriate sampler can be complicated if z is of very high dimension. More complex models, involving, say, interaction terms, spatial and/or longitudinal components, often result in high-dimensional z's.

Three basic types of simulation are available. Rejection sampling attempts to simulate i.i.d. draws exactly from $f(z|x;\theta)$. While i.i.d. draws are most preferable, this type of sampler is, at least in practice, the most limited one, since its efficiency typically declines with increasing complexity of the problem. Rejection sampling can be analytically challenging in the set-up in that it requires calculation of a certain supremum, but this can be alleviated using recent modifications [7]. Rejection sampling draws from the target distribution by thinning out samples from a candidate distribution, rejecting those that are not appropriate. Unfortunately, finding good candidate distributions gets harder and harder as the model-complexity increases and as a consequence the acceptance-rate drops. For that reason, rejection sampling works well only for low-dimensional settings.

An alternative to rejection sampling is importance sampling. In contrast to rejection sampling, importance sampling does not dispose of any of the simulations. Importance sampling draws from a suitably chosen importance distribution and accounts for the discrepancy between the target and the importance distribution by weighing the samples via importance weights. The main two advantages of importance sampling are that it produces independent samples and that is uses all of the simulation output. The disadvantage is that its performance depends heavily on the chosen importance sampler. In fact, the efficiency of the method can be extremely poor if the sampler is chosen without care [22]. If the target distribution is not too skewed, then decent importance samplers can be obtained via a multivariate Normal or t-distribution, shifted and scaled by the Laplace approximation to the mean and variance of $f(z|x;\theta)$ [24]. Importance sampling has been successfully applied to high-dimensional problems, but it is generally recommended to monitor the magnitude of the importance weights to ensure numerical stability.

The third alternative for simulation from $f(z|x;\theta)$ is to use MCMC. MCMC produces a sequence of draws that depend on one another in the sense that at each stage of the method, the sequence either moves to a new value or remains at the current one. Thus, depending on the mixing-properties of the chain, the resulting MCMC sampler can feature strong long-range dependencies. The MCMC methodology is especially appealing if the target distribution is of non-standard form, since it promises exact draws from $f(z|x;\theta)$. It is therefore conceptually also very appealing for high-dimensional simulation. The drawback of MCMC is that, as pointed out above, it produces

dependent samples which makes variance-estimation hard. Moreover, it only produces draws from the target after reaching the stationary distribution, so initial simulation is typically discarded. However, determining exactly how much has to be discarded (i.e. the amount of burn-in) is not easy.

In comparing importance sampling and MCMC, the question comes up as to when should which method be used? Unfortunately there exist hardly any general recommendations or formal comparisons between the two. Quite often, the choice is ultimately influenced by personal preference. Experience also shows that both work well when dealing with high-dimensional integration. While one produces independent samples, it also carries the burden of the importance weights. On the other hand, while the other one has no importance weights to worry about, it can take very long to converge to the stationary distribution, and the strong dependence-structure may also make it hard to calculate appropriate standard errors. An overview of importance sampling and MCMC within the context of MCEM can be found in [8].

Efficiency of Simulation

The aspect of efficient simulation has, for the most part, only played a secondary role in the literature on EM. While a strong emphasis has been placed on simulating from the *correct* distribution, only little attention has been paid to whether this can also been done in an *efficient* way. By efficient simulation, we mean simulation that produces estimates with very little variance. Clearly, the variance of the estimates can be reduced simply by increasing the size of the simulated data. However, this can become computationally intensive and time consuming in more complex problems. The question is whether we can produce samples from the correct (or near-correct) distribution, and also do so with the least possible amount of computation.

Variance reductions techniques attempt to make more efficient use of the simulated data. There exist a variety of variance reduction techniques such as antithetic variables, control variates or stratification. One particular set of methods that has received a lot of interest in the simulation literature is Quasi-Monte Carlo, which is related to Monte Carlo in that it uses simulation to approximate an intractable integral. However, in contrast to classical Monte Carlo, Quasi-Monte Carlo does not use *random* draws. In fact, Quasi-Monte Carlo produces a sequence of *deterministic* numbers with the best-possible spread in the sampling space. This sequence is also referred to as low-discrepancy sequence. There have been many examples where Quasi-Monte Carlo significantly beats classical Monte Carlo methods by a factor of 10, 100 or even 1,000 [30].

One drawback of Quasi-Monte Carlo is that it is deterministic in nature and, therefore, statistical methods do not apply for error estimation. Recent advances in *randomized* Quasi-Monte Carlo methodology [29] can overcome this drawback. Randomized Quasi-Monte Carlo combines the variance-reduction benefits of Quasi-Monte Carlo with the statistical error-estimation

properties of classical Monte Carlo. One way of generating randomized Quasi-Monte Carlo sequences is to initiate several parallel sequences from randomly chosen starting points [51]. While Quasi-Monte Carlo has been, to date, mostly used within the context of importance sampling, there exist efforts to apply its ideas to MCMC [43].

[24] proposes an automated MCEM algorithm based on randomized Quasi-Monte Carlo methods. The method uses Quasi-Monte Carlo to simulate from $f(z|x;\theta)$ based on Laplace importance sampling[3]. It also uses the ideas of randomized Quasi-Monte Carlo to measure the error of the integral-estimate in every iteration of the algorithm. The resulting Quasi-Monte Carlo EM (QMCEM) algorithm is embedded within the framework of the automated MCEM formulation proposed by [3].

Amount of Simulation

There exist two basic philosophies when it comes to choosing the simulation size for randomized EM implementations. One philosophy picks a value for the sample size m_t and holds this value fixed throughout all iterations. This philosophy is associated with *stochastic approximation* versions of EM [11] and we will get back to it later. The other philosophy increases the sample size steadily throughout all iterations. A complicating factor with the latter approach is that the sample size has to be determined anew in every iteration if the approach is supposed to result in efficient use of the simulations.

When approximating Q by $\tilde{Q}$, the Monte Carlo sample size m_t has to be increased successively as the algorithm moves along. In fact, [4] argue that MCEM will never converge if m_t is held fixed across iterations because of a persevering Monte Carlo error [10]. While earlier versions of the method choose the Monte Carlo sample sizes in a deterministic fashion before the start of the algorithm [34], the same deterministic allocation of Monte Carlo resources that works well in one problem may result in a very inefficient (or inaccurate) algorithm in another problem. Thus, data-dependent (and user-independent) sample size rules are necessary (and preferred) in order to implement MCEM in an automated way. Automated MCEM implementations have been proposed by several researchers [3, 31, 32, 8].

[3] are the first to propose an automated implementation of MCEM. Using a Taylor-series argument, they derive approximate confidence bounds for the MCEM parameter estimate under independent sampling schemes. Then, when the next update falls within this confidence bound, it is said to be swamped with Monte Carlo error, and consequently a larger sample size is needed to obtain a more accurate estimate of the Q-function. [31] and [32] build upon [3]'s method for MCMC sampling. [8] propose a new approach based on the

[3] Laplace importance sampling uses a multivariate Normal or multivariate t importance sampler shifted and scaled by the Laplace approximation to the mean and variance of the target distribution

difference in the Q-functions. Their approach has several advantages compared to the earlier implementations. First, it operates on a univariate quantity, and therefore makes it easier to incorporate more complicated sampling schemes like MCMC or Quasi-Monte Carlo under one umbrella. Second, it assures that EM's famous likelihood-ascent property holds in every iteration, at least with high probability, and thus counterproductive use of the simulation is rare. It also results in more stable variance-covariance estimates, since the final-iteration sample size is typically larger than in previous approaches. And finally, by using a one-sided confidence bound approach on the difference in the Q-functions, their method encourages parameter updates with a larger likelihood-increase than the deterministic EM algorithm and thus can result in a de facto acceleration of the method.

One question remains: If we don't want to increase the sample size in every iteration of MCEM, what alternatives do we have? One potential solution is to hold the sample size fixed and simply average over the MCEM output. This makes somewhat sense, because once the method reaches the stationary point, it fluctuates randomly about it with constant noise. Clearly, averaging is a very straightforward approach and it is likely to be very popular with many researchers who do not want to invest much effort into learning an automated implementation. However, what are the dangers in simply averaging over the MCEM output? If we start averaging too early, then our estimate is likely to be biased, since we average over early updates which are far from the solution. Bias can also happen if we start averaging too late, and, additionally, the estimate will then also be very noisy. But determining when to start averaging (and when to stop) is a problem that is at least as equally challenging as finding the right amount of simulation.

3.2 Approximation

After generating draws from the conditional distribution of the missing data, one approximates the Q-function via the empirical average $\tilde{Q}$ in (15). This is in principle straightforward to do and does not pose any major problems. However, there are modifications of the basic approach that promise several advantages. One of these modifications is to re-use the samples via importance re-weighting; another approach is to employ a stochastic approximation scheme.

[45] propose to use some of the simulated data repeatedly throughout the EM iterations by strategically re-weighting them. While this approach seems very reasonable, it has not received much popularity. One reason for this could be that it is hard to come up with automated sample size rules for a re-weighing scheme.

Another variant of the basic approach is to use a stochastic approximation version of EM [11]. That is, rather than steadily increasing the Monte Carlo sample size throughout the algorithm, there exist versions that converge with a constant (and typically small) value of m_t. Let γ_t be a sequence of positive

step sizes such that $\sum \gamma_t = \infty$ and $\sum \gamma_t^2 < \infty$ and define

$$\tilde{P}^{(t)}(\theta) = (1 - \gamma_t)\tilde{P}^{(t-1)}(\theta) + \gamma_t \tilde{Q}(\theta|\theta^{(t-1)}). \tag{17}$$

Notice that $\tilde{P}^{(t)}(\theta)$ is a convex combination of the information from the current iteration, and all the information from the past iterations. One typically initializes the method by setting $\tilde{P}^{(0)} = 0$. Notice that the recursion in (17) is similar to the stochastic approximation method of [47]. For that reason it is referred to as the Stochastic Approximation EM (SAEM) algorithm. One of the noteworthy features of this algorithm is that it converges (a.s.) with constant value of m_t [11]. It is also conceptually very appealing, since, by the recursion in (17), it makes use of all the simulated data.

Another appeal of the method is that, at least in principle, the only decision that has to be made is the choice of the step sizes γ_t. This is a one-time decision which is usually done *before* starting the algorithm. [44] show that for step sizes $\gamma_t \propto (1/t)^\alpha$, $1/2 < \alpha \leq 1$, the method converges at an (asymptotically) optimum rate (if used in conjunction with offline averaging). Thus, at first glance, the method appears to be easier to implement than the Monte Carlo EM algorithm, which requires a decision about the new value of m_t *in every iteration.* However, as it is often the case, no lunch is free. Indeed, while large step sizes (i.e. $\alpha \approx 1/2$) quickly bring the method in the neighborhood of the solution, they inflate the Monte Carlo error. On the other hand, while small step sizes (i.e. $\alpha \approx 1$) result in a fast reduction of the Monte Carlo error, they slow down the rate of convergence of the method.

[23] connects the problem of finding the right SAEM step size with EM's missing information principle. It is know that the convergence rate of EM depends on the fraction of missing-to-complete information [37]. In particular, if the fraction of missing information is large (and thus EM's convergence rate is already very slow), then it appears unwise to choose small step sizes and thereby even further slowing down SAEM. On the other hand, if EM converges fast, then a large step size introduces an unnecessary amount of extra noise, which should be avoided. [23] estimates EM's rate of convergence from the data and uses this estimate to choose a step size that balances the improvements in bias and variance of SAEM.

3.3 Maximization

The challenges in the M-step are in principle the same as those for the deterministic EM algorithm. If the M-step has no closed form solution (which is typically not the case), then one has to resort to numerical methods to maximize the Q-function. The most common approach is to use a version of Newton-Raphson, but alternative approaches also exist. The Newton-Raphson procedure has several nice features: It converges at a quadratic rate, and it is often straightforward to implement. One drawback of Newton-Raphson (as with many other optimization routines) is that it requires relatively good

starting values. Satisfactory starting values can often be found using a coarse grid search.

Another drawback of Newton-Raphson is that it requires evaluation of the Hessian matrix in every iteration. In situations when the Hessian is computationally too involved or numerically instable, quasi-Newton methods can be used instead. The methods of Davidson Fletcher-Powell (DFP) and Broyden-Fletcher-Goldfarb-Shanno (BFGS) are quasi-Newton procedures that only rely on the gradient of the objective function and are implemented in many software packages.

There exist further modifications of EM that are particularly aimed at simplifying or even accelerating its M-step. See for example the Newton-Raphson type modifications of [19], [20] or [27]. [38] on the other hand propose to break-up a complicated M-step into smaller, more tractable conditional M-steps [33].

3.4 Iteration

In this section we discuss for how many iterations one should run the method. The right number of iterations is closely connected with the choice of the stopping rule. Finding appropriate stopping rules for randomized EM versions is challenging. The deterministic EM algorithm is typically terminated if the relative change in two successive parameter estimates is small, smaller than some pre-defined threshold. The same stopping rule is not very useful for its randomized counterparts. The reason for this is that any deterministic stopping rule can be satisfied by a randomized method simply because of random chance, and not because convergence has occurred. Recognizing this, [3] recommend to apply a deterministic rule for several successive times, thereby reducing the chances of a premature stop. In the following we review alternative approaches for stopping a randomized EM algorithm.

[8] suggest terminating MCEM when the difference in likelihood functions becomes small. However, rather than directly estimating the likelihood-differences, they appeal to EM's likelihood-ascent property and instead operate on the differences in the Q-functions. This allows for an efficient implementation at no extra simulation expense. Other approaches are possible. [16] for example propose to monitor the gradient of the likelihood [15].

Choosing the right stopping rule can be extra hard for SAEM, especially when an already slowly converging EM algorithm is coupled with a small step size. Standard stopping rules do not take into account the effect of the step size or EM's convergence rate. [23] proposes a new way of monitoring SAEM. The approach is based on EM's likelihood-ascent property and measures the long-range improvements in the parameter updates over a flexible time window. It also provides a heuristic to gauge whether there still exists a significant trend in the long-range improvements based on the ideas of permutation tests.

And lastly a comment on the relationship between stopping rules and automated sample size rules: Automated sample size rules generally make it

easier to find reasonable stopping rules for MCEM, because the resulting algorithm mimics more closely a deterministic algorithm for which reasonable stopping rules are already well-established. On the other hand, consider the "quick-and-dirty" MCEM implementation via averaging of the parameter updates discussed earlier. The resulting averages will generally still show a good amount of variability (especially if we use a moving-average approach with a fixed time window). Consequently we cannot rely on deterministic stopping rules based on the change in the average parameter updates. For the same reason it is also harder to implement likelihood-based rules. Clearly, while averaging is, at least at first glance, a seemingly convenient approach and easier to implement than automated sample size rules, it does come with a whole additional package of complicating factors.

3.5 Convergence

The EM algorithm converges to the maximum of the likelihood function [6, 53]; that is, at least to a local maximum. Randomized versions of EM typically mimic that behavior under very mild regularity conditions [11, 14]. It is important to point out though that this convergence only occurs if the Monte Carlo sample size is increased successively, which again underlines the importance of automated sample size rules. (Increasing m_t is of course only necessary for MCEM; SAEM converges, as discussed earlier, with a fixed m_t.)

While randomized EM versions typically converge to a *local* maximum, there is no guarantee that this value is also the *global* optimum. The EM algorithm is a greedy method in the sense that it is attracted to the solution closest to its starting value. This can be a problem when several sub-optimal solutions exist. The mixture model, for example, is well-known to feature many local maxima, especially when the number of mixture-components is large (see again Section 2.5). This puts the additional burden on the researcher that any solution found by EM may not be the best solution and far from the true solution. One ad-hoc approach to alleviate this problem is to initialize EM from a variety of different starting values, but this approach can be burdensome if the parameter-space is of high dimension.

It is quite surprising that, despite the popularity of the EM algorithm, there have been only very few attempts at making it suitable for global optimization. Some of the exceptions include [25] and [50]. One possible reason is the disconnect in the literature between the field of statistics and that of operations research. While much of the theory for the EM algorithm has been developed in statistics, its literature has been relatively ignorant of the principles of global optimization that have been developed, in large parts, in operations research. Another reason may be a difference in language between the two fields. Either way, we hope to bridge some of the gap with this work and spark some ideas for cross-disciplinary research.

4 Global Optimization with Randomized EM

4.1 A Genetic Algorithm Version of EM (GAEM)

In the following we propose a new algorithm that combines the ideas of global optimization with that of the basic principles of the EM algorithm. This algorithm is based on earlier work by [25], who describes additional implementation variants using some of the ideas outlined in Section 3.

There exist a variety of different global optimization paradigms. One very popular approach is the concept of evolutionary computation. Evolutionary computation is associated with the groundbreaking work of [18]. Evolutionary algorithms find their inspiration from natural selection and survival of the fittest in the biological world. These algorithms weed out poor solutions, and combine good solutions with other good solutions to create new generations of even better solutions. Our algorithm combines the ideas of evolutionary computation with the principles of the randomized EM algorithm. We want to point out though that the proposed algorithm is certainly not the only possible avenue towards making EM suitable for global optimization, and different, possibly more efficient versions could be derived with additional research.

The genetic algorithm (GA) belongs to a general class of global optimization procedures that imitate the evolutionary process of nature. The basic building blocks of GA are *crossover*, *mutation* and *selection*. GAs are iterative, and each iteration is called a new generation. Starting from a parent population, two parents create offspring via crossover and mutation. The crossover operator imitates the mixing of genetic information during reproduction. The mutation operator imitates the occasional changes of genetic information due to external influences. An offspring's fitness is evaluated relative to an objective function. Offspring with the highest fitness are then selected for further reproduction. Although GA operations appear heuristic, [18] provides theoretical arguments for convergence to a high quality optimum.

We use the ideas of crossover, mutation and selection and combine them with the basic principles of the EM algorithm in the following way. Let $\theta^{(0,1)}, \theta^{(0,2)}, \ldots, \theta^{(0,R)}$ denote a set of R distinct starting values, possibly randomly chosen. Let $\{\theta^{(t,r)}\}_{t\geq 0}$, $1 \leq r \leq R$, denote a sequence of parameter updates from $\theta^{(0,r)}$ generated by randomized EM, so $\{\theta^{(t,1)}\}_{t\geq 0}, \{\theta^{(t,2)}\}_{t\geq 0}, \ldots,$ $\{\theta^{(t,R)}\}_{t\geq 0}$ denote the sequences generated by R *parallel* randomized EM algorithms. Notice that we run the sequences simultaneously via R parallel implementations. Therefore we obtain R parameter estimates in the t^{th} iteration, $\theta^{(t,1)}, \theta^{(t,2)}, \ldots, \theta^{(t,R)}$. This will be the parent population for the next generation.

Using this parent population, we create offspring via crossover and mutation. Crossover can be thought of as a swapping of genes. In our context, the information from two parameter estimates are being swapped. Consider the following simple example for illustration. Let θ^a and θ^b denote two elements

in the parent population of the form

$$\theta^a = (\theta_{a1}, \theta_{a2}, \theta_{a3}, \theta_{a4}, \theta_{a5}),$$
$$\theta^b = (\theta_{b1}, \theta_{b2}, \theta_{b3}, \theta_{b4}, \theta_{b5}),$$

and let c be a crossover point chosen randomly out of the set $\{1, 2, 3, 4, 5\}$. For instance, if $c = 2$ then the crossover of θ^a and θ^b are given by θ^a_c and θ^b_c where

$$\theta^a_c = (\theta_{a1}, \theta_{a2}, \theta_{b3}, \theta_{b4}, \theta_{b5}),$$
$$\theta^b_c = (\theta_{b1}, \theta_{b2}, \theta_{a3}, \theta_{a4}, \theta_{a5}).$$

In other words, the last three components of θ^a have been swapped for those of θ^b (and vice versa). Crossover can also be performed in other ways. For instance, rather than swapping entire component-chains, parameter-components can also be swapped one-by-one, each time using a coin-flip to decide whether component j of θ^a should be replaced by the corresponding component of θ^b.

Crossover is typically augmented by mutation, which prevents premature convergence and ensures a wider exploration of the parameter space. Mutation can be thought of as inflicting random shocks into the gene sequence. In the above example, θ^a can be mutated by first randomly selecting a mutation component $m \in \{1, 2, 3, 4, 5\}$, and then replacing component number m of θ^a with a randomly chosen value, say, θ^*. For instance, for $m = 4$ we would get the mutation of θ^a as

$$\theta^a_m = (\theta_{a1}, \theta_{a2}, \theta_{a3}, \theta^*, \theta_{a5}). \tag{18}$$

Mutation is typically not performed in every generation and on every offspring, but rather only occasionally and with a very small mutation probability p_m, e.g. $p_m = 0.1$.

Our genetic algorithm version of EM (GAEM) now proceeds as follows. After creating the parent population $\theta^{(t,1)}, \theta^{(t,2)}, \ldots, \theta^{(t,R)}$, we select pairs of parents $\{\theta^{(t,1)}, \theta^{(t,2)}\}$, $\{\theta^{(t,3)}, \theta^{(t,4)}\}$, $\ldots$, $\{\theta^{(t,R-1)}, \theta^{(t,R)}\}$ (not necessarily adjacent as in this case), and apply crossover and mutation to each pair. To that end, we select a random crossover point c (which could be different for the different pairs) and swap the pair's parameter components to obtain a crossover $\{\theta^{(t,1)}_c, \theta^{(t,2)}_c\}$ for the first pair (and similarly for the other pairs). For each element of a pair, we perform mutation with a small mutation probability, $p_m = 0.1$ to obtain pairs of offspring $\{\theta^{(t,1)}_{c,m}, \theta^{(t,2)}_{c,m}\}$, $\{\theta^{(t,3)}_{c,m}, \theta^{(t,4)}_{c,m}\}$, $\ldots$, $\{\theta^{(t,R-1)}_{c,m}, \theta^{(t,R)}_{c,m}\}$.

After creating a set of suitable offspring, the next step in the genetic algorithm is to evaluate an offspring's fitness and select the offspring with the highest fitness for further reproduction. We evaluate an offspring's fitness by appealing to the likelihood-ascent property. Recall that by (14), if

$$Q(\theta^{(t,1)}_{c,m} | \theta^{(t-1,1)}) > Q(\theta^{(t,1)} | \theta^{(t-1,1)}), \tag{19}$$

then $\theta_{c,m}^{(t,1)}$ improves upon $\theta^{(t,1)}$ and we replace $\theta^{(t,1)}$ by $\theta_{c,m}^{(t,1)}$. We also check equation (19) for the other offspring from the same pair, $\theta_{c,m}^{(t,2)}$, and exchange it for $\theta^{(t,1)}$ if it yields an even larger improvement. We repeat this for the other pairs. After checking the fitness of all offspring (and exchanging parameter estimates for fitter offspring), we continue with the next GAEM iteration, $(t) \rightarrow (t+1)$.

Notice that using (19) to evaluate an offspring's fitness has several advantages. First, by appealing to (19), we preserve the algorithm's basic likelihood ascent property. This means that despite running R parallel chains, the overall algorithm still falls within the class of generalized EM (GEM) algorithms. Moreover, notice that the Q-function is calculated *automatically* throughout every iteration of the EM algorithm, and thus does not pose extra computational effort. This is especially important in complex data/model situations based on huge databases and/or which require heavy simulations. In fact, [25] shows that for the mixture-model (6) with normal mixture densities, additional simplifications, and thus extra computational advantages, are possible. Using (19) has further advantages, especially within the context of the randomized EM implementation. Indeed, [8] provide evidence that evaluating parameter updates with respect to the difference in Q-functions in (19) can lead to a de-facto increase in the algorithm's convergence in the sense that it reaches the optimum faster.

4.2 Experiment

We apply our method to the clustering of functional data originating from an online auction database. By functional data, we mean data that arrives in the form of curves or shapes [46]. In our data, each functional observation represents the price formation process in an online auction [21].

Our database consists of 55,785 price curves from a variety of eBay auction categories ranging from golf balls and pens, to jewelry and automotive. The goal is to cluster the price curves to find patterns of similar price formation processes. To that end, we frame the curve-clustering problem within a finite mixture model context by finding a finite-dimensional representation of the infinite curve and by operating on the coefficients of this representation rather than on the original curve itself [25]. This results in 55,785 20-dimensional data-vectors.

4.3 Results

We set up our simulations as follows. First, we apply the deterministic EM algorithm 100 times from 100 randomly chosen starting values. Then, we compare the result with that of one run of GAEM. The EM algorithm in its basic form does not select the best number g of cluster, so we apply this process to different values of g, ranging from 2 to 10. Notice that as the number

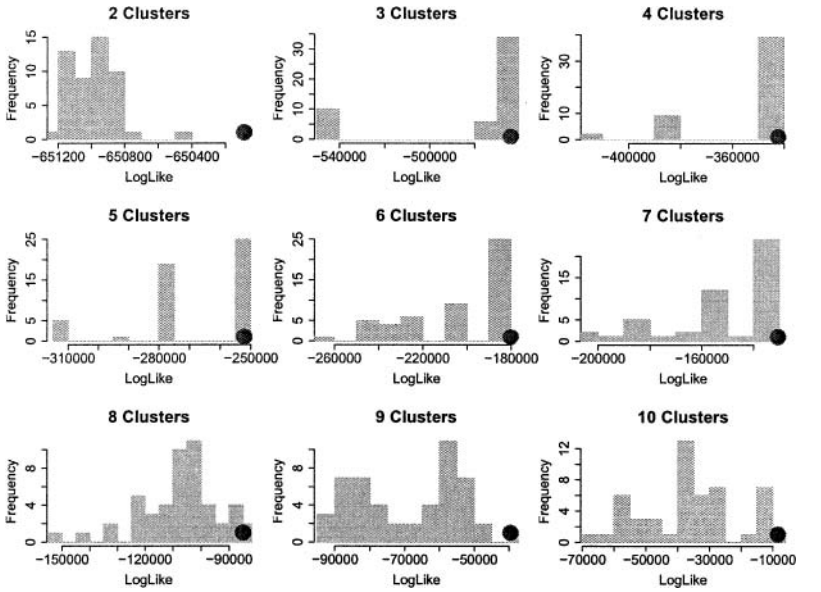


Fig. 3. Performance of GAEM relative to 100 runs of ordinary EM. The histogram shows the distribution of the 100 EM log-likelihood values. The solid dot marks the solution of GAEM.

of clusters increases, the optimization problem also becomes more challenging in that more locally optimal solutions exist.

Figure 3 shows the results. The histogram shows the distribution of the best solution found via deterministic EM. Notice that the solutions vary quite significantly. While for smaller values of g, most solutions lie in close proximity to each other, their variation increases with increasing g. The solid dot signifies the solution found by GAEM. Notice that GAEM consistently finds a solution in the top percentile of the best solutions provided by 100 runs of deterministic EM.

5 Conclusion

The numerical results show that a wedding between the EM algorithm and the principles of global search can indeed lead to fruitful progress. However, further research efforts seem necessary. Indeed, the GAEM algorithm presented in this paper is a relatively straightforward application of the ideas of evolutionary computation to that of the EM algorithm. It is likely that further research will yield more efficient and powerful variants. Moreover, we

pointed out at the beginning of this paper that alternative global optimization procedures such as the cross-entropy method promise even better global optimization performance. However, we also pointed out that these methods can encounter challenges, for instance in mixture-models, which are overcome rather elegantly within the EM framework. It seems that a closer look at how some of these methods can be combined could results in tremendous advancements within the literature on the EM algorithm but also for the area of global optimization.

References

1. S. Amari. Information geometry of the EM and EM algorithms for neural networks. *Neural Networks*, 8:1379–1408, 1995.
2. J. Booth and J. Hobert. Standard errors of prediction in generalized linear mixed models. *Journal of the American Statistical Association*, 93:262–272, 1998.
3. J. Booth and J. Hobert. Maximizing generalized linear mixed model likelihoods with an automated Monte Carlo EM algorithm. *Journal of the Royal Statistical Society B*, 61:265–285, 1999.
4. J. Booth, J. Hobert and W. Jank. A survey of Monte Carlo algorithms for maximizing the likelihood of a two-stage hierarchical model. *Statistical Modelling*, 1:333–349, 2001.
5. Z. Botev and D. Kroese. Global likelihood optimization via the cross-entropy method with an application to mixture models. In *Proceedings of the 2004 Winter Simulation Conference*, pages 529–535. IEEE Press, 2004.
6. R. Boyles. On the convergence of the EM algorithm. *Journal of the Royal Statistical Society B*, 45:47–50, 1983.
7. S. Caffo, J. Booth and A. Davison. Empirical sup rejection sampling. *Biometrika*, 89:745–754, 2002.
8. B. Caffo, W. Jank and G. Jones. Ascent-Based Monte Carlo EM. *Journal of the Royal Statistical Society, Series B*, 67:235–252, 2005.
9. G. Celeux and J. Diebolt. A stochastic approximation type EM algorithm for the mixture problem. *Stochastics and Stochastics Reports*, 41:127–146, 1992.
10. K. Chan and J. Ledolter. Monte Carlo EM estimation for time series models involving counts. *Journal of the American Statistical Association*, 90:242–252, 1995.
11. B. Delyon, M. Lavielle and E. Moulines. Convergence of a stochastic approximation version of the EM algorithm. *The Annals of Statistics*, 27:94–128, 1999.
12. A. Dempster, N. Laird and D. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B*, 39:1–22, 1977.
13. M. Feder and E. Weinstein. Parameter estimation of superimposed signals using the EM algorithm. *Acoustics, Speech, and Signal Processing*, 36:477–489, 1988.
14. G. Fort and E. Moulines. Convergence of the Monte Carlo expectation maximization for curved exponential families. *The Annals of Statistics*, 31:1220–1259, 2003.
15. M. Gu and S. Li. A stochastic approximation algorithm for maximum likelihood estimation with incomplete data. *Canadian Journal of Statistics*, 26:567–582, 1998.

16. M. Gu and H.-T. Zhu. Maximum likelihood estimation for spatial models by Markov chain Monte Carlo stochastic approximation. *Journal of the Royal Statistical Society B*, 63:339–355, 2001.
17. J. Heath, M. Fu and W. Jank. Global optimization with MRAS, Cross Entropy and the EM algorithm. Working Paper, Smith School of Business, University of Maryland, 2006.
18. J. Holland. *Adaptation in Natural and Artificial Systems.* The University of Michigan Press, Ann Arbor, MI, 1975.
19. M. Jamshidian and R. Jennrich. Conjugate gradient acceleration of the EM algorithm. *Journal of the American Statistical Association*, 88:221–228, 1993.
20. M. Jamshidian and R. Jennrich. Acceleration of the EM algorithm by using Quasi-Newton methods. *Journal of the Royal Statistical Society B*, 59:569–587, 1997.
21. W. Jank and G. Shmueli. Dynamic profiling of online auctions using curve clustering. Technical report, Smith School of Business, University of Maryland, 2003.
22. W. Jank and J. Booth. Efficiency of Monte Carlo EM and simulated maximum likelihood in two-stage hierarchical models. *Journal of Computational and Graphical Statistics*, in print, 2002.
23. W. Jank. Implementing and diagnosing the stochastic approximation EM algorithm. Technical report, University of Maryland, 2004.
24. W. Jank. Quasi-Monte Carlo Sampling to Improve the Efficiency of Monte Carlo EM. *Computational Statistics and Data Analysis*, 48:685–701, 2004.
25. W. Jank. Ascent EM for fast and global model-based clustering: An application to curve-clustering of online auctions. Technical report, University of Maryland, 2005.
26. M. Jordan and R. Jacobs. Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6:181–214, 1994.
27. K. Lange. A gradient algorithm locally equivalent to the EM algorithm. *Journal of the Royal Statistical Society B*, 57:425–437, 1995.
28. S. Lauritzen. The EM algorithm for graphical association models with missing data. *Computational Statistics and Data Analysis*, 19:191–201, 1995.
29. P. L'Ecuyer and C. Lemieux. Recent advances in randomized Quasi-Monte Carlo Methods. In M Dror, P L'Ecuyer, and F Szidarovszki, editors, *Modeling Uncertainty: An Examination of Stochastic Theory, Methods, and Applications*, pages 419–474. Kluwer Academic Publishers, 2002.
30. C. Lemieux and P. L'Ecuyer. Efficiency improvement by lattice rules for pricing asian options. In *Proceedings of the 1998 Winter Simulation Conference*, pages 579–586. IEEE Press, 1998.
31. R. Levine and G. Casella. Implementations of the Monte Carlo EM algorithm. *Journal of Computational and Graphical Statistics*, 10:422–439, 2001.
32. R. Levine and J. Fan. An automated (Markov Chain) Monte Carlo EM algorithm. *Journal of Statistical Computation and Simulation*, 74:349–359, 2004.
33. C. Liu and D. Rubin. The ECME algorithm: A simple extension of EM and ECM with faster monotone convergence. *Biometrika*, 81:633–648, 1994.
34. C. McCulloch. Maximum likelihood algorithms for generalized linear mixed models. *Journal of the American Statistical Association*, 92:162–170, 1997.
35. C. McCulloch and S. Searle. *Generalized, Linear and Mixed Models.* Wiley, New-York, 2001.

36. G. McLachlan and D. Peel. *Finite Mixture Models.* Wiley, New York, 2000.
37. X.-L. Meng. On the rate of convergence of the ECM algorithm. *The Annals of Statistics*, 22:326–339, 1994.
38. X.-L. Meng and D. Rubin. Maximum likelihood estimation via the ECM algorithm: A general framework. *Biometrika*, 80:267–278, 1993.
39. R. Neal and G. Hinton. A view of EM that justifies incremental, sparse and other variants. In M Jordan, editor, *Learning in Graphical Models*, pages 355–371, 1998.
40. S.-K. Ng and G. McLachlan. On some variants of the EM Algorithm for fitting finite mixture models. *Australian Journal of Statistics*, 32:143–161, 2003.
41. S.-K. Ng and G. McLachlan. On the choice of the number of blocks with the incremental EM algorithm for the fitting of normal mixtures. *Statistics and Computing*, 13:45–55, 2003.
42. K. Nigam, A. Mccallum, S. Thrun and T. Mitchell. Text classification from labeled and unlabeled documents using EM. *Machine Learning*, 39:103–134, 2000.
43. A. Owen and S. Tribble. A Quasi-Monte Carlo Metropolis algorithm. *Proceedings of the National Academy of Sciences*, 102:8844–8849, 2005.
44. B. Polyak and A. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal of Control and Optimization*, 30:838–855, 1992.
45. F. Quintana, J. Liu and G. delPino. Monte Carlo EM with importance reweighting and its applications in random effects models. *Computational Statistics and Data Analysis*, 29:429–444, 1999.
46. J. Ramsay and B. Silverman. *Functional Data Analysis.* Springer-Verlag, New York, 1997.
47. H. Robbins and S. Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, 22:400–407, 1951.
48. D. Rubin. EM and beyond. *Psychometrika*, 56:241–254, 1991.
49. B. Thiesson, C. Meek and D. Heckerman. Accelerating EM for large databases. *Machine Learning*, 45:279–299, 2001.
50. Y. Tu, M. Ball and W. Jank. Estimating flight departure delay distributions: A statistical approach with long-term trend and short-term pattern. Technical report, University of Maryland, 2005.
51. X. Wang and F. Hickernell. Randomized Halton sequences. *Mathematical and Computer Modelling*, 32:887–899, 2000.
52. G. Wei and M. Tanner. A Monte Carlo implementation of the EM algorithm and the poor man's data augmentation algorithms. *Journal of the American Statistical Association*, 85:699–704, 1990.
53. C. Wu. On the convergence properties of the EM algorithm. *The Annals of Statistics*, 11:95–103, 1983.

Recovering Circles and Spheres from Point Data

Christoph Witzgall, Geraldine S. Cheok, and Anthony J. Kearsley

National Institute of Standards and Technology
Gaithersburg, MD 20899
witzgall@nist.gov, cheok@nist.gov, ajk@nist.gov

Summary. Methods for fitting circles and spheres to point sets are discussed. LADAR (LAser Detection And Ranging) scanners are capable of generating "point clouds" containing the (x, y, z) coordinates of up to several millions of points reflecting the laser signals. In particular, coordinates collected off objects such as spheres may then be used to model these objects by fitting procedures. Fitting amounts to minimizing what is called here a "gauge function," which quantifies the quality of a particular fit. This work analyzes and experimentally examines the impact of the choice of three such gauge functions. One of the resulting methods, termed here as "algebraic" fitting, formulates the minimization problem as a regression. The second, referred to as "geometric" fitting, minimizes the sum of squares of the Euclidean distances of the data points from the tentative sphere. This method, based on orthogonal distance minimization, is most highly regarded and widely used. The third method represents a novel way of fitting. It is based on the directions in which the individual data points have been acquired.

Key words: algebraic fitting; circles; coordinate search; directional fitting; geometric fitting; LADAR; optimization; quasi-Newton; registration; spheres.

1 Introduction

In 1997, an article by Rorres and Romano [19] addressing an archaeological problem caught the attention of Saul Gass. In the Greek city of Corinth, the circular starting line for an ancient (circa 550 B.C.) track for foot races had been found. The other end of such a racetrack had to be marked by a turning pole [18]. The circular starting line was chosen, presumably, to equalize the distances between starting positions and the turning pole at the far end. The location of the turning pole could thus be inferred as the center of the circle passing through the starting line. This is precisely what Rorres and Romano did: they surveyed points on the starting line and fit a circle to these points.

Saul was intrigued by this problem of recovering a circle from a set of points as he had encountered circle and sphere fitting problems earlier in con-

nection with Coordinate Measuring Machines (CMMs), which are crucial to today's precision manufacturing. Their metrology is a major concern at the National Institute of Standards and Technology (NIST) [17], with which he has been associated for nearly three decades as an academic advisor. The computer-guided probe of a CMM touches an object and records to a high level of accuracy the (x, y, z) coordinates of the point on the object where contact was made. A collection of such measured points then permits an accurate assessment of the shape of the object, say, the roundness of a sphere, or its dimensions, say, the radius of a sphere. Captivated by the obvious parallels between such geometric problems and those encountered in classical Operations Research [5, 6], he explored the use of linear programming in this context [9]. He also dug deeper into the racetrack problem, using several different methods for fitting the data and, in particular, compared the widths of "annuli," the areas between two concentric circles containing all data points. He was able to achieve tighter results than those reported by Rorres and Romano. A beautifully crafted unpublished manuscript [10] summarizes this work. It is also telling that his emphasis was not so much on "how" to compute, but rather on "what" to compute, in other words, the task of "modeling" so as to capture a particular aspect of reality.

Our work aspires to follow his example. It was prompted by the rapid growth of 3D imaging technology and its applications, and the corresponding need for metrological analysis. 3D imaging systems include laser scanners and optical range cameras. The former category covers LADARs (LAser Detection And Ranging) or laser radars. Similarly to a CMM, a LADAR also determines 3D coordinates of points on an object, but does so by sending out a laser signal and analyzing its reflection back to the instrument as indicated in Figure 1. Also, a LADAR is typically a scanning device that can obtain millions of measurements in a very short time, resulting in large "point clouds" of possibly millions of data points. Applications include the monitoring of construction sites, the development of "as built" models of existing structures, mapping, visualization of hidden objects, guiding of unmanned vehicles, just to mention a few.

1.1 LADAR Technology; Point Clouds

The metrology of LADARs is a current research issue at NIST. This work also supports the development of standard test protocols for the performance evaluation of LADARs. Figure 2 shows the indoor, artifact-based LADAR facility at NIST. Outdoor test facilities are planned for instruments with ranges of 100 m and above.

Figure 3 presents a LADAR visualization of a rubble pile. The casual reader may see this picture as just what a photographic camera might produce. But this would miss the point that, once the point cloud has been captured in 3D, it can be displayed as seen from different view points. For instance, the

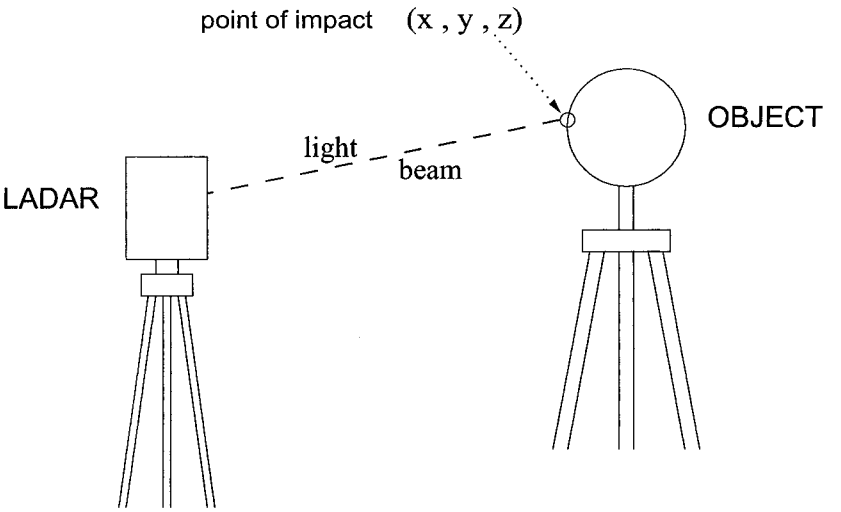


Fig. 1. Schematic of the operation of a LADAR scanner.

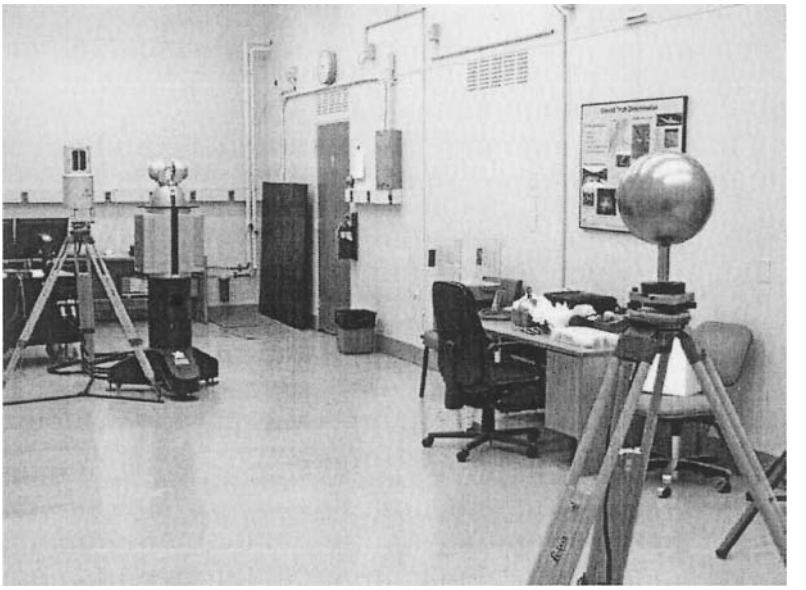


Fig. 2. Indoor artifact-based LADAR facility at NIST .

Disclaimer: Certain products are shown in this photograph. In no case does this imply recommendation or endorsement by NIST, nor does it imply that the products are necessarily the best available for the purpose.



Fig. 3. LADAR scan of a rubble pile

point cloud generated off a sphere is displayed in Figure 4 as seen sideways, that is, perpendicular to the direction of the scan.

The point cloud in Figure 4 consists of actual data points. The superimposed picture of the sphere does not depict the actual sphere but a "fitted" sphere, a sphere that in some sense best represents the point cloud, a sphere that was calculated using a fitting algorithm. The fitted sphere is expected to capture the radius and location of the actual sphere vis-à-vis the point cloud. Fitting the sphere thus provides a method for recovering the actual sphere from which the point cloud had been generated.

The term "recovery" suggests an emphasis on certain applications that are different from the ones pursued when using CMMs. LADARs are more likely to be used for locating, identifying, and recognizing an object. CMMs, on the other hand, emphasize quality control. For spheres, a typical question there would be how small an annulus between two concentric spheres would contain the data points. Also, as Figure 4 shows, the point clouds generated by

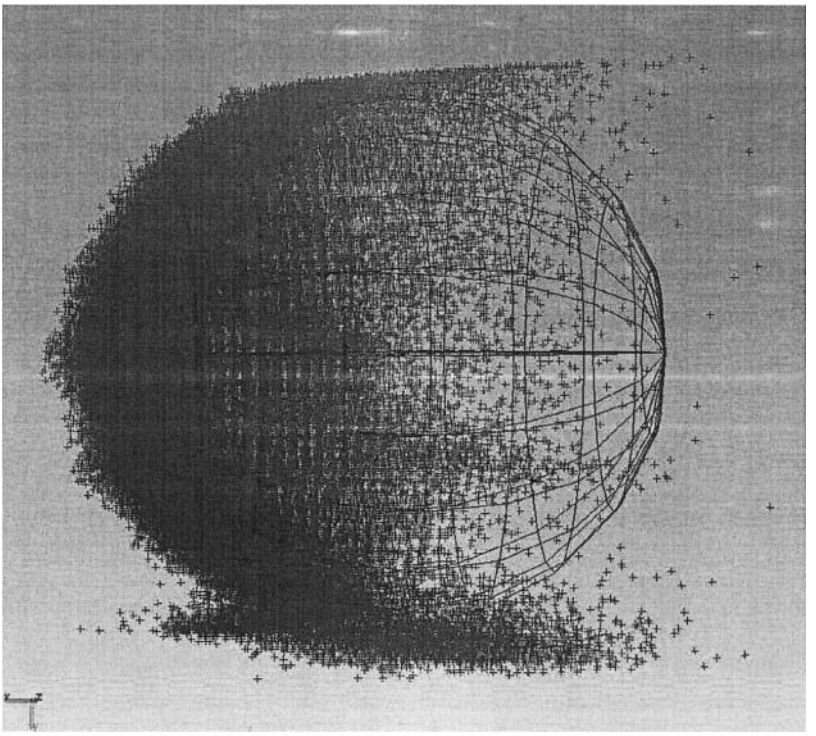


Fig. 4. Point cloud and fitted sphere.

LADARs tend to be big and "noisy," i.e., subject to significant random errors. The data generated by CMMs, on the other hand, tend to be less voluminous and less noisy.

While the emphasis in this work is on spheres, the results will typically hold for circles, too.

1.2 The Fitting Paradigm

The term "fitting" will be used here in a somewhat narrow and explicit sense. This sets it apart from other approaches to recovering scanned objects. Iterative Closest Point (ICP) [3] techniques have also been very effective. Tetrahedralization-based surface generation techniques may well be the way of the future [2].

Fitting requires the definition of a "gauge function," which expresses a measure of deviation of the point cloud as a whole from a putative geometric model such as a sphere. The parameters of the model governing its shape and size are the arguments of the gauge function. They are the variables of

an optimization problem: to minimize the deviation of data points from an object model.

Gauge functions are typically defined in two steps. First, a measure of the individual deviation

$$\Delta_i = \text{deviation}\,((x_i, y_i, z_i),\ \text{model})$$

of each data point from the putative model is defined. Then a norm $\|*\|$ is selected and applied to the vector of all deviations Δ_i in order to arrive at a single gauge value for an overall measure of deviation,

$$\boldsymbol{\Delta} \ = \ \|(\Delta_1\ ,\ \Delta_2\ ,\ \ldots\ \Delta_n)\| \ .$$

The quantity $\boldsymbol{\Delta}$ represents the desired gauge function, which depends on the model parameters as its arguments, since every one of the individual deviations Δ_i depends on those parameters.

The following are the most commonly selected norms:

- the maximum or Chebychev norm (L_∞) : $\max_i |\Delta_i|$,
- the least-squares norm (L_2) : $\sqrt{\sum_i \Delta_i^2}$,
- the sum of absolute values norm (L_1) : $\sum_i |\Delta_i|$.

Directly minimizing the sum of squared deviation is, of course, equivalent to minimizing their L_2 norm. In practice, to avoid the very large numbers that typically arise for large data sets, both the L_2 and the L_1 norms are usually averaged by the number n of data points. The resulting measures are commonly referred to as

$$\text{RMS} \ = \ \sqrt{\frac{\sum_i \Delta_i^2}{n}}\ , \quad \text{ASD} \ = \ \frac{\sum_i |\Delta_i|}{n}\ .$$

The least-squares approach, as in using the RMS, is most commonly used. At NIST, a system for testing least-squares algorithms for a collection of types of objects including spheres and circles has been developed and implemented [20], as has been orthogonal distance regression for specified functional expressions [4]. Applications of the Chebychev norm are given in [1]. Chebychev fitting of circles has been examined by Saul Gass [9, 10] and others [22, 8]. As mentioned before, the interest here is in determining the minimum annulus containing all data points. The authors apply linear programming as well as farthest point Voronoi diagrams. In our work on spheres, the desired end result consists of the "center coordinates" (x_0^*, y_0^*, z_0^*) of a fitted sphere, and perhaps its "radius" r^*. Indeed, when fitting spheres, two different tasks may be encountered:

- fitting a sphere with its radius "free" to be determined;
- fitting a sphere with a specified "fixed" radius.

1.3 Layout of the Article

The remainder of the paper is organized as follows. In Section 2, we describe physical evidence on fitting spheres collected by the second author. In Section 3, two standard methods are discussed: "algebraic fitting" (Section 3.1) and "geometric fitting" (Section 3.2). Section 4 is devoted to the geometry (Section 4.1) and the process of scan-directed fitting (Section 4.2) based on an algorithm developed by the third author.

2 Results of Experiments

At NIST, considerable experience in LADAR scanning and fitting has been gathered. Key issues are applications to registration, volume determination, and object identification.

2.1 Locating I-Beams

The following demonstration experiment [12] was designed to demonstrate the feasibility of automated placing and pick-up of an I-beam by the computer-guided crane developed at NIST. The I-beam, residing on the laboratory floor, was scanned for location and orientation, with the data in the LADAR's coordinate system. A LADAR scanner was used to determine the pose (location and orientation) and the type of an I-beam residing on the floor of a laboratory at NIST. The idea was to scan the I-beam in order to determine shape and pose within the coordinate system of the LADAR instrument.

This instrument's coordinate system then had to be related to the coordinate system of the crane, a process generally called "registration". To this end, three "target" spheres, "A","B","C", were placed in the vicinity of the I-beam

The centers of the spheres were predetermined in the coordinate system of the crane. The LADAR scan covered these spheres along with the I-beam, and the fitting process yielded center coordinates in the instrument system. Thus, there were three target locations, each of which with coordinates known in both systems. The algorithm "Procrustes" [21] [14] was employed, which combines a translation with a rotation in order to transform one system into the other, matching the coordinates at each target point as well as possible. More precisely, this transformation is chosen so as to minimize the sum of the squares of the resulting coordinate differences at each target point. It is clear that the accuracy of the fitting algorithm as applied to the target spheres is the key to a correct registration.

This successful demonstration also provided an opportunity to experiment with sphere fitting. As described in [12], several fitting algorithms were implemented and tried in addition to a commercial software package. The radius of the target spheres was specified by the manufacturer to be 76.2 mm (3

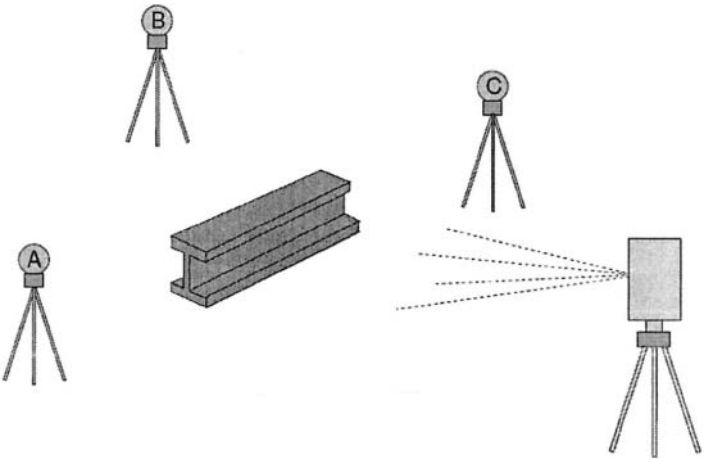


Fig. 5. Determining location and orientation of an I-beam by LADAR.

in). Thus, the fixed radius option for sphere fitting was used for the actual demonstration.

Determining the accuracy rather than the precision of such LADAR measurements would require the availability of an independent measurement system at a higher level of accuracy, which was not available at the time. Thus it was not possible to ascertain the actual centers of the target spheres with sufficient accuracy to check against the centers derived from fitting the LADAR data. The radii of the target spheres, however, were presumably known, and it could be determined how well the radii were reproduced if the free radius option of the sphere fitting algorithms were to be used. The results were disconcerting. The radii were generally underestimated. In particular, applying the commercial fitting package to target sphere "C" yielded for $n = 90$ the average radius,

$$r_{average} = 69 \text{ mm}, \; r_{actual} = 76 \text{ mm}, \; st.dev. = 3 \text{ mm},$$

a discrepancy of about 10 %. This result raised a "red flag". Could the center determination be trusted if established fitting methods produced systematic errors in the measurement of the radius? An effort was therefore started to find the reasons for such errors.

At first, suspicion centered on the quality of the target spheres which were made of styrofoam purchased at a local crafts store. It was thought that this material might permit some penetration by the scan beam, or perhaps the dimensions of the spheres were not quite accurate. These considerations led to the fabrication of a machined aluminum sphere. However, the same

discrepancies were encountered when this sphere was scanned and fitted, as will be seen in Section 2.2.

One other possible explanation concerned the LADAR instrument. Perhaps the distribution of the scan errors was not symmetric. In other words, there may have been more undershoots than overshoots, or vice versa. And finally, the instrument position with respect to the sphere may possibly matter. To check for this latter possibility, the experiment reported in Section 2.2 below was conducted.

2.2 An Additional Experiment

In this experiment, a data set was collected, and reduced to avoid boundary effects, off an aluminum sphere machined to a radius of 101.6 mm (4 in). This data set is displayed in Figure 6 together with two subsets, an upper and a lower subset into which the full set had been split as shown in Figure 7. The results of applying the commercial fitting package to these three data sets are displayed in Tables 1 and 2.

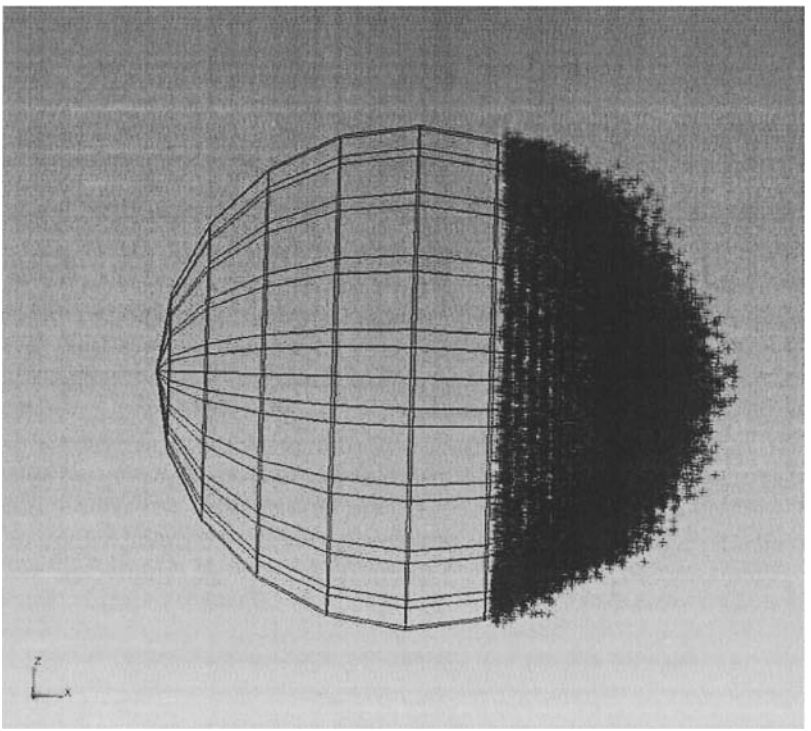


Fig. 6. Full "hemispherical" data set from aluminum sphere.

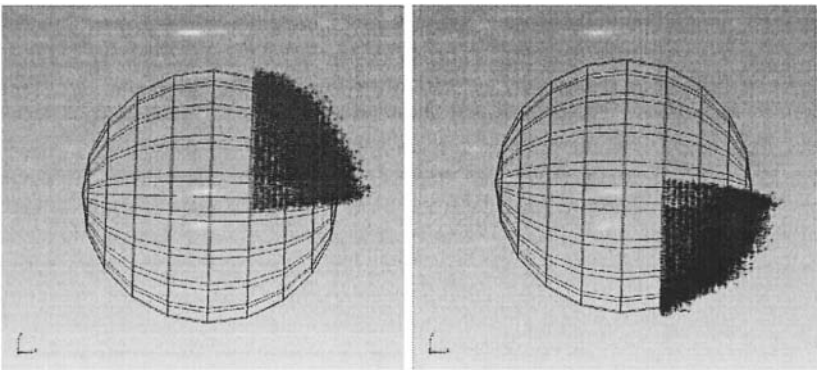


Fig. 7. Upper and lower portions of the hemispherical data set.

Table 1. Results of the experiment; variable radius.

	x	y	z	r
Full	-6254.99	-196.51	-78.85	98.41
Upper	-6258.27	-196.37	-83.02	102.36
Lower	-6258.61	-196.82	-72.61	103.66

Table 2. Results of the experiment; fixed radius.

	x	y	z	r
Full	-6259.19	-196.58	-78.87	101.6
Upper	-6257.52	-196.36	-82.55	101.6
Lower	-6256.59	-196.77	-73.98	101.6

The first observation concerns the result of fitting the full data set with the free radius option. As in the demonstration reported in Section 2.1, the radius was still underestimated:

$$r_{computed} = 98.41 \text{ mm}, \; r_{actual} = 101.6 \text{ mm},$$

but then it was overestimated for both the upper and the lower portion of the full data set. The next observation concerned the high level of sensitivity in the z-coordinate, which represents vertical elevation. Note that the same sensitivity in the z-coordinate showed up when the known radius of 101.6 mm had been kept fixed. Such variations are at odds with the fact that regions on the sphere are equivalent. Indeed, the upper and the lower data set occupy essentially symmetric positions on the sphere. Yet there is a substantial difference in fitting results.

The upper and the lower subset are, however, in a different position vis-à-vis the LADAR instrument. The angles of incidence certainly differ for these two subsets. This forces the conclusion that the instrument position has to

be taken into account when fitting. A method for this will be presented in Section 4.

3 Algebraic and Geometric Fitting

In a key paper [7], two generic concepts for defining deviations of data points were distinguished, and fitting methods based on those concepts were termed "algebraic" and "geometric," respectively. Algebraic fitting, in a broad sense, is based on an equation describing the object to be fitted, defining the deviation of a data point as the amount by which that equation is violated. In general geometric fitting, that deviation is defined as the orthogonal Euclidean distance to the object. For geometric fitting of functions, the term "orthogonal distance regression" is frequently used.

Specifically in [7], algebraic and geometric fitting methods based on least squares were examined for circles and ellipses. In this section, we will take the same approach to fitting circles and spheres.

3.1 Algebraic Fitting

Let (x_0, y_0, z_0) denote the center coordinates, and $r \geq 0$ the radius of a sphere. Then the following equation

$$(x - x_0)^2 + (y - y_0)^2 + (z - z_0)^2 - r^2 = 0$$

characterizes the points (x, y, z) of the sphere. Substituting

$$x_0 = -\frac{a}{2},\ y_0 = -\frac{b}{2},\ z_0 = -\frac{c}{2},\ r^2 = \frac{a^2 + b^2 + c^2}{4} - d, \tag{1}$$

yields an alternate equation

$$x^2 + y^2 + z^2 + ax + by + cz + d = 0$$

of the above sphere in terms of linear parameters a, b, c, d. Note that the above equation has geometric meaning only if its parameters satisfy the condition

$$\frac{a^2 + b^2 + c^2}{4} - d \geq 0, \tag{2}$$

as otherwise the resulting radius would not be a real number.

The above equations for the sphere suggest the following definition of a deviation from the sphere by a data point (x_i, y_i, z_i):

$$\Delta_i = (x_i - x_0)^2 + (y_i - y_0)^2 + (z_i - z_0)^2 - r^2, \tag{3}$$

and, equivalently,

$$\Delta_i = x_i^2 + y_i^2 + z_i^2 + ax_i + by_i + cz_i + d. \tag{4}$$

The algebraic method to be discussed here takes advantage of the fact that the deviation expression (4) is linear in its parameters. Choosing the least squares norm leads to the gauge function

$$F = \sum_i \left(x_i^2 + y_i^2 + z_i^2 + ax_i + by_i + cz_i + d\right)^2, \tag{5}$$

the minimization of which amounts to a straightforward linear regression,

$$-x_i^2 - y_i^2 - z_i^2 \sim ax_i + by_i + cz_i + d, \tag{6}$$

with an optimal solution,

$$a^*, b^*, c^*, d^*,$$

which always exists and is unique as long as atleast four data points do not lie in a single plane.

Also, the optimal parameters satisfy the condition (2). Indeed, it is well known that the data averages satisfy regression equations such as (6) exactly if the optimal parameters are inserted:

$$\frac{1}{n}\sum_i(-x_i^2 - y_i^2 - z_i^2) = a^*\frac{1}{n}\sum_i x_i + b^*\frac{1}{n}\sum_i y_i + c^*\frac{1}{n}\sum_i z_i + d^*.$$

Then

$$\begin{aligned}
&\frac{(a^*)^2 + (b^*)^2 + (c^*)^2}{4} - d^* \\
&= \frac{1}{n}\sum_i \left(x_i^2 + y_i^2 + z_i^2 + a^*x_i + b^*y_i + c^*z_i + \frac{(a^*)^2 + (b^*)^2 + (c^*)^2}{4}\right) \\
&= \frac{1}{n}\sum_i \left((x_i + \frac{a^*}{2})^2 + (y_i + \frac{b^*}{2})^2 + (z_i + \frac{c^*}{2})^2\right) \;\geq\; 0.
\end{aligned}$$

Thus, for any data set containing at least four points not in a single plane, the regression (6) yields a unique result that, moreover, represents a real sphere.

By (1), the above derivation also yields an explicit expression of the optimal radius in terms of the optimal center coordinates:

$$(r^*)^2 = \frac{\sum_i \left((x_i - x_0^*)^2 + (y_i - y_0^*)^2 + (z_i - z_0^*)^2\right)}{n}. \tag{7}$$

The expressions (3) and (4) represent the same deviation values and lead to the same gauge quantity, provided the same norm is used to combine the respective individual deviations, in this case, the least squares norm, However, the two resulting gauge functions differ in their parameters. As a positive

definite quadratic function, the former gauge function F is convex, whereas the latter gauge function

$$G = \sum_i \left((x_i - x_0)^2 + (y_i - y_0)^2 + (z_i - z_0)^2 - r^2\right)^2$$

is not. If the value of the radius has been prespecified, then this value needs to be accessible as one of the parameters, as in the above gauge function G, but not in gauge function F. As a result, regression will not work for the fixed radius option. Furthermore, there will always be minima if the radius is fixed. They may be, however, no longer unique as an example in Section 3.3 shows.

3.2 Geometric (Orthogonal) Fitting

Here, the actual Euclidean distance of a data point (x_i, y_i, z_i) from the sphere given by its center coordinates (x_0, y_0, z_0) and its radius r is specified as the deviation of this data point from this sphere:

$$\Delta_i = \sqrt{(x_i - x_0)^2 + (y_i - y_0)^2 + (z_i - z_0)^2} - r. \tag{8}$$

Following least squares, the gauge function

$$H = \sum_i \left(\sqrt{(x_i - x_0)^2 + (y_i - y_0)^2 + (z_i - z_0)^2} - r\right)^2 \tag{9}$$

characterizes the particular "geometric fitting" method examined in this paper.

By comparing results with a geometric fitting algorithm implemented by the authors, it was possible to ascertain that the commercial fitting package employed in the experiments described in Section 2 uses the geometric fitting method. In what follows, we will refer simply to "geometric fitting" regardless of which implementation has been used.

A comparative analysis of algebraic and geometric fitting has been provided by [16]. The difference in the performance of these methods is also illustrated by the example in Figure 8. Here, a LADAR scan of the sphere discussed in Section 2.1 has been fitted algebraically, and the algebraic deviations defined by (3) or (4) — the algebraic errors so to speak — have been plotted in the sequence in which the data points appear in the point cloud. It is seen that these algebraic errors are closely gathered around the zero horizontal, which is not surprising, since that is what was minimized. However, the geometric deviations defined by (8) are also displayed, exhibiting uncomfortably large swings. Both kinds of deviations show definite patterns such as the scalloped pattern of the algebraic deviations. The reasons for these patterns are not understood.

Moreover, the radius of that sphere determined by algebraic fitting falls even more short of the actual value than the one determined by geometric fitting:

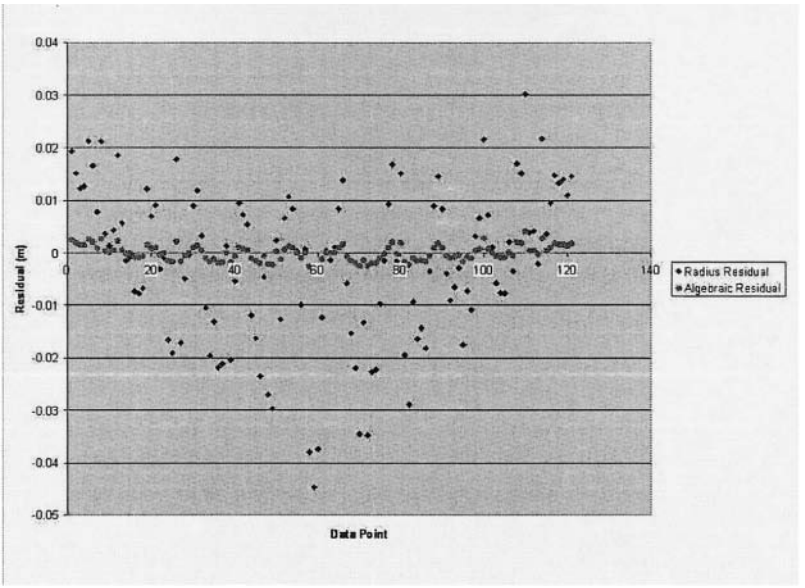


Fig. 8. Deviation display: geometric vs. algebraic fitting.

$$r_{algebraic} = 54 \text{ mm}, \; r_{geometric} = 69 \text{ mm}, \; r_{actual} = 76 \text{ mm}.$$

Experiences similar to the above, coupled with the strong intuitive appeal of orthogonal distance deviations and the general acceptance of the least squares approach, have made geometric fitting the main choice of the metrology community.

Contrary to algebraic fitting, the geometric fitting may not always have a solution. The reason is that an orthogonally best fitted plane, — in a sense a sphere with infinitely large radius —, competes with actual spheres. In such cases, the geometric fitting procedure may not converge as the radius grows ever larger. Furthermore, even if a solution exists, the solution may not be unique. Corresponding examples in the case of circles are given in Section 3.3. In addition, the gauge function H defined by Equation (9) is not differentiable in its parameters. The function displays an upward pointing cusp, wherever the center coordinates coincide with those of a data point:

$$(x_0, y_0, z_0) = (x_i, y_i, z_i).$$

3.3 Examples of Circles

Again, the following examples for circles are indicative of phenomena that pertain to spheres, as well.

The first extremely simple example shows that any fitting method with fixed radius may have different optimal solutions for geometric fitting.

Example A:

$$x_1 = -1, y_1 = 0,$$
$$x_2 = +1, y_2 = 0.$$

The second example admits four different optimal solutions in symmetric positions.

Example B:

$$\begin{aligned} x_1 &= 0, y_1 = 0, \\ x_2 &= +10, y_2 = +10, \\ x_3 &= -10, y_3 = +10, \\ x_4 &= -10, y_4 = -10, \\ x_5 &= +10, y_5 = -10. \end{aligned}$$

The center of one optimal circle is at $x_0^* = 3.892718$, $y_0^* = 0.0$, and its radius is $r^* = 12.312514$. The remaining three optimal circles are in symmetric position rotated by multiples of right angles around the origin. Algebraic fitting, on the other hand, has a unique solution, which for reasons of symmetry is centered at the origin. The optimal radius can then be calculated using (7): $r = \sqrt{160} = 12.64911$.

In order to establish optimality, the gradient and the Hessian for the parameters of the circle to the right were also computed. Up to factors of two, the gradient components compute to zero within seven digits after the decimal point, and the following Hessian results:

$$\begin{vmatrix} +0.5957504 & 0 & -0.31615945 \\ 0 & +0.06370123 & 0 \\ -0.31615945 & 0 & 1 \end{vmatrix}.$$

Since the eigenvalues of the Hessian, 0.0637, 0.4226, 1.1731, are all positive, local optimality follows. No other local optima have been found. The calculations were done with in-house geometric fitting software. The commercial fitting package has also been used, but produced a saddle point instead of a local minimum.

The third example is a data set without an optimal solution when geometric fitting is attempted. The $x-$axis is a limit solution.

Example C:

$$\begin{aligned} x_1 &= +10, y_1 = +1, \\ x_2 &= +1, y_2 = 0, \\ x_3 &= -10, y_3 = +1, \\ x_4 &= -10, y_4 = -1, \\ x_5 &= -1, y_5 = 0, \\ x_6 &= +10, y_6 = -1. \end{aligned}$$

The claim that there are no finite local minima rests on use of the in-house software, which established that absence within a very large radius around the origin.

4 Fitting in Scan Direction

In what follows, we assume that the data points are acquired in scan direction, i.e., the data point (x_i, y_i, z_i) and its intended impact point on a sphere are found on a ray, referred to as "scan ray," emanating from the origin $(0, 0, 0)$, so that the intended target point is at the intersection of the scan ray with the sphere. Of course, the scan ray may not intersect the sphere, in which case we declare a "miss;" otherwise, a "hit." In reality, one might argue, every data point would arise from a hit, as otherwise it would not have gotten into the data set. But as long as a sphere is not in best-fit position, misses are to be expected.

4.1 Geometry of Directed Deviations

We aim to determine the intersection, closest to the origin, of a sphere by a ray from the origin through a given data point (x_i, y_i, z_i). While this just amounts to solving a quadratic equation, there is also a straightforward geometric solution.

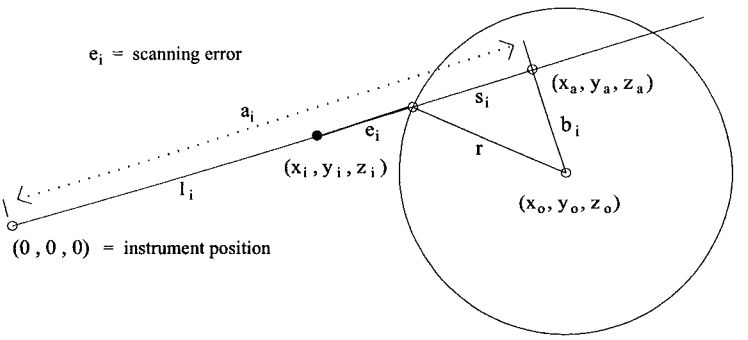


Fig. 9. Geometry of deviations in scan direction.

Consider the plane spanned by the ray and the center (x_0, y_0, z_0) of the sphere (see Figure 9). The direction of the ray is given by the direction cosines,

$$\xi_i = \frac{x_i}{\ell_i},\ \eta_i = \frac{y_i}{\ell_i},\ \zeta = \frac{z_i}{\ell_i},$$

where

$$\ell_i = \sqrt{x_i^2 + y_i^2 + z_i^2} > 0 \tag{10}$$

denotes the distance of the data point from the origin. Of interest is the orthogonal projection of the center of the sphere into the ray:

$$a_i(\xi_i, \eta_i, \zeta_i). \tag{11}$$

Here a_i denotes the distance of the projection point (11) from the origin. Using the orthogonality of the ray direction (ξ_i, η_i, ζ_i) to the difference vector

$$(a_i\xi_i, a_i\eta_i, a_i\zeta_i) - (x_0, y_0, z_0), \tag{12}$$

we find the expression

$$a_i = x_i\xi_i + y_i\eta_i + z_i\zeta_i$$

for the distance a_i.

Next we introduce the length of the difference vector (12), i.e., the distance between the sphere center and its projection into the scan ray. We have

$$b_i^2 = x_0^2 + y_0^2 + z_0^2 - a_i^2 \ \geq \ 0, \tag{13}$$

by the Pythagorean theorem. Comparison against the radius

$$b_i \leq r, \tag{14}$$

yields the condition for the scan ray hitting the sphere. If this condition is met, then the sphere center, the projection point on the ray, and the intersection point on the sphere form a right triangle with sides b_i and s_i, and the hypotenuse r,

$$s_i = \sqrt{r^2 - b_i^2}. \tag{15}$$

It is now possible to express the distance along the scan ray from the origin to the impact point on the sphere. The difference between this distance and the distance ℓ_i of the data point from the origin (10),

$$a_i - s_i - \ell_i = a_i - \ell_i - \sqrt{r^2 - b^2}, \tag{16}$$

represents, in the case of a hit, the overshoot or undershoot of the measurement with respect to the sphere.

In general, the scan ray intersects the sphere at two points, one close to the origin of the instrument, and one on the other side of the sphere. Analogous to the above derivation, we find the expression

$$a_i + s_i - \ell_i = a_i - \ell_i + \sqrt{r^2 - b^2}$$

for the distance of the intersection point farther from the origin. It may be of interest to note that the product of those two distances equals the algebraic orthogonal deviation given by Equation (3)

$$(a_i - s_i - \ell_i)(a_i + s_i - \ell_i) = (x_i - x_0)^2 + (y_i - y_0)^2 + (z_i - z_0)^2 - r^2. \tag{17}$$

Indeed, by (15),

$$(a_i - s_i - \ell_i)(a_i + s_i - \ell_i) = (a_i - \ell_i)^2 - s_i^2$$
$$= (a_i - \ell_i)^2 + b_i^2 - r^2.$$

By the Pythagorean theorem applied to the triangle formed by the data point, the projection point and the sphere center,

$$(a_i - \ell_i)^2 + b_i^2 = (x_i - x_0)^2 + (y_i - y_0)^2 + (z_i - z_0)^2,$$

which establishes (17).

4.2 Geometric Scan-Directed Fitting

The expression (16) of the presumptive measurement error of a data point (x_i, y_i, z_i) was determined under the assumption of a hit. It, therefore, fails to provide a suitable basis for a gauge function. Indeed, any gauge function based solely on hits would yield a minimum of zero for any sphere position that avoids hits altogether. For this reason, we propose to define deviations also for data points whose scan rays are missing the sphere. The version of "directional" fitting reported here thus distinguishes between two kinds of deviations according to the value of b_i as follows:

If the data point causes a hit, that is, if $b_i \leq r$ (14), then following (15),

$$\Delta_i = a_i - s_i - \ell_i = a_i - \ell_i - \sqrt{r^2 - b_i^2}. \tag{18}$$

If the data point causes a miss, that is, if $b_i > r$, then

$$\Delta_i = \sqrt{(x_i - x_0)^2 + (y_0 - y_i)^2 + (z_i - z_0)^2} - r, \tag{19}$$

which in fact represents the orthogonal distance deviation (8).

The least squares approach leads to the gauge function

$$H = \sum_i \Delta_i^2. \tag{20}$$

Minimizing this gauge function may still not provide the desired result, because the so-fitted sphere may still permit some data points to miss, and the deviations assigned to these values may cause distortions. The procedure chosen here is to delete such data points temporarily, and re-minimize. The full point cloud is then again screened for data points that cause misses of the re-minimized sphere, and these points are deleted prior to the next re-minimization. That process of deleting followed by re-minimization is repeated until there are no more misses or the number of misses has stabilized. The goal is to arrive at a sphere fitted to a stable set of data points all providing hits.

This and similar procedures are still subject to experimentation and adjustments. Our first efforts, however, were clearly successful as attested by

Table 3. Results of fitting in the scan direction.

	x	y	z	r
Full	-6258.98	-198.07	-79.18	101.29
Upper	-6259.06	-198.15	-78.90	101.22
Lower	-6259.38	-198.01	-79.12	101.60

the results given in Table 3. They indicate that the abnormalities reported in Section 2 appear indeed to be caused primarily by "modeling error," namely the choice of an unsuitable gauge function for the fitting procedure.

The gauge function (20) is not differentiable, because a perturbation of the sphere parameters may cause the deviation of a data point switch from definition (18) to definition (19), or vice versa, and such transitions are not smooth. The minimization of such a gauge function thus requires an optimizer that does not require differentiability. The method that turned out to be successful was a method based on recent research in the optimization of noisy functions. Loosely speaking, this method is designed for minimizing functions with many non-differentiabilities as long as these are "shallow," i.e., they act as perturbations of an overall differentiable function. This method also permits constraining the minimization to avoid, for instance, values of variables that define spheres that would infringe upon the instrument location.

The algorithm proceeds in two stages. Initially, a quasi-Newton method ("BFGS") is employed to solve the nonlinear programming problem at hand, where gradients of the objective function are calculated using a centered finite-difference approximation with a large finite difference initial step-length [11]. As the algorithm progresses, the finite difference step-length is decreased until its size falls below the square root of machine precision. Subsequently, a simplex-based coordinate search method is employed [15]. This coordinate search method requires no gradient calculation or approximation, and has been applied successfully in the past on difficult non-differentiable constrained optimization problems (e.g. [13]).

5 Concluding Remarks

The analyses in this paper underscore the fact that the outcome of fitting methods strongly depends on the choice of the gauge function that is minimized in order to establish the fit. Three gauge functions were discussed. The associated "algebraic" and "geometric" fitting methods are most commonly used. In a novel third approach, the deviations are measured in the direction of the scan. In all three methods, the deviations were combined into single numbers by invoking the least squares L_2 norm.

The experimental evidence presented here, although limited so far, indicates that for certain applications it may not be enough to base fitting methods on the coordinates of points alone, but that it may be necessary to take into

account the directions in which those points had been acquired. This may require a general revision of currently used approaches to approximation.

Acknowledgments

The authors are indebted to Michael Fu for providing editorial and typesetting assistance.

References

1. G.T. Anthony, B. Bittner, B.P. Butler, M.G. Cox, R. Drieschner, R. Elligsen, A.B. Forbes, H. Gross, S.A. Hannaby, P.M. Harris, and J. Kok. Chebychev reference software for the evaluation of coordinate measuring machine data. *Report EUR 15304 EN*, National Physical Laboratory, Teddington, United Kindom, 1993.
2. J. Bernal. AGGRES: a program for computing power crusts of aggregates. *NISTIR 7306, National Institute Of Standards and Technology*, Gaithersburg, MD, 2006.
3. P.J. Besl and N.D. McKay. A method for registration of 3-D shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14:239–256, 1992.
4. P.T. Boggs and J.E. Rogers. Orthogonal distance regression. *Contempory Math*, 112:183–194, 1990.
5. S-Y. Chou, T.C. Woo, and S.M. Pollock. On characterizing circularity. Dept. of Industrial and Operations Research, University of Michigan, Ann Arbor, 1994.
6. Z. Drezner, S. Steiner, and O. Weselowski. On the circle closest to a set of points. *Computers and Operations Research*, 29:637–650, 2002.
7. W. Gander, G.H. Golub, and R. Strebel. Least-squares fitting of circles and ellipses. *BIT*, 34:558–578, 1994.
8. J. Garcia-Lopez, P. A. Ramos, and J. Snoeyink. Fitting a set of points by a circle. *Discrete and Computational Geometry*, 20:389–402, 1998.
9. S.I. Gass, C. Witzgall, and H.H. Harary. Fitting circles and spheres to coordinate measuring machine data. *International Journal of Flexible Manufacturing Systems*, 10:5–25, 1998.
10. S.I. Gass. Comments on an ancient Greek racecourse: Finding minimum width annuluses. (unpublished manuscript), 1998.
11. P. Gilmore and C.T. Kelley. An implicit filtering algorithm for optimization of functions with many local minima. *SIAM Journal on Optimization*, 5:269–285, 1995.
12. D.E. Gilsinn, C. Witzgall, G.S. Cheok, and A. Lytle. Construction object identification from LADAR scans: an experimental study using I-beams. *NISTIR 7286, National Institute of Standards and Technology*, Gaithersburg, MD, 2005.
13. R. Glowinski and A.J. Kearsley. On the simulation and control of some friction constrained motions. *SIAM Journal on Optimization*, 5:681–694, 1995.
14. E. Golub and C.F. Van Loan. *Matrix Computation.* The Johns Hopkins University Press, Baltimore, MD, 1996.

15. R. M. Lewis and V. Torczon. Pattern search algorithms for bound constrained minimization. *SIAM Journal on Optimization*, 9:1082–1099, 1999.
16. B.A. Jones and R.B. Schnabel. A comparison of two sphere fitting methods. *Proceedings of the Instrumentation and Measurement Technology Subgroup of the IEEE*, 1986.
17. S.D. Phillips, B. Borchardt, and G.C. Gaskey. Measurement uncertainty consideration for coordinate measuring machines. *NISTIR 5170, National Institute of Standards and Technology*, Gaithersburg, MD, 1993.
18. D.G. Romano. Athletes and mathematics in archaic Corinth: the origins of the Greek stadium. *Memoirs of the American Philosophical Society*, 206:xiv,117, 1993.
19. C. Rorres and D.G. Romano. Finding the center of a circular starting line in an ancient Greek stadium. *SIAM Review*, 39:745–754, 1997.
20. C.M. Shakarji. Least-squares fitting algorithms of the NIST algorithmic testing system. *Journal of Research of the National Institute of Standards and Technology*, 103:633–640, 1998.
21. P.H. Schoenemann. A generalized solution of the orthogonal Procrustes problem. *Psychometrica*, 31:1–10, 1966.
22. C.K. Yap. Exact computational geometry and tolerancing metrology. *Snapshots on Computational and Discrete Geometry*, D. Avis and P. Bose, eds., McGill School of Computer Science, 3:34–48, 1994.

Why the New York Yankees Signed Johnny Damon

Lawrence Bodin

Robert H. Smith School of Business
Department of Decision and Information Technologies
University of Maryland
College Park, MD 20742
lbodin@rhmsith.umd.edu

Summary. In this paper, we apply the Analytic Hierarchy Process (AHP) to analyze why the New York Yankees signed Johnny Damon for $52 million over the next four years.

Key words: Analytic Hierarchy Process (AHP); decision-making; baseball.

1 Introduction

In the mid 1990s, Saul Gass and I introduced the teaching of the Analytic Hierarchy Process (AHP) in some MBA and undergraduate classes in the R. H. Smith School of Business at the University of Maryland. Based on our experiences, we wrote two papers on the teaching of the AHP [3], [4]. In [3], we discussed the basic topics covered in our courses and, in [4], we presented some advanced and, hopefully, interesting AHP homework exercises that we developed for these courses. Two fundamental references on AHP are [5] and [6].

For the past 10 years, I have been interested in the use of AHP in making decisions in Major League Baseball (MLB). Eddie Epstein, Director of Baseball Operations, and I developed an AHP model for ranking the players that the San Diego Padres should keep as a prelude to the 1997 Major League Baseball expansion draft [2]. The rankings were completed about two months before the expansion draft was held and were very close to the final ranking that the Padres used in the expansion draft.

In [1], I developed a model for ranking the available free agents after the conclusion of the 1998 baseball season. This paper was presented at the ISAHP meeting in Kobe, Japan in 1999. I have also been a serious fantasy baseball player since 1992. I finished 5th out of 12 in the 2004 LABR-AL sponsored by the Sports Weekly and USA Today (six regular guys like me competed against six fantasy baseball experts). I also wrote several articles for Rotoworld.com on fantasy baseball in 2005 and am writing a few columns for Rototimes.com this year (these articles are coauthored with Cameron Bennett).

Saul and I grew up in Massachusetts and have followed the Red Sox for many years. I thought that analyzing the Johnny Damon situation using the AHP would be of relevance for this book, given that Saul and I have coauthored two papers on AHP.

One of the most interesting situations in the 2005 off-season hot stove league in MLB was the decision on the team that would sign Johnny Damon, a free agent center fielder who played for the Boston Red Sox from 2001-2005. Any team other than the Red Sox who signed Damon would have to give the Red Sox two players in the 2006 draft of high school and college players. Boston would receive no other compensation.

Both the Yankees and the Red Sox needed a center fielder for the 2006 season, and Johnny Damon was the best free agent center fielder available. The Yankees had Bernie Williams and Bubba Crosby on their roster before signing Damon. Bernie Williams had played for the Yankees since 1991. Since Williams's hitting and defense have diminished the last three years, the Yankees were not comfortable having him as their full-time center fielder in 2006. The Yankees declared Crosby their starting center fielder in November 2005, but the general opinion was that Crosby or a Crosby/Williams combination was not good enough for the Yankees. On the other hand, the Red Sox had no center fielder other than Damon.

The Yankees could have solved their problem by trading for a center fielder but had few players in their minor league system that could be used in a trade. The Yankees refused to trade Robinson Cano (2b) and Chien-Ming Wang (starting pitcher) in any trade. These players were attractive to other teams, because they played well at the major league level in 2005, were young, and had low salaries.

As such, the Yankees had to resort to signing a free agent to seriously upgrade their center field situation. Any free agent other than Damon would marginally upgrade the Yankees center field position, would probably be expensive, and this player would not perform at Damon's level. The Yankees could trade some players other than Cano and Wang to get a center fielder. This player would be of reasonable ability but probably have some risk attached to him.

In this paper, we use AHP to model the Yankees' decision-making process in signing Johnny Damon. We believe this decision was made in the latter part of November, 2005 or early part of December, 2005, and that the Yankee management (George Steinbrenner, Owner, and Brian Cashman, General Manager) were key to making this decision. Damon signed with the Yankees on December 23, 2005.

MLB's luxury tax complicates the decision on expending the money to sign a free agent as expensive as Damon. A luxury tax is imposed on a team if a team's total salary for the season exceeds a threshold set by MLB and the MLB players union. Ronald Blum, a baseball writer for the Associated Press, reported the following in a December 21, 2005 article.

The Yankees owe $34,053,787 in luxury taxes for 2005. The Yankees paid $25,964,060 in luxury tax in 2004 and $3,148,962 in 2003. Boston paid $4,156,476 in luxury tax in 2005 and $3,148,962 in 2004. The Yankees exceeded the payroll threshold for the third time in 2005. Under the labor contract that began after the 2002 season, the Yankees were taxed at a 40 percent rate on their salary to players exceeding $128 million. Boston topped the $128 million threshold for the second time in 2005. Boston's luxury tax was 30 percent of their salary in 2005. Both teams will be taxed at the 40% level in 2006. In addition to their luxury taxes, New York and Boston paid over $30 million in revenue sharing in 2005.

2 The AHP Model

The three aspects of an AHP model controlled by the decision-maker are the alternatives, the criteria, and the pairwise comparisons. There are no subcriteria in this analysis. In this section, the alternatives and the criteria are described. In Section 3, the pairwise comparisons, necessary for doing an analysis using AHP, are presented. In Section 4, the results of this analysis are given and, in Section 5, some closing remarks are made.

2.1 Alternatives

The three alternatives for this analysis are now given. The identifier for each alternative used in the computer software, Expert Choice, is given in bold.

- **NOTHING:** Do nothing. Go into the 2006 season with Bernie Williams and Bubba Crosby as their center fielders.
- **ADD-CF:** Add a free agent center fielder that would cost between $1.5 million and $4 million a year.
- **DAMON:** Sign Johnny Damon.

2.2 Criteria

We now present the five criteria for this model. The title for each criteria used in the computer software, Expert Choice, is given in bold.

- **Cost:** This criterion represents the cost that that Yankees incur over and above the cost paid to Williams and Crosby.
    1. **NOTHING** has a cost of zero.
    2. **ADD-CF** has a cost of about $1.5 - $4 million/year if the cost for exceeding the Yankees' luxury tax payment is not considered.
    3. **DAMON** has a cost of $13 million/year (or $52 million for the 4 years of the contract) if the cost for exceeding the Yankees' luxury tax payment is not considered.

  Since the Yankees will be over MLB's luxury tax limit for 2006 (and the foreseeable future), the Yankees would pay a 40% penalty for alternatives **ADD-CF** and **DAMON**. Thus, **ADD-CF** would probably cost the Yankees between $2.1 and $5.6 million in 2006, and **DAMON** would cost the Yankees about $18.2 million in each of the years, 2006-2009, or $72.8 million for the duration of the Damon contract.
- **BEN/COST:** This criterion is the ratio of the perceived benefit to the Yankees in implementing a particular alternative to the cost to the Yankees for implementing the alternative. We used this criterion in [3] and [4]. In this particular context, the BEN/COST ratio for the various alternatives can be described as follows:
    1. The BEN/COST ratio for **DAMON** is quite high. The Yankees believed that the benefit they will receive for signing **DAMON** is

quite high when compared to the cost of alternative **DAMON** even though Johnny Damon is expensive. Johnny Damon is regarded as an excellent player, a strong presence on the bench, and an excellent person in the community.

2. The BEN/COST ratio for **ADD-CF** is somewhat less than **DAMON**.
3. The BEN/COST ratio for **NOTHING** is less than **ADD-CF**.

The Yankee organization operates differently than virtually any other organization in professional sports in the United States. Expressions such as "If the Boss wants it, the Boss gets it" is often heard. The Boss is Yankee owner, George Steinbrenner. Thus, BEN/COST placed more of an emphasis on the benefit that the player gives and less of an emphasis on the cost of the player. Other organizations might rank the BEN/COST ratio differently for these alternatives and place more of an emphasis on the cost of the alternatives.

- **ONFIELDS:** This criterion represents the anticipated on-field performance in the short term (next two years). The ONFIELDS for the various alternatives are as follows:
  1. The ONFIELDS for **DAMON** is high.
  2. The ONFIELDS for **ADD-CF** is significantly less than **DAMON**.
  3. The ONFIELDS for **NOTHING** is somewhat less than **ADD-CF**.
- **SUBSCRIBER:** This criterion represents the importance of the alternative in the building and maintenance of the Yankees cable subscriber base. The Yankees Entertainment and Sports regional cable TV network (YES Network) is the only provider of Yankee games to both cable network providers and games on TV. About 130 games a year are provided solely to cable network providers and 30 games a year are shown on public TV. The YES Network and the Yankees are owned by Yankee Global Enterprises LLC (YGE), a holding company. The YES Network is an important revenue generator for YGE. Because of the organization of YGE, the revenues generated by the YES Network are not subject to MLB's revenue sharing – in other words, all revenues generated by the YES Network remain with YGE and, hence, can be used to build the Yankees.

  Through 2005, the YES Network had no major competition from the New York Mets, the other major league baseball team in the New York area. Starting in 2006 or 2007, the New York Mets plan to launch their own cable network. For the first time, the YES Network will be in competition for subscribers. The New York Mets have acquired several big name players in the last two years in order to improve their team. The Mets hope that this improved team will increase the number of subscribers to their network.

  My assessment is that the Yankees believe that the signing of DAMON will help to attract subscribers and, hence, increase revenues for YGE. With this background, my evaluation of the three alternatives according to the SUBSCRIBER criterion is as follows:
  1. The SUBSCRIBER for **DAMON** is very high.
  2. The SUBSCRIBER for **ADD-CF** is low.
  3. The SUBSCRIBER for **NOTHING** is very low.

- **RED SOX:** The criterion represents the Yankees' belief that the signing of Damon will hurt the Red Sox, as well as benefiting the Yankees. It is generally believed that the Yankees entered the negotiations for the service of Johnny Damon in order to build up the cost of Johnny Damon to the Boston Red Sox. This was especially true in late November and the first couple of weeks in December. In their negotiations with Johnny Damon and his agent, the Yankees found that they could sign Damon for $52 million for 4 years. The Yankees believed that signing Damon would hurt the Red Sox by taking away their center fielder.

  The evaluation of the three alternatives according to the RED SOX criterion is as follows:

  1. The RED SOX for **DAMON** is very high. The Yankees believed that signing DAMON would hurt the Red Sox.
  2. The RED SOX for **ADD-CF** is low. The Yankees believed that signing another outfielder would hurt the Red Sox a little, since it would benefit the Yankees.
  3. The RED SOX for **NOTHING** is very low. The Yankees believed that doing nothing would not hurt the Red Sox.

2.3 Notes

If we were conducting this analysis for another team, we would use different categories. One category that we would most likely use in ONFIELDL. **ONFIELDL** is the anticipated on-field performance of the alternative for the years after the next two years.

There is a general consensus that Johnny Damon will perform quite well for the first two years of his contract, but his performance the last two years may be inferior to his first two years. Sean McAdams of the Providence Journal wrote the following on January 21, 2006: "The Yankees filled their center field hole with Damon. I'm not sure they'll love that contract in 2008-09, but for the next two years at least, he'll help them. And, of course, atop that lineup, he could score 130 runs." In other words, Johnny Damon in 2008-09 may be the equivalent of Bernie Williams today.

Furthermore, the players used in the other two alternatives may not be part of the Yankee organization in 2008-9. Thus, **ONFIELDL** was eliminated from this analysis. I believe that the Yankees acquired Damon for 2006 and 2007. If Damon is not able to perform at the 2006-2007 levels in 2008 and 2009, then the Yankees can decide to find a replacement in 2008. The Yankees would definitely have to find a replacement in either 2007 or 2008 under the alternatives **ADD-CF** and **NOTHING**.

3 AHP Analysis

The AHP analysis for this problem is now presented. The AHP tree is described in Section 3.1, the pairwise comparisons of the criteria with respect to the GOAL node is given in Section 3.2, the pairwise comparisons of the alternatives with respect to each criterion is presented in Section 3.3, and the synthesis of the results is given in Section 4.

As noted previously, a summary of the career statistics for the players involved in this analysis is presented in Appendix 1. In Appendix 1, Bubba Crosby and Bernie Williams represent alternative **NOTHING**, Jason Michaels represents alternative **ADD-CF**, and Johnny Damon represents alternative **DAMON**. Coco Crisp is used later in this paper.

3.1 AHP Tree

The AHP tree has three levels:

Level 1: **Goal Node**

The Goal Node represents the AHP analysis of the Yankees decision to sign Johnny Damon.

Level 2: **Criteria Nodes**

The five criteria become nodes on Level 2 of the AHP Tree. These criteria, described in Section 2 of this paper are the following: **COST**, **BEN/COST**, **ONFIELDS**, **SUBSCRIBER**, and **RED SOX**. Each of these nodes is connected by an edge to the Goal Node.

Level 3: **Alternatives**

The three alternatives in this problem are **NOTHING**, **ADD-CF**, and **DAMON**. Each alternative on Level 3 is connected by an edge to each of the criteria listed in Level 2. There are 15 nodes on Level 3 of the AHP tree.

We next present the pairwise comparisons for each of the criteria on Level 2 of the AHP tree and for each of the alternatives on Level 3 for each of the criteria given on Level 2.

3.2 Pairwise Comparisons for the Criteria

My evaluation of the five criteria for this problem is as follows. Criterion **ONFIELDS** was regarded as the most important criterion and moderately preferred to criterion **SUBSCRIBER**. These criteria were regarding as moderate to strongly more important than criteria **BEN/COST** and **COST**, who were regarded as equal in importance. **BEN/COST** and **COST** were regarded as moderately more important than **RED SOX**.

The pairwise comparison matrix for the criteria, using the preferences describe above, is given in Table 1. In Table 1 (as well as Tables 2-6 given in

Section 3.3), the pairwise comparisons form a reciprocal matrix. In a reciprocal matrix A(I,J) = 1/A(J,I).

Table 1. Pairwise comparisons for the criteria with respect to the GOAL node.

	COST	BEN/COST	ONFIELDS	SUBSCRIBER	RED SOX
COST	1	1	1/4	1/3	3
BEN/COST	1	1	1/4	1/3	3
ONFIELDS	4	4	1	2	6
SUBSCRIBER	3	3	1/2	1	5
RED SOX	1/3	1/3	1/6	1/5	1

3.3 Pairwise Comparisons for the Alternatives with respect to the Five Criteria.

We now present the pairwise comparison matrices for the alternatives with respect to each of the five criteria. Our reasoning in making these pairwise comparisons was given in Section 2.

Table 2. Pairwise comparison of alternatives with respect to criterion **COST**.

	NOTHING	ADD-CF	DAMON
NOTHING	1	2	9
ADD-CF	1/2	1	8
DAMON	1/9	1/8	1

Table 3. Pairwise comparison of alternatives with respect to criterion **BEN/COST**.

	NOTHING	ADD-CF	DAMON
NOTHING	1	1/2	1/3
ADD-CF	2	1	1/2
DAMON	3	2	1

Table 4. Pairwise comparison of alternatives with respect to criterion **INFIELDS**.

	NOTHING	ADD-CF	DAMON
NOTHING	1	1/2	1/8
ADD-CF	2	1	1/7
DAMON	8	7	1

Table 5. Pairwise comparison of alternatives with respect to criterion **SUBSCRIBER**.

	NOTHING	ADD-CF	DAMON
NOTHING	1	1/2	1/8
ADD-CF	2	1	1/7
DAMON	8	7	1

Table 6. Pairwise comparison of alternatives with respect to criterion **RED SOX**.

	NOTHING	ADD-CF	DAMON
NOTHING	1	1/2	1/8
ADD-CF	2	1	1/7
DAMON	8	7	1

3.4 Synthesis of the Results

We solved the AHP model described in Section 3.1-3.3 using the Expert Choice software. Expert Choice is an implementation of the AHP and the software system that Saul Gass and I used in our research and teaching. The results of this analysis are given in Table 7. It is noted that each inconsistency measure that was computed was between .01 and .04, indicating that the pairwise comparisons in Tables 1-6 were very consistent.

Table 7 can be explained as follows:

- The first row of Table 7 gives the names of the criteria for this problem.
- The second row of Table 7 gives the weights for the criteria, where the weights are determined by Expert Choice for the pairwise comparison matrix in Table 1.
- The last three elements of the first column of Table 7 give the names of the alternatives for this problem.
- Each of the last three elements of the second to sixth columns of Table 7 gives the weights for each alternative in the row that can be attributed to

the criterion listed in the column. Each weight is determined by Expert Choice for the pairwise comparison matrix given in Tables 2-6.

- The last three elements of the TOTAL column in Table 7 give the results of the synthesis step in the AHP analysis. These results give the weight each alternative earns based on the pairwise comparisons in Tables 1-6 and the AHP analysis.

Table 7. Summary of results for the AHP analysis for the pairwise comparison matrices given in Tables 1-6.

	COST	BEN/COST	ONFIELDS	SUBSCRIBER	RED SOX	TOTAL
Criteria Wgt.	.115	.115	.437	.283	.050	-----
Nothing	.589	.163	.081	.081	.081	.165
NEW-CF	.357	.287	.135	.135	.135	.191
DAMON	.054	.540	.784	.784	.784	.644

4 Analysis of the Results

Examining the results of the last three elements of column 7 of Table 7, the Yankee strategy of signing Johnny Damon was the overwhelming preferred strategy with 64.4% of the weight. Each of the other two alternatives got less than 20% of the weight.

The only criterion for which **DAMON** was inferior to the other two alternatives was the **COST** criterion. **DAMON** was overwhelmed in this criterion because this alternative is millions of dollars more expensive than the other two alternatives. **DAMON** overwhelmed the other two alternatives in the other four criteria, earning over 50% of the weight in the **BEN/COST** criterion and over 70% of the weight in the other three criteria.

If the Yankees were not able to sign Johnny Damon, then it appears that the Yankees prefer **NEW-CF** to **NOTHING**. Since the Yankees signed Johnny Damon, the Yankees did not have to worry about whether to implement alternative NEW-CF.

5 Final Remarks

a. The YES Network and Subscriber Television

I believe the **SUBSCRIBER** category introduces an interesting and important component in the analysis of baseball operations. This category indicates that baseball is a big business and faces many important and expensive decisions outside of what happens on the baseball field.

While writing this article, I was curious about how much revenue can be generated from subscriber-type networks. I was unable to get an answer to this issue. However, I believe the following quote, taken from an article by Sasha Talcott in the Boston Globe on January 31, 2006 gives some indication of the potential revenues (NESN means the New England Sports Network): "The Red Sox have plenty of experience in another medium: television. The Sox's owners (including The New York Times Co., owner of the Globe, which has a 17% stake) also control NESN, which telecasts Sox games. NESN generates tens of millions in profits annually – conveniently exempt from Major League Baseball's revenue-sharing rules. (Team financial statements are private, so NESN's exact profits are unavailable.)"

At noted above, if properly organized, subscriber TV and radio revenues are exempt from MLB's revenue sharing program. In 2005, the Yankees and the Red Sox paid over $30 million in revenue sharing as well as paying a luxury tax. The Yankees and the Red Sox believe that sheltering subscriber revenues from revenue sharing is important and have organized their businesses to take advantage of these rules.

b. What Happened to the Red Sox since Damon was Signed by the Yankees

The Red Sox were without a center fielder until January 27, 2006 when they made a trade with the Cleveland Indians for Coco Crisp. Coco Crisp is an excellent 26 year old center fielder who is recognized as a good defensive outfielder, fast, and possesses offensive statistics at about the same level as Johnny Damon (see Appendix 1 for an analysis of some of these statistics). (The Indians traded Arthur Rhodes for Jason Michaels a few days before the trade of Coco Crisp to the Red Sox. Jason Michaels is scheduled to take the place of Coco Crisp in left field for the Indians in 2006. In this paper, Jason Michaels is considered to be the player in alternative **New-CF**. Jason Michaels' batting statistics is given in Appendix 1.)

From reading the various articles on the Internet, two things became clear:

a. The Red Sox were only willing to pay Damon up to $40 million over four years.

b. The Red Sox anticipated that it was quite likely that they would lose Damon as a free agent in 2006 and targeted Crisp as a replacement for Damon as early as the latter part of the 2004 season. Crisp is now a member of the Red Sox and cannot become a free agent until after the 2009 season. The cost for Crisp over the 2006-2009 time period is estimated to be $20 million versus the $40 million for Damon or $52 million if the Red Sox were to try to match the Yankees offer. There are no luxury tax penalty costs figured into these computations.

In the context of this paper, if we were to perform the AHP analysis for the Red Sox on this problem rather than the Yankees, then the criteria would be somewhat different, the alternatives would be different, and the pairwise comparisons would be different.

c. Other Comments

This case study illustrates that baseball is a big business and teams are making multi-million dollar personnel decisions. Tools such as AHP can be valuable processes for helping an organization make these key strategic decisions. Further, it appears that subscriber television is going to become a bigger factor in these decisions in the future. Over time, I believe there will be fewer games shown on free television or on cable television without a charge. Viewers will have more games to watch, but will have to pay a fee for watching these games.

References

1. L. Bodin. Use of the Analytic Hierarchy Process in major league baseball, *Proceedings of the Fifth International Symposium on the Analytic HierarchyProcess, Kobe, Japan*: 129-134, 1999.
2. L. Bodin and E. Epstein. (1999), Who's on first – with probability .4, *Computers & Operations Research*, 27 (2): 205-215, 1999.
3. L. Bodin and S.I. Gass. On teaching the Analytic Hierarchy Process, *Computers & Operations Research*, 30 (10): 1487-1498, 2003.
4. T.L. Saaty. *The Analytic Hierarchy Process*. McGraw-Hill, New York, 1980.
5. T.L. Saaty. *Fundamentals of Decision Making and Priority Theory with the Analytic Hierarchy Process*, RWS Publications, Pittsburgh P, 1994.

APPENDIX 1: Summary Batting Statistics for the Players Mentioned in this Article

Batting statistics for the outfielders mentioned in this article are presented in Table A-1. Bubba Crosby and Bernie Williams are the two outfielders considered in alternative **NOTHING**. Jason Michaels (who signed a 1 year, $1.5 million contract this winter), is an example of an outfielder who could have been obtained under alternative **ADD-CF** to be added to the Yankees along with Crosby and Williams. Johnny Damon, Crosby, and Williams are the outfielders considered in alternative **DAMON**. Coco Crisp is the outfielder that the Boston Red Sox traded for as the replacement for Johnny Damon.

The 2005 statistics for Crosby, Michaels and Crisp are presented since they have only played in the major league for a couple of years and averaging their statistics would not properly capture the 2005 statistics. Williams's statistics are broken down into two time periods – 1993-2002 and 2003-2005 to show the difference in performance between these periods. Damon's statistics are from 1995-2004 and disregard 1994, as he was either not a starter in 1994 or did not play the entire season in 1994.

Table A-1. Summary statistics for the players considered in this article.

Year	Bubba Crosby 2005	Bernie Williams 1993-2002	Bernie Williams 2003-2005	Johnny Damon 1996-2005	Jason Michaels 2005	Coco Crisp 2005
Age-2006	29	37	37	32	29	26
G	76	139	136	151	105	145
AB	98	538	497	600	289	594
H	27	168	128	174	88	178
2B	0	32	22	32	16	42
3B	1	5	1	8	2	4
HR	1	22	16	12	4	16
RBI	6	94	66	68	31	69
R	15	98	78	104	54	86
BB	4	75	70	53	44	44
K	14	83	78	68	45	81
SB	5	12	2	27	3	15
CS	1	7	2	7	3	6
AVG	0.276	0.312	0.258	0.29	0.304	0.3
OBP	0.304	0.36	0.321	0.348	0.396	0.345
SLG	0.33	0.44	0.37	0.43	0.42	0.47
OPS	0.63	0.795	0.688	0.778	0.811	0.81

Some of the most important statistics in Table A-1 for analyzing the performance of a hitter are the age of the player, the batting average (AVG), the on-base percentage (OBP), the slugging percentage (SLG), and the OPS of the player. Other statistics Table A-1 are official at bats (AB), hits (H), two base hits or doubles (2B), three base hits or triples (3B), home runs (HR), runs batted in (RBI), runs scored (R), walks (BB), strikeouts (K), stolen bases made (SB), and attempted stolen bases not made (CS).

The age of a player can be a critical criterion in evaluating a player. It has been found that many hitters put up their career statistics when they are 26 or 27 years old, and that the hitting statistics begin to decline for players older than 27. There is an inherent danger in trading for hitters over 27 years old.

The important ratio statistics, AVG, OBP, SLG, and OPS are defined as follows:

$$AVG = H/AB$$

$$OBP = \frac{\text{Hits (H)} + \text{Walks (BB)}}{\text{Official At Bats (AB)} + \text{BB}}$$

$$SLG = \frac{\text{Hits(H)} + 2*\text{Doubles (2B)} + 3*\text{Triples (3B)} + 4*\text{Home Runs (HR)} - 2B - 3B - HR}{AB}$$

$$OPS = OBP + SLG$$

Recently, OPS has become one of the key statistics used in evaluating a player. Acccording to Ron Schandler's 2005 Baseball Forecaster, "OPS combines the two basic elements of offensive production – the ability to get on base (OBP) and the ability to advance runners (SLG)." Schander goes on to note that baseball's top batters will have an OPS over .900 whereas the worst batters will have levels under .600." Further, in some analyses, OBP includes No. of Times Hit by a Pitch.

We will now analyze each of these alternatives.

A-1 Alternative NOTHING

In 2005, Crosby is 29 and has an OPS of .630 and Williams is 37 and has an OPS of .688. William's OPS is less than .8 for the past 3 years (2003-2005) and had an OPS greater than .9 for the years (1996-2003). The Yankees were concerned about his performance for 2006 with such a rapid decline the last three years.

A-2 Alternative ADD-CF

We proposed in the alternative ADD-CF that the Yankees try to trade for a player or sign a player that costs about $1.5 million to $4 million a year. I decided to use Jason Michaels, an outfielder that was traded from the Philadelphia Phillies to the Cleveland Indians, as an example. Michaels is 29 years old, never played full time in the National League, has only 4 years of major league experience (except for the 6 games in 2001) and signed a one year contract with the Phillies for $1.5 million for the 2006 season.

Michaels has reasonable statistics. His OPS was at least equal to .779 for the past 4 years. I believe that Michaels is a hitting upgrade to Crosby and Williams if he can play center field. Michaels is not a proven quantity, however, since he never had more than 290 at bats in a season. Cleveland traded for Michaels in January 2006 as a replacement for Crisp although they plan to play Michaels in left

field rather than center field. Cleveland has Grady Sizemore, a great talent, to play center field.

I went through the other outfielders that were free agents. Jason Michaels appears to be equivalent to any free agent center fielder costing no more than $4 million a year.

A-3 Alternative DAMON

Damon was clearly the best free agent available in the winter of 2005-2006 and commanded a premium price. Damon possesses reasonable defensive skills (and center field is as much a defensive position as well as an offensive position). The only negative that can be said about Damon is that he has had problems with his throwing the past couple of years due to injuries. As noted in the body of the paper, Damon is a great team player and a positive presence in the dugout (bench).

Damon's batting average was over .300 and OPS was over .800 the past two years – very acceptable statistics for a center fielder. Damon scored over 100 runs and stole 18+ basis in every year since 1998 – excellent statistics for a leadoff hitter. Further, Damon has shown reasonable patience as a hitter. Generally, Damon's walks and strikeouts were almost equal. These statistics should repeat themselves over the 2006 and 2007 seasons, especially when you consider the hitting potential of the Yankees in 2006 and 2007.

A-4 DAMON versus CRISP

Crisp has only played full time for the past three years. Crisp's defensive skills are similar to Damon's defensive abilities. Crisp can field well and his throwing is reasonable and similar to Damon's throwing.

The Boston Globe on February 18, 2006 ran the following comparison of the aggregated statistics for his first four years of Damon's and Crisp's careers.

- Crisp has 1,626 major league at-bats. In that time, Crisp had a .287 batting average, .332 OBP, .424 SLG, .756 OPS, 35 home runs, 176 RBIs, 54 stolen bases, and 29 times caught stealing.
- Damon, had 1,623 major league at-bats in his first four years. In that time, Damon had a .272 average, .326 OBP, .395 SLG, .721 OPS, 29 homers, 157 RBIs, 65 steals and 25 times caught stealing.
- Damon, in his second, third, and fourth major league seasons, steadily increased his home run totals (6 to 8 to 18), doubles (22 to 12 to 30), and OPS (.680 to .723 to .779). Crisp's home run totals have increased quicker (3 to 15 to 16), and so have his doubles (15 to 24 to 42) and OPS (.655 to .790 to .810). Crisp tends to strike out nearly twice as many times as when he walks, is not as effective base stealer as Damon, and has more power than Damon.

Boston's hope is that Crisp, entering his 5th season, will put up statistics in the next four year similar to what Johnny Damon put up while playing for the Red Sox.

Index

algebraic, 202, 308, 322, 323, 329, 330, 393, 399, 403–407, 409, 411
algorithm, 26, 31, 34–36, 39, 42, 44, 50, 58–61, 65, 89, 134, 179–197, 210, 277, 279, 281, 282, 285, 287, 291, 292, 298–304, 341, 342, 347, 348, 367–375, 377, 378, 381–392, 396, 398–400, 405, 411–413
Analytic Hierarchy Process (AHP), 38, 121, 122, 415, 420, 425
arbitrage, 202, 210
arithmetic, 307, 308, 311, 312, 328
auction, 106, 108, 153–156, 158–160, 162, 163, 165, 167–171, 173–177, 207, 208, 210, 367, 369, 388, 391
augmentation, 285, 287, 291–295, 297–300, 302–304, 373, 392

baseball, 415, 416, 418, 423–425, 427
Ben Franklin, 115, 116, 121, 122

clustering, 221, 235, 241, 244–246, 367, 369, 370, 375, 388, 391
combinatorial auctions, 153, 158, 169, 173, 175–177
combinatorial optimization, 173, 246, 247, 270
complexity, 38, 109, 142, 173, 175, 176, 179, 180, 184–186, 192, 196, 197, 368, 372, 379
computational, 25–28, 37, 38, 65, 66, 153, 159, 160, 173, 176, 179, 180, 184, 185, 187–190, 192, 194–197, 211, 213, 214, 217, 220, 234, 247–249, 255, 256, 260, 264–267, 269, 270, 273, 352, 362, 365, 368, 371, 374–377, 380, 384, 388, 391, 392, 412, 413
connected components, 285–287, 289, 291, 292, 298
constraint, 86, 112, 113, 115, 120, 122, 125–128, 161, 164–166, 171, 173, 205, 208, 209, 212–215, 218, 220, 221, 229, 230, 235, 237, 240–243, 245, 282, 311, 314, 315, 326, 335, 336, 340, 341, 343–345
coordinate search, 393, 411
cost-benefit analysis, 121
CPLEX, 120, 173, 345

data mining, 108, 139, 235, 245
database, 58, 235, 242, 367–369, 372, 376, 377, 388, 392
decision making, 33, 44, 47, 61, 115, 120, 123, 235, 239, 415, 425
decision support, 108, 137, 147, 171
decision theory, 123
derivatives, 308, 351, 354
distribution, 39, 60, 108, 112, 130, 179, 202, 203, 207, 248, 256, 353–355, 357–362, 369, 373–375, 378–382, 389, 401
dual, 34, 36, 71, 125, 163–167, 200, 201, 205, 211, 212, 214, 219, 225, 226, 228, 230, 232, 300, 301, 304
duality, 199–201
duality theory, 125, 127
Dutch auction, 154

economic order quantity, 307, 309, 329
economic production quantity, 308, 316
Edelman, 127, 138, 145, 149
empirical, 39, 177, 179, 180, 184, 187, 195, 196, 274, 276, 364, 382, 390
English auction, 154
equilibrium, 51, 108, 170, 200, 204–206, 210
Euclidean, 39, 247, 269, 272, 277, 281, 393, 403, 405

facility location, 103
Farkas, 199–204, 207, 209, 210
Franklin square, 115, 120

geometric, 37, 40, 277, 279, 283, 307, 308, 311, 312, 328, 393, 394, 397, 399, 403, 405–408, 410, 411
global optimization, 367–369, 385, 386, 390, 391
graph, 39, 171, 180, 186, 187, 197, 228, 245, 247, 278, 285–298, 300–304, 353, 373, 391, 392

heuristic, 40, 66, 112, 211, 220, 221, 241, 242, 247, 248, 257, 262, 269–272, 274, 277–283, 335, 339, 342, 345, 349, 384, 386

integer programming, 38, 81, 90, 163, 211–213, 221, 235, 242, 301, 333, 335
inventory control, 123

LADAR, 393–396, 399–402, 405, 412
Lagrangian relaxation, 235, 237, 245
linear algebra, 42, 59, 199, 200
linear programming, 19, 23, 25, 26, 28, 30, 34, 35, 38, 41–44, 51, 59–66, 68, 69, 71, 73, 81, 84–86, 91, 97, 99, 101, 111, 112, 123, 127, 129, 142, 153, 163, 199, 200, 211–215, 220, 221, 223–225, 228–230, 234–239, 242, 245, 246, 351, 352, 394, 398, 411

magic square, 116, 117, 120
management, 20, 25, 29, 32, 33, 41, 44, 45, 47–51, 53, 55, 56, 58, 59, 62, 64–66, 68, 71, 72, 81, 107, 108, 112, 123–129, 131–137, 139, 140, 143, 144, 147, 149, 176, 179, 199, 211, 220, 221, 235, 242, 270, 307, 325, 346, 349, 351, 352, 365, 416, 431
manufacturing, 37, 65, 107, 112, 123, 124, 127–129, 131, 133, 139, 307, 394, 412
Markov, 41, 112, 203, 251, 252, 378
Markov chain, 203, 251, 252
mathematical programming, 20, 25, 35, 65, 66, 115, 122, 123, 125, 132, 197, 220, 221, 246, 333, 343
matroid, 223
maximization, 367, 369, 372, 378, 383, 390
minimization, 86, 167, 213, 221, 393, 404, 410, 411, 413
modeling, 20, 23, 26, 31–33, 40, 41, 44–47, 49, 51–53, 55, 58, 60–62, 64, 66, 75, 92, 95, 108, 111, 112, 115, 122–125, 132–134, 142, 270, 334, 391, 394, 411
Monte Carlo, 64, 149, 251, 351–353, 364, 365, 367, 375, 377, 378, 380–383, 385, 392
multi-criteria, 63, 235, 239
multi-criteria decision making, 235, 239
multi-objective programming, 235, 238

network, 27, 30, 40, 41, 62, 103, 108, 156, 179, 180, 185–197, 221, 235, 237, 245, 270, 281, 285, 300, 335, 349, 351, 352, 362, 364, 365, 367, 373, 390
noisy traveling salesman problem, 247, 248, 269

operations research, 19–21, 23, 28, 30, 32, 33, 43, 44, 47, 52–56, 58–69, 71, 73–75, 94, 96, 98, 101, 102, 104, 107, 108, 115, 120, 122–124, 129, 133–140, 144–146, 149, 153, 176, 197, 199, 220, 221, 235, 245, 246, 269, 270, 282, 335, 349, 351, 352, 365, 367, 368, 385, 394, 412, 425
optimization, 26, 35, 37, 41, 54, 57, 59, 61, 66, 68, 71, 73, 77–79, 81, 98,

103, 106, 108, 112, 126, 160–167, 171, 173–175, 221, 235, 246, 247, 254, 270, 271, 304, 307, 311, 314, 315, 318, 325–327, 329, 334, 342, 349, 364, 365, 367–369, 372–374, 378, 383, 385, 386, 389–393, 398, 411–413, 430

parametric programming, 34, 35, 59, 235, 236, 238, 241, 243, 245, 246
partitioning, 106, 161, 174, 183–185, 192–197, 211–215, 220, 221, 235, 241–246, 333, 335, 336, 339, 342–344, 349
PERT, 351, 352, 364, 365
polynomial, 181, 185, 192, 193, 196, 236, 237, 242, 245, 246
polytope, 212, 220, 223–225, 234
pricing, 108, 139, 153, 156, 163, 167, 391
primal, 165, 200, 201, 205, 226, 228, 232
probability, 30, 146, 202, 203, 207, 210, 238, 245, 246, 248, 251
project management, 29, 47, 48, 351, 365

queueing, 30, 40, 103–105, 110, 130, 131
queueing theory, 103, 123, 125, 129, 130, 132

rationality, 208–210
regression, 190, 191, 194, 195, 221, 353, 393, 398, 403–405, 412
risk neutral probability, 202

scheduling, 112, 128, 129, 134, 139, 179, 211–214, 220, 221, 282, 333–335, 337–343, 345, 348, 349, 352, 365
sensitivity analysis, 48, 351, 354, 365
shortest path, 179–185, 187, 188, 192, 193, 195–197, 237, 279, 359, 360
simplex method, 25–27, 56, 61, 85, 88, 223, 225, 226, 235, 236, 238, 245
simulation, 28, 31, 40, 51, 59, 61, 64, 66, 82, 139, 149, 175, 251, 349, 351–353, 355, 364, 365, 375, 377–382, 384, 390, 391
spheres, 35, 65
strongly connected component, 285, 286, 291–294, 296, 298–300
supply chain management, 33, 65, 176

tableau, 34, 223, 225, 226, 228–232
tour, 247–251, 254, 255, 257, 259, 262, 264, 269, 271, 273, 274, 276, 279–281
traveling salesman problem, 44, 247, 248, 269–271, 281

Printed in the United States