INFORMS Journal on Computing
Vol. 25, No. 1, Winter 2013, pp. BR1–BR10
ISSN 0899-1499 | EISSN 1526-5528 | 2501
http://dx.doi.org/10.1287/ijoc.2013.1.br

Book Reviews

Harvey J. Greenberg, Editor
Professor Emeritus, University of Colorado Denver, hjgreenberg@gmail.com

The *INFORMS Journal on Computing* (IJOC) reviews books on subjects at the interface between operations research and computer science. We welcome books on theory, applications, computer systems, and generally any subject covered by an IJOC Area, or any combination of these. This includes both printed and electronic books. In addition, we consider comparative reviews—several books on one relevant topic. Team reviews are also possible, particularly for a large, broad-scope book, such as an encyclopedia. For further information, please visit www.informs.org/Pubs/IJOC/Book-Reviews.

In this issue, we review three books. The first uses the paradigm of the traveling salesman problem to bring the limits of computation to life for algorithm design and analysis. The reviewer, Gábor Pataki, received his Ph.D. in Algorithms, Combinatorics, and Optimization from Carnegie Mellon University in 1996. He joined the faculty of the Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, in 2000. His research is in integer and convex programming with recent papers on the complexity of the classical branch-and-bound algorithm, the closedness of the linear image of a closed convex cone, and semidefinite programs that behave badly from the standpoint of duality. Among his publications, Gábor wrote "Teaching Integer Programming Formulations Using the Traveling Salesman Problem" [*SIAM Review* 45-1, 2003], directly applicable to the essence of this book. His review is both informative and insightful. He concludes, "the next time a student asks me 'Why are you guys so crazy about research?', I will just answer, 'Read this book."'

The second book deals with e-commerce. The reviewer, Steve Kimbrough, is Professor of Operations and Information Management at The Wharton School, University of Pennsylvania. He is a recognized expert in many things centered around concepts and methods for knowledge-based decision support, including e-commerce. As early as 1997, Steve was invited to be the keynote speaker at the SOBU Symposium on Electronic Commerce, Tilburg University. His work with DSS is extensive, including heading up the Coast Guard's Knowledge-Based Support Systems project for ten years, as well as several other related sponsored research projects. Within e-commerce, he has a long, ongoing research stream applying formal logic to business messaging with the aim of developing a formal language for business communication. His teaching and research continue to include this growing area, making him eminently qualified to review this book. Steve points out, "The book focuses on understanding complex, multi-attribute choice tasks in the online consumer purchasing context." He concludes, "I believe that this book has its greatest value as a resource and point of departure for those with an interest in doing research in this worthy area."

The third book deals with finance models, which uses Matlab for exercises. The reviewer is James Morris, Emeritus Professor of Finance in the School of Business at the University of Colorado Denver. After receiving his Ph.D. in Finance from the University of California, Berkeley, Jim joined the faculty at The Wharton School, followed by a range of visits in the U.S. and Europe. Joining UCD in 1982, he continued to teach financial modeling and publish in a variety of journals. During 2006–09, Jim was Director of the UCD M.S.-Finance Program and offers us a perspective as teacher, researcher, and curriculum designer. He notes that there could be a problem with finance majors in a business school who have insufficient mathematical background and mathematics majors who have insufficient knowledge of options, but he concludes, "In summary, *Monte Carlo Simulation with Applications to Finance* is well done."

In Pursuit of the Traveling Salesman: Mathematics at the Limits of Computation, by W. J. Cook, Princeton University Press, 2011. See http://press.princeton.edu/. ▷

Reviewed by: Gábor Pataki, Department of Statistics and Operations Research, University of North Carolina Chapel Hill, gabor@unc.edu.

How do you find a tour of minimum length that visits each city in a given list exactly once, and returns to the start? Not just a fascinating mathematical puzzle, the traveling salesman problem (TSP) finds uses in areas as diverse as logistics, data mining, and computational biology. Its study has seen remarkable successes in the past 60-odd years. Starting with the optimal solution of a 49-city USA instance in 1954 by Dantzig, Fulkerson, and Johnson, the current record among optimally solved instances has 85,900 cities. For the world-TSP with over 1.9 million cities, a solution proven to be within 0.047% of the optimum is available. These computational successes were driven by advances in our theoretical understanding of the polytope underlying the TSP, of approximation algorithms and heuristics.

Bill Cook of Georgia Tech is a leading authority on the subject; his Concorde code, written primarily with David Applegate, Bob Bixby, and Vašek Chvátal, currently holds the record in finding optimal solutions to TSP instances. Whereas technical, in-depth treatments of the topic have been available (see, for instance, [1, 2, 3]), this book is geared to the layperson with a solid high-school mathematics background. It aims to pique the reader's interest with the hope that some will take up research in the area, and deliver a new leap in our understanding of the TSP. The book's coverage of the subject is intuitive, and largely geometric, with many excellent illustrations. There are few equations. Still, there is sufficient mathematical detail to provide a good start to readers interested in a more technical treatment. The style is congenial, breezy, and entertaining; many anecdotes and pop culture references are included. Even seasoned researchers will find the book a truly enjoyable read, and it can serve as an ideal basis for a college level freshman seminar.

The appetizer of Chapter 1 takes us on a tour of TSP computation from the 1954 study of Dantzig and his colleagues, through a 1962 Proctor & Gamble challenge with a $10,000 prize, and to current records. It introduces the reader to the dichotomy between good and bad algorithms through the concepts of polynomial time computability, NP completeness, and the potential consequences of finding a polynomial algorithm for the TSP (including the end-of-world scenario in a 2001 science fiction novel by Charles Stross).

Chapter 2 covers the origins. We learn that actual salesmen (numbering approximately 350,000 in the United States in 1900), traveling judges (among them the young Abraham Lincoln), and traveling preachers were major consumers of more or less well designed tours. The origins of the Hamiltonian cycle problem are also treated here, through the Icosian game (a tour-finding game through the corners of the dodecahedron), and Tait's incorrect proof of the existence of Hamiltonian cycles in 3-connected, 3-regular graphs, and its connection to the 4-colour theorem. After a discussion of the contrast between Hamiltonian cycles and Eulerian tours, we arrive at the first true mathematical lecture on the TSP (more precisely, on the shortest Hamiltonian path problem), given by Menger in 1930. The chapter concludes with a review of the famous TSP constant, and its connection to Mahalanobis' study on the optimal inspection of jute crop in 1930s India.

Chapter 3 presents applications, some straightforward (in routing, and logistics), some much less so: TSP tours help find the optimal movement of telescopes to scan celestial objects and of laser beams to create artwork; they identify genome orders, create aesthetically pleasing tours of musical collections, and even help compose music.

Chapter 4 describes heuristics to construct good tours. Starting with the pegs-and-strings approach employed to find the optimal solution in the Dantzig et al. computation, we are walked through the nearest-neighbour algorithm and its relatives, Christofides' heuristic, the Lin-Kernighan heuristic of 1973, its striking improvement by Helsgaun in 1998, and others. All presented algorithms were run on a 42-city version of the Dantzig-Fulkerson-Johnson instance; pictures of the evolving tours provide exceptional illustrations. The chapter finishes with an overview of ant colony optimization, and genetic algorithms.

Chapter 5 covers the history, geometry, and duality of linear programming, and its use in attacking salesman problems. It shows how the TSP can be formulated as an integer program; this is the first place where the subtour inequalities appear. This chapter also contains a nice coverage of the Jünger-Pulleyblank bounding technique for geometric TSP instances, its relation to duality, and the 4/3 conjecture on the gap between the subtour LP relaxation and the optimal tour.

Cutting planes are the subject of Chapter 6, which starts with a step-by-step description of how Dantzig, and his colleagues proved the optimality of their 49-city tour. Separation algorithms are then covered for comb inequalities, clique trees, and Letchford's domino parity constraints, the latter significantly contributing to the solution of the record 85,900 city instance. Edmonds' "glimpse of heaven"—his polynomial-time perfect-matching algorithm—follows, and the tantalizing (if unlikely) possibility that a similar algorithm may exist for the TSP; these topics naturally lead to the equivalence of separation and optimization.

Likening the search for an integral point in a polyhedron to seeking a needle in a haystack, the addition of cutting planes to removing excess hay, and the process of branching to splitting the haystack, Chapter 7 presents an intuitive, and well-illustrated history, and description of branch-and-bound, and of branch-and-cut, the result of its combination with cutting planes.

Chapter 8 charts the progress of TSP computation from 1954 to the present day. As the sizes of the solved problems increase, each one is accompanied by a fascinating research story. We learn of the importance of parallel computing (the solution of the record 85,900 city instance took the equivalent of 136 years of computing on a single machine), and how the Applegate et al. team provided easy-to-verify proofs of the optimality of its tours. The Mona Lisa TSP, the world-TSP, and the star-TSP (with 100 thousand, 1.9 million plus, and 526 million plus cities, respectively) remain the top three challenges. The reader will likely be surprised to learn that the last one, with the best known solution "only" within 0.419% of optimal, is considered 10 years behind the world-TSP, for which there is a 0.047% solution available.

What is an algorithm, a "good" algorithm, and the "best" algorithm for the TSP from the theoretical viewpoint? Chapter 9 delves deeper into complexity theory, starting with the $1 million Clay Institute Prize for settling the P versus NP question. We learn about Turing machines, and Edmonds' quest to have polynomial time accepted for "good" in the world of algorithms. (Those who always took the equivalence of these for granted will be surprised to learn just how arduous his campaign was.) We are then led through the theory of NP-completeness, and approximability and inapproximability results. A survey of challenges less daunting than deciding whether P equals NP follows: beating the 60-year old record of Held and Karp's $O(n^2 2^n)$ algorithm, the 3/2 approximation ratio of Christofides' heuristic and the 220/219 inapproximability ratio of Papadimitriou and Vempala.

In the future we may end up solving TSPs without "Turing-style" computers; the chapter ends discussing various other models of computing, such as DNA computing (in vitro and in vivo), optical, and quantum computing, and even how time travel may get involved. Nonetheless, Bill concludes that the pegs-and-strings approach from the 1800s (also used by several modern research groups) would still beat any of these alternative methods.

Chapter 10 turns to the question of how sentient creatures (humans without recourse to computers, chimpanzees, pigeons, and the like) solve small salesman problems. All prove to be remarkably adept. In one experiment, 7-year old children routinely found solutions within 9.4% of optimal in 15-city instances and the TSP and related puzzles also appear in clinical tests in psychology. In another study a group sought aesthetic tours and another group, short tours; the results ended up rather close, with a member of the first team as overall winner. Still, the problem sizes considered here are small, and no proof of optimality (or near optimality) is supplied by either children, or pigeons; thus the author concludes that in a man vs machine TSP competition the computer is likely to come out on top.

"The salesman creating art" is taken up in Chapter 11. The paintings of Julian Lethbridge and Philip Galanter vividly capture the partitioning of the plane by a salesman tour. Mathematician Robert Bosch employs optimization more heavily in his TSP art: suitable constraints enforce symmetry of the TSP curve, or that certain pairs of points lie on opposite sides, "bending the curve to the artist's will." In other art pieces, tours or tour-like curves display more complex images. Artist Eric J. Morales creates striking human portraits from long, meandering, nonintersecting lines. The team of Robert Bosch and Craig Kaplan wields TSP tours to render classical paintings: the 100 thousand city Mona Lisa TSP and the 140 thousand Birth of Venus TSP are based on their work. The art of mathematician Jaroslav Nešetřil and painter Jiři Načeradský is covered next, and finally, we pay a visit to the Bonn Arithmeum, dedicated to computing, art, and music. The design of VLSI chips can be significantly improved using discrete optimization—the resulting layouts also give rise to attractive images that form part of the museum's collection.

Chapter 12 brings the conclusion: the TSP is addictive, as researchers attest, and more than an ingenious puzzle, or even a practical optimization problem. Unsolvable instances of any NP-hard optimization problem are likely to remain. How do we then still push the limits? By pulling out all stops (a memorable story from Chapter 4 is how Bill and David Applegate wrote a program that in turn wrote a program to examine 5-opt, 6-opt, etc. tour-improving exchanges. The output for 8-opt is a 17 million line plus long C-code). Bill advocates adapting this leave-no-stone-unturned attack on the TSP to other problems as well—in other words, the approach of "bashing on regardless."

As I mentioned before, the book is ideal for a layperson's independent study, for an enjoyable read for just about anyone, and for a first year seminar. Because the chapters are rather loosely coupled, I think some of them can be used independently to give students a look at topics that go beyond the main subject of a course. For instance, in an undergraduate course on discrete mathematics Chapter 5 can serve as a quick and friendly introduction to linear programming; in an introductory course on optimization (where the TSP may not otherwise be included) one could devote some lectures to covering parts of Chapters 3 through 8. Finally, the next time a student asks me "Why are you guys so crazy about research?", I will just answer, "Read this book."

References

[1] D. L. Applegate, R. E. Bixby, V. Chvátal, and W. J. Cook. *The Traveling Salesman Problem: A Computational Study.* Princeton University Press, 2007.

[2] G. Gutin and A. P. Punnen. *The Traveling Salesman Problem and Its Variations.* Kluwer, 2007.

[3] E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and D. B. Shmoys. *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization.* Wiley, 1985.