Book Review

"The Traveling Salesman Problem"

David L. Applegate, Robert E. Bixby, Vašek Chvátal and William J. Cook

Patrick D. Bangert
algorithmica technologies GmbH
Ausser der Schleifmühle 67
28203 Bremen
Germany

The traveling salesman problem (TSP) is a problem that all of us face regularly: Given a list of places and the knowledge of how far each pair of these places is separated, we are asked to construct a journey through these places and back again such that the total distance traveled is the shortest possible. This issue is the subject of the recent book by Applegate, Bixby, Chvátal and Cook.

The TSP is a, if not the, key problem of applied science today. In my personal opinion, this has four reasons: (1) The problem's appealing simple statement continues to motivate thousands of people to work on it and continues to motivate the public to learn a little mathematics, (2) many advances and challenges of applied mathematics and computer science are directly connected with it (for example the branch-and-bound techniques and the P = NP problem), (3) the TSP and its (heuristic) solutions serve as a basis for many extensions that are problems of actual practice and have relevance in everyday life.

It was with great pleasure and expectation that I received the book and started to read it. Based on the advertising, I expected to learn the secrets of the record setting program concorde that the authors wrote to solve TSP instances as well as to learn about new methods to solve a number (practical) problems intimately related to the TSP – a number of interesting problems can be obtained, for example, by adding constraints to the salesman such as allowing him a finite capacity of transporting goods.

The authors of the book are four of the world leaders in the research surrounding the TSP and have worked on it for decades. This extreme focus on this problem unfortunately shows in the book. While the authors mention the public appeal and the connection to other mathematical and computer scientific issues, they completely ignore the connections to industrial problems in their treatment. This is a central omission as much of the benefit gained in the industry due to computational methods has been discovered in connection with the TSP.

One may excuse this based on the inherent restrictions the authors have placed on themselves. Apart from a lengthy and very well written introduction to the TSP, the book is essentially a guided tour through the methods underlying the computer program concorde that the authors have written to solve TSP instances. Thus, the authors never claim to offer a comprehensive or even unbiased account of research on the TSP. To a reader familiar with the literature on the TSP, the book reads like an anthology of the papers written by the authors in the past few decades.

Inherent in such an anthology is the uncertainty as to the audience. Given that the topic is the collection of methods behind a specific computer program, one could imagine several possible audiences: (1) Students who want to enter the area, (2) expert researchers on the TSP, (3) experts in other areas who want to use TSP methods for their challenges and (4) practical computer scientists or programmers who want to learn the ins and outs of concorde, the computer program of the authors.

While the authors make an attempt at addressing each of these four groups, the book as a whole is written for none of these. It is far too advanced and not general or comprehensive enough for the student, it does not provide enough detail (especially proofs) for the TSP expert, it does not provide enough background and explanation for experts from other areas and it does not explain the implementation issues for the programmer (except in one instance where these are treated in unnecessarily exhaustive detail).

Even though the authors do not really say it, the book comes in three parts. The first 127 pages are an introduction to and history of the TSP. This part is well written and fairly complete. The exceptions to the completeness are primarily two: (1) The applications in the industry and (2) heuristic methods. The authors give a number of applications that are all academic research applications but not applications that impact the general public. Heuristic methods are mentioned but deserve far more treatment due to the fact that TSP instances in practice must inherently be solved heuristically because of the time consuming nature of the exact algorithms.

The second part occupies the next 296 pages and covers the branch-and-cut methods of the authors in great detail but not exhaustively as many proofs are missing and also the implementation issues. The authors make a solid attempt at explaining the concepts but this part is a difficult read. For the non-initiated, this part is impossible to understand without first reading more basic textbooks.

The authors advertise this part as being an explanation of the methods behind concorde. The reader who expects anything like a manual is sorely disappointed however. While some of the algorithms involved are presented in a typesetting style common for algorithms in computer science, it is quite intransparent how one might implement them in a programming language. The presentation therefore does not fulfill the usual standards set for the presentation of algorithms.

The next 64 pages cover one heuristic approach to finding good but not necessarily optimal traveling salesman tours. This heuristic, the Lin-Kernighan heuristic of Helsgaun, is very good at solving the TSP and deserves this discussion that is very well written and quite understandable. Being a problem specific algorithm however, this heuristic is not immediately capable of solving related problems such as a capacitated vehicle routing problem (CVRP) that assigns amounts of goods to be delivered at each of the TSP stops and restricts the maximum amount of goods that the traveling salesman can transport. It is these related problems that are important in practice and that are solved by other TSP heuristics that can be extended. The omission of these more general heuristics is a major downside of this book as it leaves the reader with no methods at all to attack problems that arise in real life.

The last 52 pages do not really make a fourth part but are essentially a conclusion and outlook. Here, the authors describe the computational experiments made and the successes they have been able to achieve with concorde. They also provide hints as to the future of TSP research and leave the reader with a nice challenge in the form of a few very large TSP instances yet to be solved.

This being said, the achievements of the authors are staggering. They have managed to set record after record with their methods and their computer program. Their last (and current) record is the solution to a TSP instance that includes 85,900 cities from around the globe that took about 286 days of processor time to achieve. Given the vast number of possible journeys (85,900! / 2), finding the one optimal journey in this large but realistically finite amount of time, is quite a result. More than that, the program

concorde can run on a parallel computer thus reducing the time that the user has to wait for the solution.

The authors provide concorde for free and open-source on the internet at http://www.tsp.gatech.edu/. The reader also has access to the TSPLIB, a collection of TSP instances, on the internet. This makes it possible for anyone to experiment with the TSP and concorde. The source code is commented but not documented. That is to say that most functions have a note next to them that says what that function does but there is no detailed program structure description that would tell the user how the functions fit together. Studying the source and experimenting with it is therefore possible but very time consuming.

It should be mentioned that this is a very notable exception in research – usually the researchers do not provide the world with all the resources of their work. Either the programs or the data behind the programs, if not both, are usually kept from the general public. The public, then, has no choice but to simply believe the results claimed. The authors of this book have chosen to share everything and allow us to check, to modify and to improve upon their work. This deserves our gratitude and respect and generates hope that, in the future, more of the leading researchers will share the full details of their work.

This brings us to the question of recommending the book. The general first part and the concluding two chapters of the book can be safely recommended to anyone to read. They are well written, interesting and challenging. Almost all of the actual content of these chapters can however easily be read on the internet in a variety of sites written about the TSP and the TSPLIB by various serious researchers; not least among them is the website of the authors (http://www.tsp.gatech.edu/).

The third part describing the Lin-Kernighan heuristic is readable and interesting but long-winded. This is however very good reading for anyone who wants to actually compute a very good but not necessarily optimal journey for a TSP with little runtime. However this is helpful only if the problem is precisely the TSP and not an extension of it.

The large second part is cryptic and cannot really be recommended except to a person who has a background in TSP-related algorithms and wants to know about recent advances but not about the proofs or the implementation issues.

Relative to similar literature, the book comes at a very modest price of 26.95 GBP or 45.00 USD. Given this price, the book is worth its money relative to similar books. However, given the complexity of the presentation in the major parts of the book and the critical omissions, it is my personal opinion that the book is not worth the investment in time needed to understand its content.