We can only see a short distance ahead,
but we can see plenty there
that needs to be done.
—Alan Turing, Mathematician

PROLOGUE

Imagine you visit a friend in the beautiful city of Augsburg in southern Germany. It is summer and you set yourselves the challenge to visit all 127 'Biergärten' (beer gardens) on a single day. (If you don't like beer, or if you don't have friends in Augsburg, consider visiting all coffee shops in Vancouver, Canada.) Can this be done? If so, which route should you take? Clearly, your chances of reaching your goal may depend on finding a short round trip that takes you to all 127 places.

As you arrive at Biergarten No. 42, your friend gives you the following puzzle, offering to pay for all your drinks if you can solve it before the night is over: 'Last week my friends Anne, Carl, Eva, Gustaf and I went out for dinner every night, Monday to Friday. I missed the meal on Friday because I was visiting my sister and her family. But otherwise, every one of us had selected a restaurant for a particular night and served as a host for that dinner. Overall, the following restaurants were selected: a French bistro, a sushi bar, a pizzeria, a Greek restaurant, and the Brauhaus. Eva took us out on Wednesday. The Friday dinner was at the Brauhaus. Carl, who doesn't eat sushi, was the first host. Gustaf had selected the bistro for the night before one of the friends took everyone to the pizzeria. Tell me, who selected which restaurant for which night?'

There are various approaches for solving these problems. Given the huge number of possible round trips through the Biergärten, or assignments of weekdays, hosts, and restaurants, systematic enumeration (i.e., trying out all possibilities) is probably not a realistic option. Some people would take a more sophisticated approach and eliminate certain assignments or partial tours through careful reasoning, while systematically searching over the remaining alternatives.

But most of us would probably take a rather different approach in practice: starting with a rough and somewhat arbitrary first guess, small changes are repeatedly performed on a given tour or assignment, with the goal of improving its quality or of getting closer to a feasible solution. This latter type of approach is known as *stochastic local search (SLS)* and plays a very important role in solving combinatorial problems like the ones illustrated above. (It may be noted that the logical puzzle and the shortest round trip problem can be seen as instances of the Propositional Satisfiability and Travelling Salesman Problems, which will be more formally introduced in Chapter 1 and used throughout this book.)

Why Stochastic Local Search?

There are many reasons for studying stochastic local search (SLS) methods. As illustrated above, SLS is closely related to a very natural approach in human problem solving. Many SLS methods are surprisingly simple, and the respective algorithms are rather easy to understand, communicate and implement. Yet, these algorithms can often solve computationally hard problems very effectively and robustly. SLS methods are also typically quite general and flexible. The same SLS methods have been found to work well for a broad range of different combinatorial problems, and existing algorithms can often be modified quite naturally and easily to solve variants of a given problem. This makes SLS methods particularly attractive for solving real-world problems, which are often not completely or correctly specified at the beginning of a project and may consequently undergo numerous revisions or modifications before all relevant aspects of the given application situation are captured.

Another reason for the popularity of SLS lies in the fact that this computational approach to problem solving facilitates an explorative approach to algorithm design. Furthermore, as we will discuss in more detail in Chapter 2, many prominent and successful SLS methods are inspired by natural phenomena, which gives them an additional intellectual appeal.

For these (and many other) reasons, SLS methods are among the most prominent and successful techniques for solving computationally hard problems in many areas of computer science (specifically artificial intelligence) and operations research; they are also widely used for solving combinatorial problems in other disciplines, including engineering, physics, management science and bioinformatics.

The academic interest in SLS methods can be traced back to the beginnings of modern computing. In operations research, local search algorithms were developed and described in the 1950s, and in artificial intelligence, SLS methods have been studied since the early days of the field, in the 1960s. To date, the study

of SLS algorithms falls into the intersection of algorithmics, statistics, artificial intelligence, operations research, and numerous application areas. At the same time, SLS methods play a prominent role in these fields and are rapidly becoming part of the respective mainstream academic curricula.

About this Book

'Stochastic Local Search: Foundations and Applications' was primarily written for researchers, students and practitioners with an interest in efficient heuristic methods for solving hard combinatorial problems. In particular, it is geared towards academic and industry researchers in computer science, artificial intelligence, operations research, and engineering, as an introduction to and overview of the field or as a reference text; towards graduate students in computer science, operations research, mathematics, or engineering, as well as towards senior undergraduate students with some background in computer science and mathematics, as primary or supplementary text for a course, or for self-study; and towards practitioners, who need to solve combinatorial problems for practical applications, as a reference text or as an introduction to and overview of the field.

The main goal of this book is to provide its readers with important components of a scientific approach to the design and application of SLS methods, and to present them with a broad, yet detailed view on the general concepts and specific instances of SLS methods, including aspects of their development, analysis, and application. More specifically, we aim to give our readers access to detailed knowledge on the most prominent and successful SLS techniques; to facilitate an understanding of the relationships, the characteristic similarities and differences between existing methods; to introduce and discuss basic and advanced aspects of the empirical analysis of SLS algorithms; and to give hands-on knowledge on the application of some of the most widely used SLS methods to a variety of combinatorial problems.

Stochastic search algorithms are being studied by a large number of researchers from different communities, many of which have quite different views on the topic or specific aspects of it. While striving for a balanced and objective presentation, this book provides a view on stochastic local search that is based on our background and experience. This is reflected, for instance, in the specific choice of our formal definition of stochastic local search (Chapter 1), in the GLSM model for hybrid SLS methods (Chapter 3), the extensive and in-depth coverage of empirical analysis and search space structure (Chapters 4 and 5), as well as in the selection of algorithms and problems we cover in varying degree of detail (particularly in Chapters 9 and 10). There are rational reasons for most

– if not all – of these choices; nevertheless, in many cases, equally defensible alternative decisions could have been made.

Clearly, some topics would benefit from broader and deeper coverage. However, even relatively large book projects are subject to certain resource limitations in both time and space, and it is our hope that our choices of the material and its presentation will make this book useful for the previously stated purposes.

Structure and Supplementary Materials

The main body of this book consists of two parts. **Part 1**, which comprises Chapters 1 to 5, covers the foundations of the study of stochastic local search algorithms, including:

- fundamental concepts, definitions, and terminology (Chapter 1),
- an introduction to a broad range of important SLS methods and their most relevant variants (Chapter 2),
- a conceptual and formal model that facilitates the development and understanding of hybrid SLS methods (Chapter 3),
- a methodical approach for the empirical analysis of SLS methods and other randomised algorithms (Chapter 4), and
- features and properties of the spaces searched by SLS algorithms and their impact on SLS behaviour (Chapter 5).

The material from the first two chapters provides the basis for all other aspects of SLS algorithms covered in this book; Chapters 1 and 2 should therefore be read before any other chapters and in their natural sequence. Chapters 3, 4, and 5 are quite independent from each other and expand the foundations of SLS in different directions. Chapter 3 complements Chapter 2; since it discusses some of the more complex SLS methods in a different light, it can be very useful for reviewing and deepening the understanding of these practically very relevant methods. The scope of Chapter 4 extends substantially beyond the empirical analysis of SLS algorithms; although most of the material covered in the subsequent chapters does not directly depend on the concepts and methods from Chapter 4, we strongly believe that anyone involved in the design and application of SLS algorithms should be familiar at least with the basic issues and approaches discussed there. Chapter 5 in some sense covers the most advanced material presented in this book; it should be useful to readers interested in a deeper knowledge of the

factors and reasons underlying SLS behaviour and performance, but reading it is not a prerequisite to understanding any of the material covered in the other chapters.

Part 2 comprises Chapters 6 to 10, which present, in varying degree of scope and detail, SLS algorithms for a number of well-known and widely studied combinatorial problems. Except for Chapter 7, which should be read after Chapter 6 since it builds on much of the material covered there, all chapters of this second part are basically independent of each other and can be studied in any combination and order. Chapters 6 to 8 provide a reasonable coverage of the most prominent and successful SLS methods for the respective problems and discuss the respective algorithms in a relatively detailed way. Chapters 9 and 10 are of a more introductory nature; their focus lies on a small number of SLS algorithms for the respective combinatorial problems that have been selected primarily based on their performance and general interest. In particular, the five main sections of Chapter 10 are independent of each other and can be studied in any combination and order.

'In Depth' Sections. Additional, clearly marked 'In Depth' sections are included in various chapters. These provide additional material that expands or complements the main body of the respective chapter, often at a more technical or detailed level. These sections are generally not required for understanding the main text, but in many cases they should be helpful for obtaining a deeper understanding of important concepts and issues.

Further Readings. Towards the end of each chapter, a 'Further Readings and Related Work' section provides additional references and pointers to literature on related topics. In the case of subjects for which there is a large body of literature, these represent only a small selection of references deemed especially relevant and/or accessible by the authors. These references should provide good starting points for the reader interested in a broader and deeper knowledge of the respective topic.

Chapter Summaries. Each chapter closes with a summary section that briefly reviews the most relevant concepts and ideas covered in the respective chapter. The purpose of this summary is to provide the reader with a high-level overview of the material presented in the chapter, and to point out connections (and differences) between the respective concepts and approaches. Together with the chapter introductions and exercises, these summaries facilitate rapid reviewing of previously studied or known material.

Exercises. Each chapter is accompanied by a collection of exercises, classified according to their degree of difficulty as 'easy', 'medium' and 'hard'. This classification is only approximate and does not necessarily reflect the anticipated amount of time needed for producing a solution; although an exercise marked as 'easy' may be relatively straightforward to solve, it may still require a substantial amount of time until the details of the solution are worked out and written down. The exercises cover the material presented in the respective chapter and are intended to facilitate a deeper understanding of the subject matter. They include theoretical questions as well as hands-on implementation and experimentation exercises.

References and Bibliography. References to the technical and research literature are provided throughout the book, particularly in the previously mentioned 'Further Readings and Related Work' sections. These give rise to an extensive bibliography that covers much of the most relevant literature on SLS algorithms and related topics, with a particular emphasis on recent publications.

Glossary and Index. The glossary contains brief explanations of important technical terms useful throughout the book. In conjunction with the extensive and thoroughly compiled index, the glossary particularly facilitates using this book as a reference book or for self-study.

Webpage and Supplementary Materials. Supplementary materials are provided from the book webpage at `www.sls-book.net`. These include slide sets that may be useful in the context of courses that use the book as a primary or supplementary text (see also Section 'Suggested Uses' below), as well as reference implementations of some of the SLS algorithms discussed in this book (needed for some of the hands-on exercises and useful for further practical experience) and some educational tools, for example, for the empirical analysis of SLS behaviour.

Suggested Uses

This book was designed for various types of uses. As a whole, it is intended to be used as a reference book for researchers and practitioners or as the primary text for a specialised graduate or upper-level undergraduate course on stochastic search algorithms; furthermore, parts of it can be used as primary reading or supplementary material for modules of more general courses in artificial intelligence, algorithms, operations research, combinatorial problem solving, empirical methods in computer science, etc. The following specific suggestions reflect our own experience, including the use of parts of this book by students, researchers, and

course instructors at the University of British Columbia (Vancouver, Canada) and Darmstadt University of Technology (Darmstadt, Germany).

General introduction to SLS methods, particularly for self-study. Chapters 1 and 2; Sections 3.1 to 3.3 and 3.6; Sections 4.1 to 4.3 and 4.6; Section 5.8; any one or two sections from Chapter 10. For more advanced self-study, the remaining materials can be added as desired; particularly the remaining sections of Chapters 4 as well as Chapters 6 and 8 are highly recommended.

Graduate Course on SLS methods/stochastic search. Chapters 1 and 2; Sections 3.1 to 3.3 and 3.6; Chapter 4; Sections 5.1 to 5.3 and 5.8; Chapters 6 and 7 without the sections on CSP and MAX-CSP; Chapter 8; and any two sections from Chapter 10. Depending on the precise format, focus and level of the course, this selection may be expanded in various ways, for example, by additionally covering Section 9.1 and any one other section from Chapter 9. For a general course on stochastic search methods, an additional module on randomised systematic search algorithms should be included (a sample set of slides for such a module is available from `www.sls-book.net`).

SLS Module(s) in a general AI course. Parts of Chapters 1 and 2; Sections 3.1 to 3.3 and 3.6; Sections 4.1 to 4.3 and 4.6; parts of Chapter 6; and possibly parts of Chapters 8, 9, or 10. The selections from Chapters 1, 2, 6 and 8 to 10 will naturally be based on the prerequisite knowledge of the students as well as the format, level and other modules of the course. A minimal subset for a module of about two lectures in an undergraduate course would mainly take parts of Chapters 1 and 2 and illustrate the working principles of SLS methods using example applications described in Part 2.

SLS Module(s) in a general algorithms course. Parts of Chapters 1 and 2; Sections 3.1 to 3.3 and 3.6; Sections 4.1 to 4.3 and 4.6; Sections 5.1 to 5.3 and 5.8; parts of Chapters 6 and 8; and possibly one or more sections from Chapter 10. The precise balance between these components will naturally depend on the exact nature of the course, particularly on its focus on theoretical or practical aspects of problem solving. In the context of strongly practically oriented algorithms courses, the in-depth sections in Chapters 4, 6 and 8 may be of particular interest.

SLS Module(s) in a discrete optimisation course. Parts of Chapters 1 and 2; Sections 3.1 to 3.3 and 3.6; Chapter 4; parts of Chapter 8 and 9; and any one or two sections from Chapter 10. Additional material, particularly from Chapters 6 and 7, can be used to further expand and complement this selection.

Parts of this book can also be used as primary or supplementary material for specialised graduate courses on SAT, CSP, TSP, scheduling and empirical methods in computing.

The Making of SLS:FA

The process of creating this book is in many ways related to the subject material discussed therein. Not unlike the fundamental approach of local search, it involved navigating a huge space of possibilities in an iterative manner. This process was initiated in 1998, when both, H. H. and T. S. were finishing their Ph.D. theses at the Computer Science Department of Darmstadt University of Technology, and the idea of combining materials from both theses into a comprehensive book on Stochastic Local Search first arose. Five years and about 650 pages later, we reached the end of this search trajectory. The result of a myriad of construction, perturbation and evaluation steps is this book. Interestingly and perhaps not too surprisingly, both, the writing process and its end result turned out to be very different from what we had originally imagined.

Although it would be hard to precisely define the objective(s) being optimised through the writing process, it took us through many situations that closely resemble those of a stochastic local search algorithm trying to solve a challenging instance of a hard combinatorial problem. There were phases of rapid progress and stagnation; we encountered (and overcame) numerous local minima; and along the way, we had to make many decisions based on very limited local information, various forms of heuristic guidance, and some degree of experience.

Random, or at least completely unforeseen and unpredictable, factors played a large role in this local search process. Rather trivial sources of randomness, such as hardware and software glitches, were complemented by more fundamental stochastic influences, such as the random thoughts and ideas that on warm summer nights seem to preferably lurk around the Biergärten, always looking for a receptive mind, or the random person sticking their head into the office door, causing the more organised ideas to fly apart in a hurry. Without these random influences, and the circumstances conducive to them, this book could not have been created in its present form.

At the same time, this book has been shaped by many other factors and influences. These include the places and circumstances under which part of the work was done. (Some of the more interesting places where parts of the book have been written include a log cabin on Sechelt Inlet, the beautiful and tranquil Nitobe Garden, a grassy spot near the top of Whyte Islet in Howe Sound, and the wild and remote inlets of the Pacific Northwest, onboard the Nautilus Explorer.) More importantly, they include a huge and diverse amount of interaction with

friends and family, mentors, colleagues, students and our publishers, who provided crucial guidance, diversification, evaluation and general support. Finally, especially during the final phase of the process, our work on this book was largely driven by Hofstadter's Law: 'It always takes longer than you expect, even when you take into account Hofstadter's Law.' [Hofstadter, 1979], the significance and effects of which can hardly be overestimated.

As a consequence, it would be foolish to believe that our stochastic local search process has led us into a global optimum. However, we feel that, largely thanks to the previously mentioned factors and influences, in the process of creating this book we managed to avoid and escape from many low-quality local optima, and achieved an end result that we hope will be useful to those who study it. In this context, we are deeply grateful towards those who contributed directly and indirectly to this work, and who provided us with guidance and support in our local — and global — search.

High-level guidance is of central importance in any effective search process; in our case, there are several people who played a key role in shaping our approach to scientific research and who provided crucial support during various stages of our academic careers. First and foremost, we thank Wolfgang Bibel, our former advisor and 'Doktorvater', for providing a highly supportive and stimulating academic environment in which we could freely pursue our research interests, and whose encouragement and substantial support was highly significant in getting this project underway. Furthermore, H. H. gratefully acknowledges the ongoing and invaluable support from his academic mentors and colleagues, Alan Mackworth and Anne Condon, who also played an important role during the early stages of writing this book. T. S. would especially like to thank Marco Dorigo for the pleasure of joint research and for his support in many senses.

On the other side, we have received more specific guidance on the contents of this book from a number of colleagues, students and fellow SLS researchers. Their detailed comments led to improvements in various parts of this book and helped to significantly reduce the number of errors. (Obviously, the responsibility for those errors that we managed to hide well enough to escape their vigilance rests solely with us.) In this context, we especially thank (in alphabetical order) Markus Aderhold, Christian Blum, Marco Chiarandini, Anne Condon, Irina Dumitrescu, Frank Hutter, David Johnson, Olivier Martin, Luis Paquete, Marco Pranzo, Tommaso Schiavinotto, Kevin Smyth, Dan Tulpan and Maxwell Young. We also acknowledge helpful comments by Craig Boutilier, Rina Dechter, Jin-Kao Hao, Keld Helsgaun, Kalev Kask, Henry Kautz, Janek Klawe, Lucas Lessing, Elena Marchiori, David Poole, Rubén Ruiz García, Alena Shmygelska and Dave Tompkins. Special thanks go to David Woodruff, Toby Walsh, Celso Ribeiro and Peter Merz, whose detailed comments provided valuable guidance in improving the presentation of our work.

In addition, we gratefully acknowledge the interesting and stimulating discussions on the topics of this book that we shared with many of our co-authors, colleagues, students and fellow researchers at TUD and UBC, as well as at conferences, workshops, tutorials and seminars. It is their encouragement, enthusiasm and continuing interest that provided much of the background and motivation for this work.

The staff at Morgan Kaufmann, Elsevier, Kolam and Dartmouth Publishing have been instrumental in the realisation of this book in many ways; we deeply appreciate their expertise and friendly support throughout the various stages of this project. We are particularly grateful to Denise Penrose, Senior Editor at Morgan Kaufmann, whose enthusiasm for this project and patience in dealing with the adverse effects of Hofstadter's Law (as well as with her authors' more peculiar wishes and ideas) played a key role in creating this book. Simon Crump, Publishing Services Manager at Elsevier, and Jamey Stegmaier, Project Manager at Kolam USA, have been similarly instrumental during the production stages, and we gratefully acknowledge their help and support. We also thank Jessica Meehan and her team at Dartmouth Publishing, who produced many of the figures, as well as Lori Newhouse and Calum Ross for copyediting and proofreading the book, and Robert Swanson for creating the index. Many thanks also to Emilia Thiuri and Valerie Witte, for their help during the draft stages, throughout the reviewing process and during production, and to Brian Grimm, marketing manager at Morgan Kaufmann, for substantially increasing the visibility of our work. H. H. also wishes to thank Valerie McRae for her help with proofreading the manuscript in various draft stages, and for much appreciated moral and administrative support.

Finally, we thank our families who provided the stable and stimulating environment that formed the starting point of our personal and intellectual development, and who shape and accompany the trajectories of our lives in a unique and special way. H. H. expresses his deepest gratitude to Sonia and Jehannine for being his partners in adventure, joy, and sorrow, and his parents, siblings and extended family for their affection and diversifying influence. T. S. especially thanks his wife Maria José for sharing her life with him, Alexander for all his curiosity and love, and his parents for their continuous care and support.

This book has been shaped by many factors and influences, but first and foremost it is the product of our joint research interests and activities, which co-evolved over the past seven years into an immensely fruitful and satisfying collaboration and, more importantly, into a close friendship.