Prologue

Today, stochastic local search (SLS) algorithms belong to the standard methods for solving hard combinatorial problems from various areas of Artificial Intelligence (AI) and Operations Research (OR). Some of the most successful and powerful algorithms for prominent problems like SAT (the Satisfiability Problem in Propositional Logic), CSP (the Constraint Satisfaction Problem), TSP (the Traveling Salesperson Problem), or QAP (the Quadratic Assignment Problem), are based on stochastic local search. Many stochastic local search methods like stochastic hill-climbing, Tabu Search, and Simulated Annealing can be generically applied to a broad range of combinatorial decision and optimisation problems, including problems from practically important application areas, like planning and scheduling. Other prominent and successful algorithmic schemes, like Evolutionary Algorithms or Ant Colony Optimisation, are also based on the concept of stochastic local search.

In AI, the success story of stochastic local search is tightly linked to its application to SAT. SAT is a conceptually simple combinatorial decision problem which plays a prominent role in complexity theory, because of its status as *the* prototypical $\mathcal{NP}$-complete problem. All $\mathcal{NP}$-complete problems can be polynomially translated into SAT – these reductions are often used in the context of $\mathcal{NP}$-completeness proofs. However, recent successes in efficiently solving SAT-encoded planning problems using modern SLS algorithms suggest that SAT-reductions might provide a practically viable way for solving hard combinatorial problems from application domains. SAT has a number of advantages over other $\mathcal{NP}$-complete problems. Its conceptual simplicity facilitates the development and implementation of SLS algorithms, including hardware realisations. Furthermore, SAT encodings are often well-known and easily realised. Finally, SLS algorithms for SAT can typically be generalised to CSP, which allows somewhat more concise and natural representations, in a rather straightforward way. One of the big advantages of SLS algorithms is the fact that, different from most systematic search methods, the algorithms for solving decision problems can equally well be used to solve the related optimisation problems.

Considering these arguments, the research presented in this thesis focusses mainly on SLS algorithms for SAT. However, most of the methods and models we developed and used are not restricted to SAT, but merely applied to SAT as an exemplary problem domain. For this reason, two of the main contributions of this work, the novel methodology for empirically evaluating Las Vegas algorithms (Chapter 2 and the GLSM model for local search (Chapter 3), are presented in a more general context before they are applied to

SAT. Interestingly, we found substantial evidence that the characterisation results regarding the behaviour of some prominent SLS algorithms for SAT (Chapter 5) also apply to SLS algorithms for other combinatorial problems, like CSP.

Since the early 1990s, when simple SLS algorithms where shown to outperform the best systematic algorithms for SAT on a number of domains, SLS algorithms have become increasingly popular in AI. The general interest in SLS algorithms is documented by a growing body of research. Studies on SLS algorithms play an important role at all major AI conferences, such as the International Joint Conferences on Artificial Intelligence (IJCAI), the European Conferences on Artificial Intelligence (ECAI), and the conferences of the American Association for Artificial Intelligence (AAAI). Furthermore, research on SLS algorithms is regularly presented at more specialised AI conferences, such as the International Conferences on Principles and Practice of Constraint Programming (CP) and the Conferences on Uncertainty in Artificial Intelligence (UAI). There is also a substantial number of articles on SLS algorithms in prestigious journals, like the Artificial Intelligence Journal or the Journal on Artificial Intelligence Research. SLS algorithms are also prominent in the Operations Research community, which began to study and apply them extensively before they became as popular in AI as they are today. Generally, stochastic local search can be considered as a young, but prominent and rapidly evolving research area in computer science in general and within AI in particular.

The research described in this thesis has been done between spring 1996 and autumn 1998 at the Computer Science Department of the Darmstadt University of Technology (Technische Universität Darmstadt, TUD) in Darmstadt, Germany. This work is based on my earlier experience with local search algorithms for SAT which I acquired during my M.Sc. studies (German: *Diplomstudiengang*) and builds on research issues which evolved in the context of my master's thesis (German: *Diplomarbeit*). A significant part of the work was done during a one-year research visit at the Laboratory for Computational Intelligence at the Computer Science Department of the University of British Columbia (UBC) in Vancouver, Canada which was made possible by a scholarship from the German National Merit Foundation (Studienstiftung des Deutschen Volkes).

The results of this Ph.D. thesis have partly been presented at international conferences and in research colloquia at TUD and UBC. Furthermore, some of the methods and results presented in Chapters 2 and 5, although original work of the author, appeared in joint publications with Thomas Stützle (TUD). Specifically, a part of Chapter 2 was presented at the 14th Conference for Uncertainty in Artifical Intelligence [HS98a]; some of the results presented in Chapter 5 are covered in a journal article which is currently under preparation and also appear in a poster presentation at the Fourth International Conference on Principles and Practice of Constraint Programming [HS98b].

This thesis was written in English to make its results directly accessible to the international scientific community and to allow its being reviewed by Prof. Bart Selman (Cornell University, USA), one of the leading experts in the field of stochastic local search.

Outline of the Thesis

The thesis is divided into nine major sections: this Prologue, Chapters 1–7, and an Epilogue. Each of the Chapters 1–7 begins with a short introduction and ends with a discussion of related work and a summary of its main results and contributions. Thus, the individual chapters, although their contents are closely related to each other, are self-contained to a certain degree. This structure reflects the fact that each chapter focusses on different research issues and methodological approaches. It has been chosen to facilitate selective reading of this thesis, as in the case of a reader who is mainly interested in the specific subtopic covered by one of the chapters.

In Chapter 1, we give the background and motivation for the work presented in this thesis. The stage is set by giving a brief introduction to AI problems, logic, and search; then, we discuss systematic versus local search and informally introduce stochastic local search methods. Next, we introduce and discuss the concept of generic problem solving using SAT as an intermediate problem domain, and stochastic local search methods for solving these problems. This is followed by an introduction into $\mathcal{NP}$-complete and $\mathcal{NP}$-hard problems and the methods used for solving these in practice. Based on formal definitions of the syntax and semantics of Propositional Logic, we then define the tightly related Satisfiability and Validity problems and discuss their complexity. Finally, we formally introduce the finite discrete Constraint Satisfaction problem and discuss its relation to SAT.

Chapter 2 covers the empirical methodology which is later used for evaluating stochastic local search algorithms and for characterising their behaviour. The previously existing methodology in this field suffers from several weaknesses, which can lead to misinterpretations and false conclusions. We discuss these weaknesses and develop a novel methodology for empirically evaluating stochastic local search algorithms in the more general framework of Las Vegas Algorithms. Our new approach allows a considerably more detailed analysis of these algorithms' behaviour without requiring additional overhead in data aquisition. At the same time, it is not restricted to local search algorithms or particular problem domains, but can be applied to any algorithm of Las Vegas type.

Chapter 3 introduces the GLSM (Generalised Local Search Machine) model, a novel formal framework for local search algorithms based on the concept of an extended non-deterministic finite state machine. This new model is based on the intuition that adequate local search algorithms are usually obtained by combining simple search strategies. The GLSM model allows the adequate representation of most modern SLS algorithms for SAT and other combinatorial problems; additionally, it provides conceptually simple ways for extending and re-combining these into new, hybrid local search schemes and facilitates their implementation. In the context of this work, GLSM models are used for formalising SLS algorithms for SAT (Chapter 4), for improving their behaviour (Chapter 5), and, to some extent, for analysing the search space structure their performance depends on (Chapter 6). However, these applications merely indicate the full power of the GLSM model and its extensions.

The main topic of Chapter 4 is a comprehensive comparison of several modern SLS algorithms for SAT, including some of the latest and most successful SAT algorithms. We formalise and discuss these algorithms, which are instances of the GSAT and WalkSAT architectures, within the GLSM framework and empirically evaluate their performance using a newly compiled suite of benchmark problems. This benchmark suite comprises test-sets sampled from random problem distributions as well as SAT-encoded instances from other problem domains. Our comparative performance study shows that while some algorithms are clearly dominated by others, there is no single best algorithm. Instead, there is a group of algorithms the performance ranking of which depends on the problem domain and instance size. Nevertheless, we show that essentially the same problem instances tend to be hard for all algorithms considered here. Therefore, it makes sense to assume an intrinsic hardness of problem instances. Surprisingly, the top-performing algorithms show an exceedingly poor performance for a number of problem instances from different domains. This indicates that the corresponding local search processes suffer from early stagnation, which severly compromises their performance, especially for hard structured problem instances.

In Chapter 5, we analyse SLS behaviour in more detail, using theoretical argument as well as empirical methodology for characterising the behaviour of several prominent algorithms of the WalkSAT and GSAT families. For the first time, we prove that while GWSAT is approximately complete (*i.e.*, for arbitrarily long runs, the probability of solving a soluble problem instance approaches one), this does not hold for modern WalkSAT variants, like WalkSAT+tabu, Novelty, and R-Novelty. This explains the stagnation behaviour observed for these algorithms in Chapter 4. Based on this result, we show how to extend these algorithms such that approximately complete behaviour is achieved. This leads to new variants of Novelty and R-Novelty for which we observe significantly improved performance. We also develop a functional characterisation of the empirically observed SLS behaviour. This model shows that when applied to hard problem instances from various domains, all the algorithms we analysed exhibit approximately exponential run-time distributions. This novel and rather surprising result has interesting implications regarding the effectivity of using random restarts and suggests a novel interpretation of SLS behaviour. As another consequence, these algorithms can be optimally parallelised by using multiple independent runs, a particular simple and attractive parallelisation scheme. We also discuss the more general portfolio approach and its potential benefits in the context of SLS algorithms for SAT.

In Chapter 6, the main objective is to improve our understanding of SLS behaviour by analysing features of the search spaces underlying the local search process. While some of our analyses refine and extend previous work, such as the influence of the number of solutions on SLS performance, others are based on features which have not been studied before, such as the distribution of objective function values across the search space or local minima branching along SLS trajectories. Our analysis shows that, while the number of solutions has a dominating influence on SLS performance, other factors, such as the standard deviation of the objective function and the branching of local minima have also significant effects. Regarding the structure of local minima regions, we show that these resemble brittle canyon structures rather than compact basins, which explains the effectivity of simple escape mechanims like random walk. Interestingly, the local minimum branching is considerably lower for SAT-encoded instances from other domains than for randomly generated instances. Our analysis does not explain all the observed performance differences, but it improves on the previous understanding of SLS behaviour. While we focus on SLS algorithms for SAT, the same empirical methodology can be used to study search space structure and its influence on SLS performance for other domains of combinatorial problems.

Chapter 7 provides an initial investigation of the influence of encoding strategies on search space structure and SLS performance. Based on two case studies, covering Constraint Satisfaction Problems (CSPs) and the Hamilton Circuit Problem (HCP), we exemplify how some of the methods developed before (in particular, in Chapters 4–6) can be applied in this context. We study different SAT-encodings for CSP and HCP and analyse the differences between compact encodings, (which minimise the search space size) and sparse encodings, as well as the effects of eliminating symmetric solutions. Furthermore, for CSP, we compare the performance of state-of-the-art SLS algorithms for SAT applied to SAT-encoded CSP instances with the performance of conceptually related SLS algorithms for CSP, which are directly applied to the un-encoded problem instances. While the scope of our investigation is too limited to produce definitive answers to the problems we address, our results suggest some testable hypotheses on the effectivity of encoding strategies and the generic problem solving approach from Chapter 1. Furthermore, our study demonstrates how the techniques developed for the evaluation of SLS performance and the analysis of search space structure can be used for studying the impact of SAT-encodings on search space structure and SLS behaviour.

The work concludes with an Epilogue, where we summarise its main contributions and point out open questions as well as directions for further research.

Acknowledgements

This work would have been impossible without the support of many people. First and foremost, I would like to thank my family, especially my parents, my sister, and my brother, who sent me on my way and provided a stable and stimulating environment for my personal and intellectual development. Furthermore, I gratefully acknowledge the constant support of Sonia M. Ziesche, my loving and caring companion for many years now, who influenced my life and work in many ways, and whose commenting and proof-reading of this thesis was an important contribution to its final form. I am greatly indepted to my friend and colleague Thomas Stützle, whom I would like to thank for the many fruitful discussions, the joint work on research projects, and the fun we had while I was working on my Ph.D. project.

I would like to thank Prof. Wolfgang Bibel, my thesis supervisor, for his generous support and valuable advice during my work on this thesis. I am very grateful that I had the opportunity to work as a member of his Intellectics Group at TUD for the past few years; during this time, I enjoyed great freedom in pursuing my research interests and gathered a wealth of invaluable experience in many aspects of academic life. I also thank the members (and ex-members) of the Intellectics Group, in particular Maria Tiedemann, Steffen Hölldobler, Antje Strohmeier, Olaf Steinmann, Thomas Rath, and Stephan Schmitt, for their support and many interesting discussions.

Furthermore, I would like to thank David Poole, Craig Boutillier, Ronen Brafman, Alan Mackworth, and all the members of the Laboratory for Computational Intelligence for their hospitality, their support and the many interesting and helpful discussions we had during my stay at UBC. My special thanks to Valerie McRae for helping me find my way at UBC and in Vancouver, for sharing many thoughts and impressions, and for bringing some sunshine in my day when the skies were clouded.

I am also indepted to Prof. Hermann K.-G. Walter, who supported my work on the SALIERI Computer Music Project during the past years. Although being not directly related to this thesis, these activities have been very important for me in many ways. I would also like to thank my colleagues and friends, Jürgen Kilian and Kai Renz from the SALIERI Project Group, as well as Keith A. Hamel from UBC for the exciting joint work and the good times we had.

Finally, I gratefully acknowledge the help and support of Bart Selman, David McAllester, and Henry Kautz, who inspired my work with their research, generously shared their knowledge and resources, and helped to clarify and improve my ideas and presentation. I would also like to thank the German National Merit Foundation (Studienstiftung des deutschen Volkes), who sponsored my stay at UBC and, through their summer universities, provided many unique opportunities for meeting interesting people, as well as for forming and exchanging ideas. In this context, I would like to thank Prof. Manfred Kluge, Prof. Friedrich Beck, Gerhard Vöpel, Marius Böger, and Dr. Max Brocker for their help and support.

As undoubtedly, this list of acknowledgements is essentially incomplete, I would like to thank all the people who supported my work in many different ways and accompanied me on my way thus far.