

Symmetry-Breaking and Local Search: A Case Study

Andrea Roli

Dipartimento di Scienze
Università degli Studi "G.D'Annunzio"
Pescara – Italy
a.roli@unich.it

Abstract

Symmetry-breaking has been proved to be very effective when combined with complete solvers. Conversely, it has been conjectured that the use of symmetry-breaking constraints has negative effect on local search-based solvers. This work presents an attempt to model the effect of symmetry-breaking on the search landscape explored by local search. The results, on the one hand, exclude that symmetry-breaking constraints negatively affect the topology of the search space. On the other hand, they strongly suggest that symmetry-breaking perturbs the configuration of local and global optima basins of attraction, making global optima more difficult to be reached.

1 Introduction

Symmetry-breaking has been proved to be very effective when combined with complete solvers [Crawford *et al.*, 1996; Puget, 2002]. This can be explained by observing that symmetry-breaking constraints considerably reduce the search space. Nevertheless, the use of symmetry-breaking constraints seem to have opposite effect on local search-based solvers, despite the search space reduction. In [Prestwich, 2001; 2002] some examples of this phenomenon are reported. When the problem is modeled with symmetry-breaking constraints, the search cost¹ is higher than the one corresponding to the model with symmetries. In [Prestwich, 2002], and also in [Ebner *et al.*, 2001] in a related context, it is suggested to use models that maximize the number of symmetries, contrarily to the case of complete solvers.

An important point to investigate is *why*, despite the search space reduction – that is often dramatic – local search-based techniques seem to be penalized by the introduction of symmetry-breaking constraints. Simplifying, the reasons why this happens can be either that (i) symmetry-breaking constraints modify the topology of the search landscape (e.g., by making the space not connected), or that (ii) they reduce the basins of attraction of global optima, or both.

In this paper, we introduce a model to study the effect of symmetry-breaking constraints on the search landscape ex-

plored by local search. This enables us to provide an abstraction of the system we want to investigate and to formulate general hypotheses, not bound to a particular problem. In the following sections, we will formulate more formally the search process performed by local search-based techniques and the conjectures we want to verify/falsify². The model we present introduces some simplifications and should be considered as a case study on the topic.

The structure of the paper is the following. In section 2, we define the model we use, along with the conjectures to be tested. Section 3 reports the experimental results concerning the previously defined hypotheses. Finally, we conclude in section 4 with a brief discussion and an outlook to the future.

2 A model for symmetry-breaking and local search

The aim of this section is to provide a model of the phenomenon we want to study. First of all, we will define the execution of a local search algorithm as a ‘walk’ along a graph. This enables us to outline important features of the search space and to relate them with the symmetry-breaking constraints. Then, we will introduce the kind of symmetries and the class of instances we consider.

Local search can be used to attack both constraint satisfaction problems and combinatorial optimization problems. The main operative difference is that, in the former case, the algorithm succeeds only if it finds a feasible solution to the original problem, which means that the local search algorithm must find a global optimum in the search landscape. In the latter case, the goal is usually to find, at least, a near-optimal solution. In the following, global optima can be considered either solutions better than any other solution or solutions belonging to the set of ‘acceptable solutions’, i.e., solutions considered good to the application at hand.

2.1 Local search and search graph

The local search process can be viewed as the exploration of a landscape aimed at finding a global optimum.

The *search landscape* is defined by a triple:

$$\mathcal{L} = (S, \mathcal{N}, F)$$

²According to [Popper, 2002], scientific empirical theories can only be falsified.

¹Runtime and number of variable flips




Figure 1: Example of undirected graph representing a search landscape. Each node is associated with a solution s_i and its corresponding fitness value $F(s_i)$. Arcs represent transitions between states by means of φ . Undirected arcs correspond to symmetric neighborhood structures.

where:

- S is the set of solutions (or states);
- \mathcal{N} is the neighborhood function $\mathcal{N} : S \rightarrow 2^S$ that defines the neighborhood structure, by assigning to every $s \in S$ a set of states $\mathcal{N}(s) \subseteq S$.
- F is the objective function $F: S \rightarrow \mathbb{R}^+$.

The search landscape can be interpreted as a graph (see Fig. 1) in which nodes are solutions (labeled with their fitness value) and arcs represent the neighborhood relation between states.

The neighborhood function \mathcal{N} implicitly defines an operator φ which takes a state s_1 and transforms it into another state $s_2 \in \mathcal{N}(s_1)$. Conversely, given an operator φ , it is possible to define a neighborhood of a variable $s_1 \in S$: $\mathcal{N}_\varphi(s_1) = \{s_2 \in S \setminus \{s_1\} \mid s_2 \text{ can be obtained by one application of } \varphi \text{ on } s_1\}$

Usually, the operator is *symmetric*: if s_1 is a neighbor of s_2 then s_2 is a neighbor of s_1 . In a graph representation (such as the one depicted in Fig. 1) undirected arcs represent symmetric neighborhood structures. A desirable property of the neighborhood structure is to allow a path from every pair of nodes (i.e., the neighborhood is strongly optimally connected) or at least from any node to an optimum (i.e., the neighborhood is weakly optimally connected). Nevertheless, there are some exceptions of effective neighborhood structures which do not enjoy this property [Nowicki and Smutnicki, 1996].

The search process of local search methods can be seen as the evolution in (discrete) time of a discrete dynamical system [Bar-Yam, 1997; Devaney, 1989]. The algorithm starts from an initial state (the initial solution) and describes a trajectory in the state space, that is defined by the search graph. The system dynamics depends on the strategy used; simple algorithms generate a trajectory composed of two parts: a *transient* phase followed by an *attractor* (a fixed point, a cycle or a complex attractor). Algorithms with advanced strategies generate more complex trajectories which can not be subdivided

in those two phases. The characteristics of the trajectory outline the behavior of the algorithm and its effectiveness with respect to the instance it is tackling. It is worth underlining that the dynamics is the result of the combination of problem representation, algorithm and instance. In fact, the problem representation defines the search landscape (and, in particular, the neighborhood structure defines the *topology* of the search landscape); the algorithm describes the strategy used to explore the landscape and, finally, the actual search space characteristics are defined by the instance to be solved.

For instance, let us consider a deterministic version of the Iterative Improvement local search (DII). The trajectory starts from a point s_0 , exhaustively explores its neighborhood, picks the neighboring state s' with minimal objective function value³ and, if s' is better than s_0 , it moves from s_0 to s' . Then this process is repeated, until a minimum \hat{s} (either local or global) is found. The trajectory does not move further and we say that the system has reached a fixed point (\hat{s}).

The relations between local search-based algorithms and dynamical systems is beyond the scope of this paper. Nevertheless, in the following, we will make extensive use of the notion of *basin of attraction* (BA), that stems from dynamical system theory. We will initially consider the case of deterministic systems, then we will relax this hypothesis and extend the definition to stochastic systems.

Definition Given a deterministic algorithm \mathcal{A} , the basin of attraction $\mathcal{B}(\mathcal{A}|s)$ of a point s , is defined as the set of states that, taken as initial states, give origin to trajectories converging to point s . The cardinality of a basin of attraction represents its size (in this context, we always deal with finite spaces).

Given the set S^* of the global optima, the union of the BA of global optima $I^* = \bigcup_{i \in S^*} \mathcal{B}(\mathcal{A}|i)$ represents the set of desirable initial states of the search. Indeed, a search starting from $s \in I^*$ will eventually find an optimal solution. Since it is usually not possible to construct an initial solution that is guaranteed to be in I^* , the ratio $|I^*|/|S|$ can be taken as an indicator of the probability to find an optimal solution. On the extreme case, if we start from a random solution, the probability to find a global optimum is exactly $|I^*|/|S|$. Therefore, the higher this ratio, the higher the probability of success of the algorithm.

In the case of stochastic local search, we may define a probabilistic basin of attraction, as a generalization of the previous case.

Definition Given a (stochastic) algorithm \mathcal{A} , the basin of attraction $\mathcal{B}(\mathcal{A}|s;p^*)$ of a point s , is defined as the set of states that, taken as initial states, give origin to trajectories converging to point s with probability $p \geq p^*$. Also in this case, we define the union of the BA of global optima: $I^*(p) = \bigcup_{i \in S^*} \mathcal{B}(\mathcal{A}|i;p)$. For simplicity, in the following we will write $\mathcal{B}(s;p^*)$ instead of $\mathcal{B}(\mathcal{A}|s;p^*)$ when the algorithm involved is clear from the context.

This definition includes the previous one as a special case.

³Ties are broken by enforcing a lexicographic order of states.

Indeed, if $p^* = 1$ we are interested in finding the states that will eventually converge to s . It is also important to note that if $p_1 > p_2$, then $\mathcal{B}(s; p_1) \subseteq \mathcal{B}(s; p_2)$.

Given a local search algorithm \mathcal{A} , the topology and structure of the search landscape determine the effectiveness of \mathcal{A} . In particular, the reachability of a global optimum is the key issue. Therefore, the characteristics of the BA of optimal solutions are of dramatic importance. The graph properties that affect the structure of BA are (i) the graph topology and (ii) the graph shape, namely, the number and distribution of local optima, the auto-correlation of the landscape, the presence/absence of plateaus, etc⁴. Hence, the goal of the algorithm designer is to have an algorithm \mathcal{A} such that the resulting total BA of global optima $I^*(p)$ is as large as possible, given a reasonably high value of p . It is clear that this is an *a posteriori* property, since it is the result of the application of a particular algorithm to a particular problem instance. However, this property can be used to explain why an algorithm is or is not effective on a problem instance, or a class of instances.

Finally, we would like to remark that, while the search graph topology is only dependent on the neighborhood structure, the BA and other related landscape characteristics (hereinafter referred to as search graph *shape*) depend also on the particular algorithm used.

In the following, we give the definition of a class of instances characterized by controlled properties of the search graph. This will lead us to the study of the effect of symmetry-breaking constraints on the structure of the search graph and, consequently, on the behavior of local search algorithms.

2.2 A case study

In the following, we define the case study we analyze. We will first describe the search landscape that has some similarities with random landscapes and NK-landscapes [Kauffman, 1995; 1993]. Then, we will define the specific class of symmetries we consider in the model. Such a symmetry class is the one of permutations over problem variables.

The search graph

The class of problems we consider are defined over binary variables $x_i \in \{0, 1\}$, $i = 1, \dots, n$. Since our goal is to directly analyze the effect of symmetry-breaking constraints on the properties of the search graph, we directly define a model to construct a search landscape, abstracting from the specific problem we may deal with. Studies concerning random landscapes and NK-landscapes [Kauffman, 1995; 1993] can be found in the literature. In this work, we consider a landscape, in which the objective function is defined as follows. All the complete assignments are considered feasible (this is typically the case when local search is applied to problems defined on binary variables). Every assignment is given a value randomly chosen in a range $[r_1, r_2]$. Since usually the landscape has some degree of correlation, i.e., neighboring

solutions have objective values that are close, we introduce a distribution on the objective values such that some correlation is guaranteed. In practice, the distribution used is such that the higher the distance between two solutions, the larger the range of the difference of their objective values.

The topology of the graph is defined by the neighborhood structure. In our case, we use the neighborhood defined on unitary Hamming distance, that is the most used neighborhood for binary variables. Therefore, since n are the possible flips of an assignment, each node of the graph is connected with n other nodes. It is important to observe that this neighborhood structure generates a graph with a uniform degree (i.e., number of node edges) and this value is n . The impact of graph properties on system behavior has been recently received attention, as witnessed by the wide spectrum of publications on the subject [Watts and Strogatz, 1998; Adamic and Huberman, 2000].

Permutation symmetries

An important category of symmetries that can occur in combinatorial problems is the one of *permutation symmetries* [Gallian, 2001; Aloul *et al.*, 2002; Crawford *et al.*, 1996]. A permutation of a (finite) set Z is a function from Z to Z that is both injective and surjective. It is customary to use as set Z a (finite) set of naturals $\{1, 2, \dots, n\}$. A permutation can be expressed in the so called cycle notation, which is a composition of disjoint cycles of the form $(z_1 \dots z_m)$, $z_i \in \{1, 2, \dots, n\}$. For example, given $Z = \{1, 2, \dots, 6\}$, a permutation can be the composition of a 2-cycle and a 3-cycle such as $(1\ 2)(3\ 4\ 6)$, which means that 1 is mapped into 2 (and viceversa) and there is a cycle such that 3 is mapped into 4, 4 mapped into 6 and 6 into 3. Element 5 is mapped into itself. An important theorem [Gallian, 2001] states that every permutation in $\{1, 2, \dots, n\}$, $n > 1$, is either a 2-cycle or a product of 2-cycles⁵. Therefore, every possible permutation can be expressed as a composition of pairs $(z_i\ z_j)$, $i, j \in \{1, 2, \dots, n\}$.

We restrict our case study to permutation symmetries over a subset of problem variables. These symmetries can be expressed as a combination of 2-cycles. Moreover, we also add *phase shifts* in analogy with satisfiability problems. Phase shifts are such that exchanging a literal x_i with \bar{x}_i keeps the satisfying property of the assignments. We generated instances with n binary variables and m symmetries, each being either a 2-cycle or a phase shift. The m symmetries are randomly generated (without repetition), with the aim of covering a wide spectrum of cycle and phase shift combinations. The choice of random generation of cycles is motivated by the observation that, except for very specific cases occurring in structured instances, there is no regularity of permutation symmetries across the instances of a benchmark. The symmetries introduced by a 2-cycle $(i\ j)$ are simply cut by enforcing the constraint $x_i \leq x_j$ (we did not apply any reasoning to strengthen the constraints). Phase shifts (e.g., on variable x_l) are cut by posting the constraint $x_l = 0$. Symmetry-breaking constraints are combined and enforced in such a way that the resulting feasible solutions are the symmetry class representatives.

⁴We forward the interested reader to specific literature on the topic [Hordijk, 1996; Merz and Freisleben, 1999; Reeves, 1999; Stadler, 1996].

⁵The decomposition is not unique.

The effect of the symmetry-breaking constraints on the search graph can be directly analyzed by comparing the properties of the search graph with and without symmetry-breaking constraints. Since the search space is completely enumerated, only small size instances can be considered. In the next section, we present and discuss experimental results on the case study defined.

3 Experimental results

In [Prestwich, 2001; 2002] it has been conjectured that symmetry-breaking has a negative effect on local search-based techniques. From the standpoint of the model of local search process defined in section 2.1, we can formulate the hypothesis that symmetry-breaking constraints perturb the search graph in such a way that the algorithm effectiveness is reduced. The perturbation on the search graph can be of two kinds:

- topology perturbation
- *shape* perturbation (BA, local/global optima, etc.)

The negative effect could be originated from one or both the kinds of perturbation. In order to test this, we analyze the modification introduced in the search graph by symmetry-breaking constraints.

3.1 Topology perturbation

The topology of the search graph corresponding to the model with symmetries is highly regular (indeed, it is an hypercube) and each node has a degree equal to n (see an example in Fig.2). The graph is connected, therefore there is a path between any node and any global optimum. More important, given the structure of the graph, there are many alternative paths connecting every pair of nodes. Such a regular and redundant structure is particularly suitable for local search. ‘Irregular’ graph topologies, such as scale-free and small-world graphs [Watts and Strogatz, 1998; Strogatz, 1998; Barabasi, 2002], strongly affect the graph exploration process. As a consequence, it is natural to look for special topologies in graphs derived by models with symmetry-breaking constraints. An example of the resulting search graph is drawn in Fig.3, that is the counterpart of the one depicted in Fig.2. The outcome of our simulations, in which we tested graphs up to 10 variables⁶, can be summarized as follows. In the search graph related to the model with symmetry-breaking constraints, we observe that:

1. Node degree varies among nodes;
2. node degree frequency has a bell-shape, analogous to the Gaussian distribution characterizing random graphs (a typical case is reported in Fig.4);
3. the width of the bell-curve increases as the number of symmetries increases and it is independent of the relative number of 2-cycles and phase shifts.

Effect 1 is not surprising, since the constraints cut some parts of the search space, so they are very likely to introduce

⁶Since the analysis of the graph is complete, the instance size is strongly limited.




Figure 2: Search graph corresponding to an instance of size 5. The node degree is the same for every node and it is equal to 5.




Figure 3: Search graph corresponding to an instance of size 5, with the following symmetries: $(1\ 2)$, $(1\ 4)$ and a phase shift involving variable x_5 . The maximum node degree is 4 and the minimum is 2.

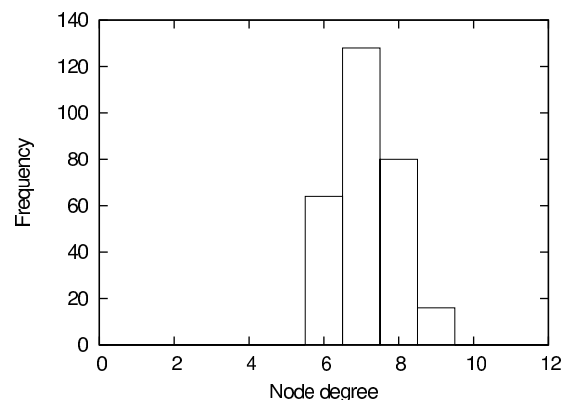


Figure 4: Node degree frequency of a search graph corresponding to an instance of size 10, with the following symmetries: $(8\ 10)$, $(6\ 8)$, $(3\ 4)$, $(6\ 10)$, $(7\ 9)$. The maximum node degree is 9 and the minimum is 6.

differences in the node degree. However, the second effect is extremely important, since it definitely exclude the possibility that symmetry-breaking constraints introduce a graph topology that negatively affect local search behavior – at least in our model. In fact, the node degree frequency is such that the reachability of global optima is not dramatically perturbed, as can be also confirmed by observing the properties of random graphs [Newman *et al.*, 2001]. In fact, such graphs are connected, the average path length connecting any pair of nodes is quite short, there are many alternative paths connecting any pair of nodes and there are no bottlenecks. Finally, effect 3 just shows that, even if the number of symmetries increases, the graph topology keeps its basic characteristics. It is important to note that the graph is still connected and there are no disconnected regions.

Concluding, we can definitely reject the conjecture that local search effectiveness and efficiency are reduced by topological modification of the search space.

3.2 Basin of attraction perturbation

The second way symmetry-breaking affects the search space is by perturbing its shape. In particular, we studied how the basins of attraction of global optima vary as a consequence of symmetry-breaking constraints. The primary effect of symmetry-breaking is to reduce the number of optimal solutions, which seems to negatively influence the efficiency of local search [Clark *et al.*, 1996]. Nevertheless, this effect should be counterbalanced by the parallel reduction of local minima⁷.

Our experiments show that the relative size of basins of attraction of global optima is reduced in the model with symmetry-breaking constraints. We compared the relative size of global optima BA of instances of different size and different number of symmetries. Table 1 reports the statistics concerning the search graph with and without symmetries (results are averaged over 20 instances). The first two columns of the table report the number of variables and the number of symmetries, respectively. The third column is computed as follows. The ratio of the total size of the global optima BA and the number of feasible states is computed for both the model with symmetry-breaking constraints ($\Omega_s^* = |I_s^*|/|S_s|$) and without them ($\Omega^* = |I^*|/|S|$). The reported value is the average of the ratio Ω_s^*/Ω^* (along with the standard deviation in the fourth column). Very important are columns five and six, in which we reported the percentage of times Ω_s^* is smaller (col.5) or larger (col.6) than Ω^* . $\Omega_s^* < \Omega^*$ means that the fraction of nodes from which the algorithm can reach a global optimum is higher in the model without symmetry-breaking constraints. For example, consider the row corresponding to 10 variables and 1 symmetry: in the 60% of the instances, the percentage of states leading to a global optimum is lower in the model with symmetry-breaking constraints, and in the 30% is higher (for the remaining 10% is equal). When the total size of global optima BA is higher in the model with symmetries, it means that symmetry-breaking

constraints reduce the probability of reaching the global optima (at least in the deterministic case). We can observe that, in most of the cases, $\Omega_s^* \leq \Omega^*$. Therefore, despite the search space reduction achieved by symmetry-breaking constraints, a deterministic local search, such as DII, has not a higher probability of reaching a global optimum. Indeed, in most of the cases the inequality is strict.

The analysis we made only considered the deterministic case, however the results are an important clue about the general effect of symmetry-breaking constraints also concerning the stochastic case. In fact, even stochastic local search methods have a strong greedy component, that characterizes DII. Therefore, even if they can move away from ‘wrong’ basins of attraction, their efficiency is lessened by the reduction of global optima BA.

In summary, the experimental results enable us to reject the hypothesis that symmetry-breaking affects local search by perturbing the topology of the search space. Furthermore, we have some interesting clues concerning the effect of symmetry-breaking on the shape of basins of attraction.

4 Discussion and future work

In this work, we have presented a model to study the effect of symmetry-breaking constraints on local search and we have exemplified this methodology on a case study. The search landscape explored by a local search algorithm can be seen as a labeled graph and the algorithm execution as a path on this graph. Moreover, since relevant properties of the search process can be modeled via dynamical system theory, concepts such as attractors and basin of attractions can be used to characterize the algorithm execution. This standpoint enables us to study the effect of symmetry-breaking constraints on local search by analyzing their impact on the search landscape and, in particular, on its topology and its shape. We have reported a preliminary study based on an example of symmetric problem, characterized by permutation symmetries. Results exclude that symmetry-breaking constraints perturb the search graph topology in such a way that local search is penalized (at least for this case study). We have also found some interesting results indicating that the global optima basin of attractions are reduced in the model with symmetry-breaking constraints. This may be the reason why these constraints have a negative effect on local search behavior.

The problem model we used, while quite general and plausible as an abstraction of a real problem, has some limits, though. First of all, since the instances are randomly generated and only expressed by the composition of 2-cycles and phase shifts, the model does not capture the generality of cases that can occur in real-world problems. Furthermore, the experiments should be extended to study in detail the structure of the search graph, for instance by considering the reduction of number of optimal solutions. We should also experience with real problems and larger size instances and taking into account diverse local search algorithms, also stochastic.

An important observation concerns the rejection of the hypothesis on the perturbation of the search graph topology. This result enables us to conjecture that robust and

⁷In some models, symmetry-breaking constraints introduce new local optima [Prestwich, 2002]. This is not the case for our model, that should be considered as a ‘best’ case.

adaptive local search methods, such as Variable Neighborhood Search⁸, might not be negatively affected by symmetry-breaking constraints.

The problem of dealing with symmetries when a local search technique is applied is of considerable relevance, since stochastic approximate algorithms are often used to attack real-world problems. Very few works in the literature tackled this issue and there is a lot of room for future research. Besides studies on models, practitioners may need guidelines on both modeling and algorithm design. For instance, it would be extremely interesting trying to apply local search on models with symmetry-breaking constraints and to use these constraints in a tabu search fashion or by relaxing them whenever a diversification step is required. Finally, it would be of practical relevance the discovery of criteria to decide when using a model with symmetry-breaking constraints, or – at the other extreme – a super-symmetric model.

Acknowledgments

I thank Michela Milano for discussions and suggestions on this work. I also thank the anonymous reviewers for helpful comments.

References

- [Adamic and Huberman, 2000] L. A. Adamic and B. A. Huberman. Power-law distribution of the world wide web. *Science*, 287, 2000.
- [Aloul *et al.*, 2002] F. A. Aloul, A. Ramani, I. L. Markov, and K. A. Sakallah. Solving difficult SAT instances in the presence of symmetry. In *SAT 2002 – Fifth International Symposium on the Theory and Applications of Satisfiability Testing*, Cincinnati, Ohio, USA, 2002.
- [Bar–Yam, 1997] Y. Bar–Yam. *Dynamics of Complex Systems*. Studies in nonlinearity. Addison–Wesley, 1997.
- [Barabasi, 2002] A.-L. Barabasi. *Linked: The New Science of Networks*. Perseus Books Group, 2002.
- [Blum and Roli, 2003] C. Blum and A. Roli. Metaheuristics in combinatorial optimization: Overview and conceptual comparison. *ACM Computing Surveys*, 35(3):268–308, September 2003.
- [Clark *et al.*, 1996] D. A. Clark, J. Frank, I. P. Gent, E. MacIntyre, N. Tomov, and T. Walsh. Local search and the number of solutions. In *Principle and Practice of Constraint Programming – CP96*, volume 1118 of *Lecture Notes in Computer Science*, pages 119–133. Springer-Verlag, 1996.
- [Crawford *et al.*, 1996] James Crawford, Matthew L. Ginsberg, Eugene Luck, and Amitabha Roy. Symmetry-breaking predicates for search problems. In *KR'96: Principles of Knowledge Representation and Reasoning*, pages 148–159. Morgan Kaufmann, San Francisco, California, 1996.
- [Devaney, 1989] R. L. Devaney. *An introduction to chaotic dynamical systems*. Addison–Wesley, second edition, 1989.
- [Ebner *et al.*, 2001] M. Ebner, M. Shackleton, and R. Shipman. How neutral networks influence evolvability. *Complexity*, 7(2):19–33, 2001.
- [Gallian, 2001] J. A. Gallian. *Contemporary Abstract Algebra*. Houghton Mifflin, 2001.
- [Hordijk, 1996] W. Hordijk. A measure of landscapes. *Evolutionary Computation*, 4(4):335–360, 1996.
- [Kauffman, 1993] S. A. Kauffman. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press, 1993.
- [Kauffman, 1995] S. A. Kauffman. *At home in the universe*. Oxford University Press, New York, 1995.
- [Merz and Freisleben, 1999] P. Merz and B. Freisleben. Fitness landscapes and memetic algorithm design. In David Corne, Marco Dorigo, and Fred Glover, editors, *New Ideas in Optimization*, Advanced topics in computer science series. McGraw-Hill, 1999.
- [Newman *et al.*, 2001] M.E.J. Newman, S.H. Strogatz, and D.J. Watts. Random graphs with arbitrary degree distributions and their applications. *Phys. Rev. E*, 64, 2001.
- [Nowicki and Smutnicki, 1996] E. Nowicki and C. Smutnicki. A fast taboo search algorithm for the job-shop problem. *Management Science*, 42(2):797–813, 1996.
- [Popper, 2002] K. Popper. *The Logic of Scientific Discovery*. Routledge, 2002. First edition in English 1977.
- [Prestwich, 2001] S. Prestwich. First-solution search with symmetry breaking and implied constraints. In *Proc. of CP'01 Workshop on Modelling and Problem Formulation*, 2001.
- [Prestwich, 2002] S. Prestwich. Supersymmetric modeling for local search. In *Proc. of SymCon'02 Workshop on Symmetry and Constraint Satisfaction Problems*, 2002.
- [Puget, 2002] J. F. Puget. Symmetry breaking revisited. In *Principle and Practice of Constraint Programming – CPO2*, volume 2470 of *Lecture Notes in Computer Science*, pages 446–451. Springer-Verlag, 2002.
- [Reeves, 1999] C. R. Reeves. Landscapes, operators and heuristic search. *Annals of Operations Research*, 86:473–490, 1999.
- [Stadler, 1996] P. F. Stadler. Landscapes and their correlation functions. *Journal of Mathematical Chemistry*, 20:1–45, 1996. Also available as SFI preprint 95-07-067.
- [Strogatz, 1998] S.H. Strogatz. Exploring complex networks. *Nature*, 410:268–276, 1998.
- [Watts and Strogatz, 1998] D.J. Watts and S.H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.

⁸For a survey on local search algorithms and metaheuristics in general, see [Blum and Roli, 2003].

Table 1: Statistics of the relative size of total basins of attraction of global optima. Statistics are computed over 20 randomly generated instances (see description of the model in section 2). $\Omega_s^* = |I_s^*|/|S_s|$ and $\Omega^* = |I^*|/|S|$, where I_s^* and S_s (resp. I^* and S) are the total size of global optima BA and the number of feasible states in the search space in the model with symmetry-breaking constraints (resp. without).

n	m	$\langle \Omega_s^*/\Omega^* \rangle$	$\sigma(\Omega_s^*/\Omega^*)$	frac. $\Omega_s^* < \Omega^*$	frac. $\Omega_s^* > \Omega^*$
5	1	0.98	0.03	0.50	0.05
5	2	1.00	0.09	0.45	0.15
5	3	0.97	0.05	0.45	0.05
5	4	1.01	0.11	0.20	0.30
6	1	1.00	0.08	0.50	0.20
6	2	0.96	0.08	0.45	0.15
6	3	1.02	0.11	0.25	0.35
6	4	1.01	0.15	0.45	0.20
7	1	0.99	0.04	0.40	0.15
7	2	0.98	0.08	0.60	0.10
7	3	0.98	0.06	0.65	0.25
7	4	1.00	0.13	0.60	0.30
7	5	0.96	0.14	0.55	0.15
8	1	0.99	0.02	0.50	0.15
8	2	0.99	0.05	0.60	0.35
8	3	1.00	0.12	0.50	0.35
8	4	0.97	0.06	0.75	0.10
8	5	1.12	0.70	0.55	0.15
8	6	0.96	0.08	0.50	0.35
9	1	1.01	0.03	0.40	0.35
9	2	0.99	0.06	0.55	0.25
9	3	0.99	0.07	0.60	0.35
9	4	0.97	0.07	0.55	0.20
9	5	0.96	0.18	0.65	0.20
9	6	1.01	0.12	0.55	0.35
9	7	0.99	0.09	0.50	0.35
9	8	1.01	0.20	0.55	0.35
9	9	1.03	0.19	0.45	0.40
10	1	0.99	0.03	0.60	0.30
10	2	0.99	0.03	0.50	0.40
10	3	0.97	0.05	0.80	0.20
10	4	0.98	0.06	0.65	0.35
10	5	0.98	0.09	0.55	0.40
10	6	1.02	0.16	0.55	0.25
10	7	1.12	0.28	0.30	0.60
10	8	0.95	0.14	0.50	0.25
10	9	1.16	0.71	0.65	0.30
10	10	1.00	0.15	0.60	0.30