

The Ising Model: Simple Evolutionary Algorithms as Adaptation Schemes

Patrick Briest, Dimo Brockhoff, Bastian Degener, Matthias Englert,
Christian Gunia, Oliver Heering, Thomas Jansen, Michael Leifhelm,
Kai Plociennik, Heiko Röglin, Andrea Schweer, Dirk Sudholt,
Stefan Tannenbaum, and Ingo Wegener*

FB Informatik, LS2, Univ. Dortmund, 44221 Dortmund,
Germany

<firstname.lastname>@uni-dortmund.de

Abstract. The investigation of evolutionary algorithms as adaptation schemes has a long history starting with Holland (1975). The Ising model from physics leads to a variety of different problem instances and it is interesting to investigate how simple evolutionary algorithms cope with these problems. A theoretical analysis is known only for the Ising model on the ring and partially for the Ising model on the two-dimensional torus. Here, the two-dimensional torus, the d -dimensional hypercube, and graphs consisting of two cliques connected by some bridges are investigated experimentally. Many hypotheses are confirmed by rigorous statistical tests.

1 Introduction

Holland (1975) has designed genetic algorithms (GAs) as adaptation systems. Today, evolutionary algorithms (EAs) and GAs are mainly applied as optimization algorithms. In order to understand how GAs and EAs work it makes sense to investigate their behavior on problems which are easy from the perspective of optimization but not so easy from the perspective of adaptation.

Naudts and Naudts (1998) have introduced the Ising model into this discussion. It is based on a model due to Ising (1925) to study the theory of ferromagnetism. In its most general form, the model consists of an undirected graph $G = (V, E)$ and a weight function $w: E \rightarrow \mathbb{R}$. Each vertex $i \in V$ has a positive or negative spin $s_i \in \{-1, +1\}$. The contribution of the edge $e = \{i, j\}$ equals $f_s(e) := s_i \cdot s_j \cdot w(e)$. The fitness $f(s)$ of the state s equals the sum of all $f_s(e)$, $e \in E$, and has to be maximized. This general problem is NP-hard and good heuristics have been designed by Pelikan and Goldberg (2003).

Here, we are interested in the adaptation capabilities of simple EAs and consider the simple case of $w(e) = 1$ for all $e \in E$. Then it is convenient to apply an

* Supported in part by the Deutsche Forschungsgemeinschaft (DFG) as part of the Collaborative Research Center “Computational Intelligence” (SFB 531).

affine transformation that leads to the state space $\{0, 1\}^n$ instead of $\{-1, +1\}^n$. The state s describes a coloring of G by the colors 0 and 1 and the fitness function counts the number of monochromatic edges, i. e., edges $\{i, j\}$ where $s_i = s_j$. This leads to a trivial optimization problem since the monochromatic colorings 0^n and 1^n are the only optimal colorings for connected graphs. Connected monochromatic subgraphs can be considered as building blocks (schemata of high fitness). The difficulty is the property of spin-flip symmetry, i. e., $f(s) = f(\bar{s})$ for all s where \bar{s} is the bitwise complement of s . This causes a problem called synchronization problem by van Hoyweghen, Goldberg, and Naudts (2002), i. e., 0-colored building blocks compete with 1-colored building blocks.

We only investigate randomized local search (RLS) and the well-known (1+1) EA. Comparisons of these two algorithms have been presented by Garnier, Kallel, and Schoenauer (1999) and Ladret (2004). These algorithms are discussed in Section 2. In Section 3, we consider graphs consisting of two cliques of equal size which are connected by some bridge edges. It is interesting to investigate how the number and the structure of the bridges help to find a monochromatic coloring. The Ising model on the ring has been analyzed intensively by Fischer and Wegener (2004) and there are some results for the two-dimensional torus by Fischer (2004). We consider the two-dimensional torus more intensively (Section 4) and investigate the torus with the smallest side length, namely the hypercube (Section 5).

Preliminary experiments have been performed to get some insight which has led to clearly formulated hypotheses. These hypotheses have been tested by independent experiments. We use different statistical tests depending on the type of hypothesis under consideration. The statistical tests are carried out using the software SPSS (Version 11.5, see www.spss.com). Regression analysis is done using gnuplot (Version 3.7, see www.gnuplot.info), using the standard settings.

The statistical tests used are the Wilcoxon signed rank test (WSRT), the Mann-Whitney test (MWT), and the binomial test. The former two tests are nonparametric tests to confirm the hypothesis that one random variable “systematically” produces larger values than another one. WSRT pairs the samples, takes their differences, ranks the absolute values and considers the sum of the ranks belonging to positive differences. MWT ranks all values and compares the rank sums for both samples. The binomial test is a test to confirm the hypothesis that the parameter p of a binomially distributed random variable is larger than some given constant p^* or that the difference of the parameters p and q of two binomially distributed random variables is larger than some given constant c .

In many cases, the hypotheses are confirmed on a good significance level. More precisely, we state our results in the form “Hypothesis H has been confirmed on the significance level α ” which means that the probability of producing our data if H is not true is bounded by α using the result of the SPSS package rounded to 3 digits. Hence, smaller values of α correspond to better results. The dimension n of the search space can be arbitrarily large and, therefore, no experiment can show evidence for all n . In all cases, our results are restricted to the

problem dimensions considered in the experiments. Results are called significant, if $\alpha \leq 0.05$, very significant if $\alpha \leq 0.01$, and highly significant, if $\alpha \leq 0.001$.

2 Randomized Local Search and the (1+1) EA

RLS and the (1+1) EA work on “populations” of size 1, i. e., with a single individual or search point x . The first search point x is chosen uniformly at random. Then the offspring x' is obtained from x by mutation. Selection chooses x' iff $f(x') \geq f(x)$. The mutation operator of RLS flips a single bit which is chosen uniformly at random. The mutation operator of the (1+1) EA flips each bit independently from the others with probability $1/n$.

RLS is a hill climber with a small neighborhood and can get stuck in local optima. The (1+1) EA always finds the optimum in expected finite time since each individual y has a positive probability to be produced as offspring of x . It is a matter of taste whether the (1+1) EA is considered as hill climber. No worsenings are accepted but big jumps are possible.

We restrict our investigations to these simple algorithms. One reason is that they reflect the so-called “game of life.” Vertices of the graph correspond to “individuals” of some “society” and edges model relations and influences between them. The problem is to estimate the time until a stable situation is obtained and whether in this stable situation all individuals are synchronized, i. e., they have the same color. All our results imply results on independent multiple runs which seem to be more useful than larger populations without crossover where one has to ensure diversity to “approximate independency.” Crossover may be helpful as shown for the ring (van Hoyweghen, Naudts, and Goldberg (2002), Fischer and Wegener (2004)). Two-point crossover is suitable on rings since it may exchange subblocks. For general graphs, only graph-specific crossover operators reflect the graph structure, e. g., exchanging rectangles for the two-dimensional torus. In this paper, we do not analyze such algorithms.

3 Partially Connected Cliques

Partially connected cliques consist of two disjoint equal-size cliques which are connected by m edges called bridges. We are interested in the success probability p of algorithm $A \in \{\text{RLS}, (1+1) \text{ EA}\}$ to find an optimal stable, i. e., a monochromatic search point before producing a stable unsynchronized search point where both cliques are monochromatic but of different colors. Theory is easy in the two extreme cases (apply the methods of Droste, Jansen, and Wegener (2002)). If $m = 0$, $p = 1/2$ and a stable situation is reached on average in $\Theta(n \log n)$ steps. If the stable search point is not an optimal one, it takes exponential time until the (1+1) EA finds an optimum. If all bridges exist, i. e., $m = n^2/4$, p is exponentially close to 1 and the expected time to produce an optimal search point is $\Theta(n \log n)$.

It seems to be obvious that many bridges are better than few bridges. Which “structure” of the bridges is useful? We distinguish three structures, namely the

random choice of m bridges (r) where the set of bridges is chosen uniformly at random, the concentrated choice (c) where all bridges touch only $\lceil m^{1/2} \rceil$ vertices of each clique, and a uniform construction (u): let $v_0, \dots, v_{n/2-1}$ be the vertices of the first clique and $w_0, \dots, w_{n/2-1}$ the vertices of the second clique. Then structure (u) is constructed by connecting v_i with $w_i, \dots, w_{(i+d-1) \bmod n/2}$ where $d \in \{\lfloor 2m/n \rfloor, \lceil 2m/n \rceil\}$ for all $0 \leq i \leq n/2-1$. Our results are based on 2000 runs for each scenario (n, s, q, A) where $n = 20i, 3 \leq i \leq 10$, or $n \in \{300, 400, 600\}$, $s \in \{r, u, c\}$ describes the bridge structure, $q = i/500, 0 \leq i \leq 500$, is the fraction of chosen bridges among all possible bridges, and $A \in \{\text{RLS}, (1+1) \text{ EA}\}$.

Result 1 (WSRT): *For all considered $n > 60$, RLS is more successful on structure u than on structure r (highly significant). This holds for the (1+1) EA only for $n \leq 100$ (significant) while the opposite is true for $n = 300$ (significant) and $n \in \{400, 600\}$ (highly significant).*

Tables and plots of the results lead to the impression that there are “obviously” no differences between the two structures. Large numbers of runs and well-chosen statistical tests can confirm small differences as highly significant. Binomial tests (for each of the 501 q -values) have confirmed (significant) that the success probabilities differ by less than 0.05. More precisely, 480 tests have confirmed this (but 4.2% outliers are expected for $\alpha = 0.05$) if $(n, A) \neq (600, (1+1) \text{ EA})$.

Result 2 (binomial tests): *For all cases where $(n, A) \neq (600, (1+1) \text{ EA})$ the difference between the success probability of the bridge structures r and u is less than 0.05 (significant).*

From the viewpoint of a single vertex, structure r increases the variance for the number of adjacent bridges. RLS suffers from vertices with a small number of bridges while the (1+1) EA can profit from a few vertices with many bridges in steps flipping many bits.

Result 3 (WSRT): *For all considered n , the success probability of the (1+1) EA on structure c is larger than on the structures r and u (highly significant).*

This can be explained as follows: The vertices with many bridges try to obtain the same color and then force the cliques to synchronize. For RLS, there is less “communication” between many vertices. For small m , the vertices with bridges have almost no influence. Moreover, if there are enough vertices without bridges, they may decide the colors of the cliques. Indeed, for each (n, q) we could decide (significant) whether RLS is better on c than on r (or u which gives the same results). Then, for fixed n and increasing q , one can observe four phases. First, RLS is better on r , then better on c , again better on r , and finally there are enough bridges to obtain a success probability close to 1 for both structures, see Figure 1.

Result 4 (WSRT): *For each considered n , there are values $\alpha < \beta < \gamma$ with the following properties. If $q \in (\alpha, \beta)$, RLS is more successful on c than on r , and,*




Fig. 1. The observed success rates for RLS and the two specified bridge structures and $n = 600$.

if $q < \alpha$ or $q \in (\beta, \gamma)$, RLS is less successful on c than on r (all statements highly significant). If $q > \gamma$, the success probability of RLS is close to 1 for both structures.

Result 5 (WSRT): For all considered (n, s) , the $(1+1)$ EA has a larger success probability than RLS (highly significant).

The $(1+1)$ EA may flip a 0-vertex and a 1-vertex of the same clique. Then the bridge edges alone decide about acceptance of the new search point. This strengthens the influence of the bridges. The reader may not be surprised by the results since he or she believes that the $(1+1)$ EA outperforms RLS almost always. However, Fischer and Wegener (2004) have proved that RLS is faster on the ring than the $(1+1)$ EA.

Based on results of Wegener and Witt (2004) on bimodal quadratic functions one may conjecture that the very early steps decide which local optimum will be reached. This implies that, after these steps, we are able to predict whether the run is successful. We have investigated prediction rules which are applied after only $n^{3/4}$ steps, more precisely the value $n^{3/4}$ is replaced by the closest integer. Rule $R1$ predicts a success iff the majority color is given to at least $n/2 + n^{3/4}/8$ vertices while $R2$ predicts a success iff more than half of the bridges are monochromatic. The experiments are based on 1000 runs of the $(1+1)$ EA and each pair (n, q) , $n = 128i$, $1 \leq i \leq 8$, $q = 2^{-j}$, $0 \leq j \leq 9$. The results for neighbored values of n are combined.

Result 6 (WSRT): For the considered n , the success probability of the prediction rules is increasing with n for each q (highly significant for $n \leq 768$, very significant for $n > 768$).

One may conjecture that the success probabilities of the prediction rules always converge to a value close to 1.

The last issue is to investigate the expected time until a stable situation is obtained. This time equals $\Theta(n \log n)$ in the extreme cases of no bridges or all bridges. Based on all the experiments one might conjecture that this bound holds for all cases. Such a general statement cannot be confirmed by experiments. We present a bridge structure (see Figure 2) which slows down the synchronization process. Let $v_1, \dots, v_{n/4}$ be vertices of the first clique and connect v_i to $n/2 - 2(i + 1)$ bridges (which lead to arbitrary vertices of the other clique).




Fig. 2. An example for the specified bridge structure and $n = 20$. All clique edges are omitted to improve readability. The vertices of the first clique, shown on the left, are enumerated clockwise starting with the rightmost vertex. The coloring describes the special search point s_3 .

Note that the bridge density q is approximately $1/4$. The special search point s_i corresponds to the coloring where exactly the vertices v_1, \dots, v_i and the vertices of the other clique are 0-colored, see Figure 2.

Proposition 1. *Starting with $s_i, 1 \leq i \leq n/8$, RLS and the (1+1) EA need on average $\Theta(n^2)$ steps to reach a stable situation.*

Sketch of Proof. For RLS, the only accepted step is to flip the color of v_{i+1} . This takes on average $\Theta(n)$ steps and we reach s_{i+1} . For the (1+1) EA, a step can be accepted only if $\Omega(n^{1/2})$ bits flip (this probability is exponentially small) or if v_{i+1} and only further vertices in $\{v_1, \dots, v_{n/4}\}$ flip. Either we flip $v_{i+1}, v_{i+2}, \dots, v_{n/4}$ one after another or we flip v_{i+1} and one of the vertices v_{i+2} and v_{i+3} in one step. The expected waiting time in both cases equals $\Theta(n^2)$. \square

This proposition proves that there are situations which imply an expected waiting time of $\Theta(n^2)$. This does not disprove the conjecture since it may be very unlikely to reach such a special situation, more precisely, a search point $s_i, 1 \leq i \leq n/8$, or its bitwise complement. We have performed for each $n = 128j, 1 \leq j \leq 16$, 10000 runs of the (1+1) EA. Let p_n be the fraction of runs reaching a special situation. In our experiment, $0.1115 \leq p_n \leq 0.1278$ and there is no tendency of increasing or decreasing p_n -values. Let q_n be the true probability that a run reaches a special situation.

Result 7 (binomial test): *For the considered n , $0.10 \leq q_n \leq 0.14$ (highly significant).*

Hence, we should expect larger synchronization times on special partially connected cliques.

4 The Ring and the Two-Dimensional Torus

Fischer and Wegener (2004) have shown that it takes an expected number of $O(n^3)$ steps until RLS or the (1+1) EA finds the optimum of the Ising model on the ring. There are situations where this number is $\Theta(n^3)$ and experiments confirm that it is almost sure to run in such a situation. For each non-optimal search point, there are accepted 1-bit flips but they do not necessarily improve the fitness.

The two-dimensional torus T_n is defined on $n = k^2$ vertices $(i, j), 0 \leq i, j \leq k - 1$. Each vertex (i, j) has the four neighbors $(i - 1, j), (i + 1, j), (i, j - 1)$, and $(i, j + 1)$ where -1 is identified with $k - 1$ and k is identified with 0 . In T_n^* , the torus with diagonals, (i, j) has four more neighbors: $(i - 1, j - 1), (i - 1, j + 1), (i + 1, j - 1)$, and $(i + 1, j + 1)$. A vertical ring of width w consists of all vertices $(i, j), i' \leq i < (i' + w), 0 \leq j \leq k - 1$. In a similar way, we can define horizontal rings. For T_n^* , diagonal rings of width w consist of all vertices (i, j) with $0 \leq i, j \leq k - 1$ such that $(i + j) \bmod k \in \{c \bmod k, \dots, (c + w - 1) \bmod k\}$ (or $(i - j) \bmod k \in \{c \bmod k, \dots, (c + w - 1) \bmod k\}$, resp.) holds for some fixed c (see Figure 3). Differently colored rings describe stable situations where special groups of “many” vertices have to flip for an acceptable step.




Fig. 3. Examples for vertical and diagonal rings. Vertices at the ring borders are marked with dots.

Our results are based on experiments with $k = 5i, 2 \leq i \leq 10$. The dimension of the search space is $n = k^2$ which equals 2500 for $i = 10$. Runs are stopped in stable situations. The number of runs was 2000 ($2 \leq i \leq 4$), 1000 ($5 \leq i \leq 6$), 500 ($i = 7$), and 250 ($8 \leq i \leq 10$).

One conjecture is that the probability of producing a non-optimal stable situation does not depend much on n , see Figure 4.

Result 8 (*regression analysis with functions $an + b$*): The percentages p_n and p_n^* of runs producing a non-optimal stable situation are described best by a linear function with almost vanishing linear term, namely $a = -0.001 \dots$ and $a^* = 0.004 \dots$.

Nevertheless, the values of p_n and p_n^* are varying with n (without clear structure). It is easier to break rings of small width than to break rings of large width. Our conjecture is that it is likely to obtain rings whose width is linear with respect to k .

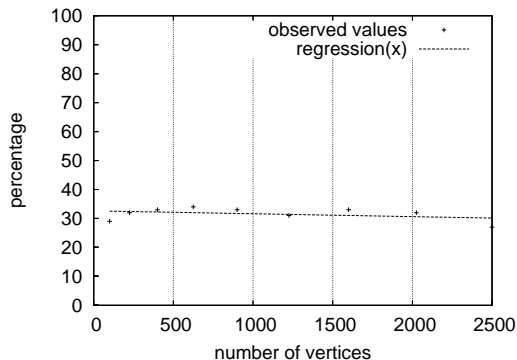


Fig. 4. The percentages of runs where the (1+1) EA reached a ring for T_n .

Result 9 (*regression analysis with functions $ak+b$*): If stable rings are produced, their width is best described by $0.407k - 0.108$ for T_n and $0.407k + 0.171$ for T_n^* .

This confirms the conjecture that stable rings typically lead to very stable suboptimal solutions. By symmetry, the probability of horizontal rings is the same as for vertical rings. Diagonal rings differ from these rings. For a horizontal ring of width $w \geq 4$, there are $2k$ ring vertices with a neighbor outside the ring. This number increases to $4k$ for diagonal rings, see Figure 3. There is more information crossing the borders of diagonal rings than in the horizontal case which should decrease the probability of constructing diagonal rings.

Result 10 (*WSRT*): For all considered n , the probability to create a diagonal ring is smaller than the probability to create a horizontal ring (*highly significant*).

Fischer (2004) has proved that the expected time to create a stable situation is bounded above by $O(n^3) = O(k^6)$. It is interesting that our data leads to estimates of smaller magnitudes. Compared with the ring or the one-dimensional torus a stable situation is found more efficiently (smaller side length), but there is a non-negligible probability of producing a stable non-optimal situation.

Result 11 (*regression analysis with polynomials of degree 5*): The time to find a stable situation for T_n is best described by

$$-0.001k^5 + 0.123k^4 + 28k^3 - 741k^2 + 73k + 73.$$

The essential term is the $O(k^3) = O(n^{3/2})$ term. A regression with polynomials of degree 3 leads to almost the same error. Similar results hold for the two subcases of obtaining an optimal search point and of obtaining a stable non-optimal situation. It seems that the second case needs a little longer.

5 Hypercubes

One may investigate d -dimensional tori for increasing d . We investigate the extreme case of tori with the minimal side length of 2. This is equivalent to the

hypercube $\{0, 1\}^d$ with edges between Hamming neighbors. The dimension of the search space equals $n = 2^d$. We have performed 1000 runs for $d \in \{6, \dots, 22\}$ and RLS and the (1+1) EA. The runs were stopped in situations without 1-bit flips improving the fitness. Such search points are called stable. The first conjecture is that the (1+1) EA is better than RLS.

Result 12 (MWT): *For the considered d , the probability of finding an optimum as first stable search point is larger for the (1+1) EA than for RLS (highly significant).*

Result 13 (MWT): *Only runs not finding the optimum are considered. For the considered d , the fitness of the first stable search point is for the (1+1) EA larger than for the RLS (highly significant for odd $d \geq 9$, significant for $d = 12$, very significant for $d \in \{14, 16\}$ and highly significant for $d \in \{18, 20, 22\}$).*

However, the (1+1) EA needs more time for the better results.

Result 14 (MWT): *For the considered d , RLS needs less time to find a stable search point than the (1+1) EA (highly significant).*

One may conjecture that the success probability, i. e., the probability of finding an optimal search point as first stable search point, is decreasing with d . We have to be careful since there is an essential difference between odd d and even d . Each vertex has d neighbors and a flip of vertex v is accepted iff at most $d/2$ neighbors of v share their color with v . For even d , there are plateaus of constant fitness and the algorithm may finally escape from this plateau to a higher fitness level. Hence, we first compare only even values of d and only odd values of d and, afterwards, we compare odd values of d with even values of d .

Result 15 (MWT): *For all considered d , the probability of finding an optimal point for dimension $d + 2$ is smaller than for dimension d (highly significant for most d . For the (1+1) EA, only very significant for $d \in \{11, 12, 15\}$ and only significant for $d \in \{18, 20\}$, and not significant for $d = 6$. For RLS, only very significant for $d \in \{12, 17\}$, only significant for $d = 18$ and not significant for $d \in \{15, 16, 19, 20\}$).*

For the large values of d , the success probabilities are decreasing slower than before and it is more difficult to obtain significant results.

Another conjecture is that the algorithms produce better results for even d than for odd d . In order to check whether it is less likely to find an optimum for odd d , we compare the data for d with the union of the data for $d - 1$ and $d + 1$, which is a way to consider the mean of the average values for $d - 1$ and $d + 1$. We refer to this mean as d' .

Result 16 (MWT): *The success probability for odd d is smaller than for d' and for even d larger than for d' (highly significant for RLS, highly significant for the (1+1) EA for $d \leq 16$, very significant for $d \in [17, 20]$ and significant for $d = 21$).*

Similar significance levels are obtained for the following results.

Result 17 (MWT): *For increasing values of even d , the approximation ratio (quotient of the fitness of the first stable solution and the fitness of an optimum) is increasing. The same holds for odd d . The approximation ratio for odd d is smaller than for d' and for even d larger than for d' .*

Conclusions

The Ising model offers a variety of different scenarios to investigate a “game of life” with different relations between the individuals of the considered society. Each individual prefers to be synchronized with its “neighbors” and the question is whether the first stable situation is one where all individuals are synchronized. Even the three scenarios of partially connected cliques, the two-dimensional torus, and the hypercube lead to surprising observations. We have described and discussed those results which have been confirmed at least as significant. The tables with all the experimental data and the plots of these data contain much more information. They are not shown due to space limitations. In some situations, our statistical tests have led to results which were not observed in the tables and plots.

References

1. Droste, S., Jansen, T., and Wegener, I. (2002). On the analysis of the (1+1) evolutionary algorithm. *Theoretical Computer Science* 276, 51–81.
2. Fischer, S. (2004). A polynomial upper bound for a mutation-based algorithm on the two-dimensional Ising model. *Proc. of the Genetic and Evolutionary Computation Conference (GECCO 2004)*, LNCS 3102, 1100–1112.
3. Fischer, S. and Wegener, I. (2004). The Ising model on the ring: Mutation versus recombination. *Proc. of the Genetic and Evolutionary Computation Conference (GECCO 2004)*, LNCS 3102, 1113–1124.
4. Garnier, J., Kallel, L., and Schoenauer, M. (1999). Rigorous hitting times for binary mutations, *Evolutionary Computation*, Vol. 7, 173–203.
5. Holland, J. H. (1975). *Adaptation in Natural and Artificial Systems*. Univ. of Michigan, MI.
6. Ising, E. (1925). Beiträge zur Theorie des Ferromagnetismus. *Z. Physik* 31, 235–288.
7. Ladret, V. (2004). Asymptotic hitting time for a simple evolutionary model of protein folding. Technical report, <http://arxiv.org/abs/math.PR/0308237>.
8. Naudts, B. and Naudts, J. (1998). The effect of spin-flip symmetry on the performance of the simple GA. *Proc. of the 5th Conf. on Parallel Problem Solving from Nature (PPSN-V)*, LNCS 1488, 67–76.
9. Pelikan, M. and Goldberg, D. E. (2003). Hierarchical BOA solves Ising spin glasses and MAXSAT. *Proc. of the Genetic and Evolutionary Computation Conference (GECCO 2003)*, LNCS 2724, 1271–1282.
10. Van Hoyweghen, C., Goldberg, D. E., and Naudts, B. (2002). From twomax to the Ising model: Easy and hard symmetrical problems. *Proc. of the Genetic and Evolutionary Computation Conference (GECCO 2002)*, 626–633. Morgan Kaufmann.
11. Van Hoyweghen, C., Naudts, B., and Goldberg, D. E. (2002). Spin-flip symmetry and synchronization. *Evolutionary Computation* 10, 317–344.
12. Wegener, I. and Witt, C. (2004). On the analysis of a simple evolutionary algorithm on quadratic pseudo-boolean functions. To appear in *Journal of Discrete Algorithms*.