1

Searching for Ground States of Ising Spin Glasses with Hierarchical BOA and Cluster Exact Approximation

Martin Pelikan[1] and Alexander K. Hartmann[2]

[1] Dept. of Math and Computer Science, 320 CCB
University of Missouri at St. Louis
One University Blvd.
St. Louis, MO 63121
pelikan@cs.umsl.edu

[2] Institut für Theoretische Physik
Universität Göttingen
Friedrich-Hund-Platz 1
37077 Göttingen, Germany
hartmann@physik.uni-goettingen.de

Summary. This chapter applies the hierarchical Bayesian optimization algorithm (hBOA) to the problem of finding ground states of Ising spin glasses with $\pm J$ and Gaussian couplings in two and three dimensions. The performance of hBOA is compared to that of the simple genetic algorithm (GA) and the univariate marginal distribution algorithm (UMDA). The performance of all tested algorithms is improved by incorporating a deterministic hill climber based on single-bit flips. The results show that hBOA significantly outperforms GA and UMDA on a broad spectrum of spin glass instances. Cluster exact approximation (CEA) is then described, which is capable of making large updates of spin-glass configurations by using a max-flow algorithm. The results show that CEA enables all tested algorithms to solve larger spin glass instances although hBOA still outperforms other compared algorithms.

Key words: Hierarchical BOA, genetic algorithm, cluster exact approximation, spin glass, ground states, scalability.

1.1 Introduction

Ising spin glasses are prototypical models for disordered systems and have played a central role in statistical physics during the last three decades [5, 20, 8, 40]. Examples for experimental realizations of spin glasses are metals with magnetic impurities, e.g. gold with a small fraction of iron added. The interaction between two magnetic spins, sitting each on an iron atom, is modulated by the conducting material leading to an effective interaction between these spins which varies with the distance in magnitude *and* in sign. Due to the random distribution of the spins, competing interactions arise, where it is not possible to satisfy all interactions at the same time. This effect is called *frustration*. At low temperatures, this leads to an ordered low-temperature phase exhibiting the so-called spin-glass order, which is characterized by frozen local moments, but which does not exhibit a global magnetization. This spin glass-phase poses a challenging, unsolved problem in theoretical physics because the nature of its phase-space structure is still not well understood [26], despite three decades of intensive research. Anyway, it is widely believed that spin glasses exhibit a rough energy landscape leading to this intrinsic complexity.

Spin glasses represent also a large class of challenging problems for optimization algorithms [16, 17, 18] where the task is to minimize energy of a given spin glass instance [22, 25, 15, 38, 29, 9]. States with the lowest energy are called *ground states* and thus the problem of minimizing energy of spin glass instances can be formulated as the problem of finding ground states of these instances. There are two main challenges that must be tackled to find ground states of spin glasses efficiently and reliably: (1) There are many local optima in the energy landscape (the number of local optima may grow faster than exponentially with problem size). (2) The local minima are often surrounded by high-energy configurations, which make it difficult for local operators to escape the local optimum once they get trapped in it.

The purpose of this chapter is to apply the hierarchical Bayesian optimization algorithm (hBOA) [28, 27] to a broad spectrum of instances of the problem of finding ground states of Ising spin glasses. The performance of hBOA is compared to that of the simple genetic algorithm (GA) and the univariate marginal distribution algorithm (UMDA). We also describe cluster exact approximation (CEA) [14], which provides an efficient method to perform large updates of spin glass configurations to decrease their energy. CEA is then incorporated into all tested algorithms to improve their performance and the resulting hybrids are tested on a number of spin glass problem instances.

The chapter is organized as follows. Section 1.2 describes the problem of finding ground states of Ising spin glasses and discusses existing methods for solving instances of this problem. Section 1.3 briefly describes algorithms hBOA, GA and UMDA; additionally, the section describes a deterministic hill climber, which is incorporated into all tested algorithms to improve their performance. Section 1.4 presents initial experiments with hBOA, GA and UMDA on a broad range of Ising spin glass instances. Section 1.5 describes cluster exact approximation (CEA). Section 1.6 incorporates CEA into hBOA and GA to improve their performance and allow efficient solution of large spin glass instances. Finally, Section 1.7 summarizes and concludes the chapter.

1.2 Ising Spin Glass

A very simple model to describe a finite-dimensional Ising spin glass is typically arranged on a regular 2D or 3D grid where each node i corresponds to a spin s_i and each edge $\langle i, j\rangle$ corresponds to a coupling between two spins. Each edge has a real value associated with it that defines the relationship between the two connected spins. To approximate the behavior of the large-scale system, periodic boundary conditions are often used that introduce a coupling between the first and the last element along each dimension.

For the classical Ising model, each spin s_i can be in one of two states: $s_i = +1$ or $s_i = -1$. Note that this simplification corresponds to highly anisotropic systems, which do indeed exist in some experimental situations. Nevertheless, the two-state Ising model comprises all basic effects also found for models with more degrees of freedom. A specific set of coupling constants define a spin glass instance. Each possible setting of the states of all spins is called a spin configuration.

Given a set of coupling constants $J_{i,j}$, and a configuration of spins $C = \{s_i\}$, the energy can be computed as

$$E(C) = \sum_{\langle i,j\rangle} s_i J_{i,j} s_j \ , \tag{1.1}$$

where $i, j \in \{0, 1, \ldots, n-1\}$ and $\langle i, j\rangle$ denote the nearest neighbors in the underlying grid (allowed edges).

Given a set of coupling constants, the usual task in statistical physics is to integrate a known function over all possible configurations of spins, where the configurations are distributed according to the Boltzmann distribution; that is, the probability of encountering a configuration C at temperature T is given by

$$p(C) = \frac{\exp\left(-E(C)/T\right)}{\sum_{\tilde{C}} \exp\left(-E(\tilde{C})/T\right)} \ , \tag{1.2}$$

where the sum in the denominator runs over all possible spin configurations.

From the physics point of view, it is interesting to know the ground states (configurations associated with the minimum possible energy). Finding extremal energies then corresponds to sampling the Boltzmann distribution with temperature approaching 0 and thus the problem of finding ground states is simpler *a priori* than integration over a wide range of temperatures. However, most of the conventional methods based on sampling the above Boltzmann distribution (see Equation 1.2) fail to find the ground states because they get often trapped in a local minimum.

The problem of finding ground states for a given set of coupling constants is a typical optimization problem, where the task is to find an optimal configuration of spins that minimizes energy (see Figure 1.1). Although polynomial-time deterministic methods exist for both types of 2D spin glasses [4, 2, 10, 11],

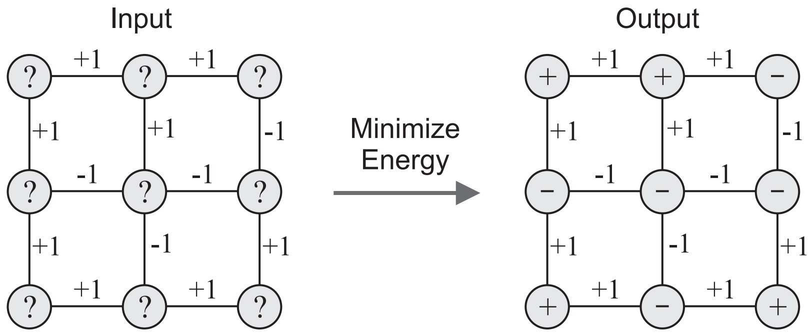

Fig. 1.1. Finding ground states of Ising spin glasses on a 2D system with no periodic boundary conditions. The left-hand side shows the input, which consists of a 2D grid with coupling constants between adjacent spins. The right-hand side shows the output, which consists of one of the ground states (the spin configurations with the minimum energy). The energy of the ground state shown in this figure is -10.

most algorithms based on local search operators, including a (1+1) evolution strategy, conventional Monte Carlo simulations, and Monte Carlo simulations with Wang-Landau [39] or multicanonical sampling [3], scale exponentially [7] and are thus impractical for solving this class of problems. The origin for this slowdown is due to the suppressed relaxation times in the Monte Carlo simulations in the vicinity of the extremal energies because of the enormous number of local optima in the energy landscape. Recombination-based genetic algorithms succeed if recombination is performed in a way that interacting spins are located close to each other in the representation; k-point crossover with a rather small k can then be used so that the linkage between contiguous blocks of bits is preserved (unlike with uniform crossover, for instance). However, the behavior of such specialized representations and variation operators cannot be generalized to similar slowly equilibrating problems that exhibit different energy landscapes, such as protein folding or polymer dynamics.

In order to obtain a quantitative understanding of the disorder in a spin glass system introduced by the random spin-spin couplings, one generally analyzes a large set of random spin glass instances for a given distribution of the spin-spin couplings. For each spin glass instance, the optimization algorithm is applied and results are analyzed to obtain a measure of computational complexity. Here we first consider two types of initial spin-spin coupling distributions, the $\pm J$ spin glass and the Gaussian spin glass.

1.2.1 $\pm J$ spin glass

In the $\pm J$ spin glass, each spin-spin coupling constant is set randomly to either $+1$ or -1 with equal probability. Energy minimization in this case can be transformed into a constraint satisfaction problem, where the constraints relate spins connected by a coupling constant. For any two connected neighbors i and j, if $J_{i,j} = +1$, then the constraint requires spins i and j to be different, whereas if $J_{i,j} = -1$, then the constraint requires spins i and j to be the same. Energy is minimized when the number of satisfied constraints is maximized.

1.2.2 Gaussian spin glass

In the Gaussian spin glass, coupling constants are generated according to a zero-mean Gaussian distribution with variance one. For real-valued couplings, energy minimization can be casted as a constraint satisfaction problem with weighted constraints. Gaussian spin glass instances are usually easier for methods based on local operators (such as MCMC), whereas for recombination-based algorithms (such as hBOA), Gaussian spin glass instances tend to be slightly more difficult than $\pm J$ instances.

1.3 Compared Algorithms

This section briefly describes the algorithms compared in this chapter: (1) The hierarchical Bayesian optimization algorithm (hBOA), (2) the genetic algorithm (GA), and (3) the univariate marginal distribution algorithm (UMDA). Additionally, the section describes the deterministic hill climber that is used in all compared algorithms to locally improve candidate solutions.

1.3.1 Hierarchical Bayesian optimization algorithm (hBOA)

The hierarchical Bayesian optimization algorithm (see Chapter **??** of this book and references [28, 27]) evolves a population of candidate solutions. Each candidate solution is represented by an n-bit binary string where each bit specifies the value of one of the n spins (0 represents state -1 and 1 represents state $+1$). The population is initially generated at random according to a uniform distribution over all n-bit strings.

Each iteration starts by selecting a population of promising solutions using any common selection method of genetic and evolutionary algorithms, such as tournament and truncation selection. In this chapter, binary tournament selection is used. Binary tournament selection selects one solution at a time by first choosing two random candidate solutions from the current population and then selecting the best solution out of this subset. Random tournaments are repeated until the selected population has the same size as the original population and thus each candidate solution is expected to participate in two tournaments.

New solutions are generated by building a Bayesian network with decision trees for the selected solutions and sampling the built Bayesian network.

To ensure useful diversity maintenance, the new candidate solutions are incorporated into the original population using restricted tournament replacement (RTR) [13]. RTR incorporates the new candidate solutions into the original population one solution at a time. For each new candidate solution X, a random subset of candidate solutions in the original population is first selected. The solution that is closest to X is then selected from this subset. The new candidate solution replaces the selected closest solution in the selected window if it is better. Here we measure the distance of two solutions by counting the number of bits where solutions differ.

The run is terminated when termination criteria are met; for example, the run can be terminated when the global optimum is found or when the number of iterations reaches some threshold.

1.3.2 Genetic algorithm (GA)

The genetic algorithm (GA) [19, 12] also evolves a population of candidate solutions represented by fixed-length binary strings. The first population is generated at random.

Each iteration starts by selecting promising solutions from the current population. New solutions are created by applying variation operators to the population of selected solutions. Specifically, crossover is used to exchange bits and pieces between pairs of candidate solutions and mutation is used to perturb the resulting solutions. Here we use one-point crossover that exchanges the tails of two binary strings starting in a randomly chosen position with probability p_c. As a mutation operator, bit-flip mutation is used, which flips each bit with a specified probability p_m (p_m is usually small). The new candidate solutions are incorporated into the original population using restricted tournament replacement (RTR) [13]. The run is terminated when termination criteria are met.

1.3.3 Univariate marginal distribution algorithm (UMDA)

The univariate marginal distribution algorithm (UMDA) [23] also evolves a population of candidate solutions represented by binary strings, starting with a random population.

Each iteration starts by selection. Then, the probability vector is learned that stores the proportion of 1s in each position of the selected population. Each bit of a new candidate solution is then set to 1 with the probability equal to the proportion of 1s in this position; otherwise, the bit is set to 0. Consequently, the variation operator of UMDA preserves the proportions of 1s in each position while decorrelating different string positions. The new candidate solutions are incorporated into the original population using restricted tournament replacement (RTR) [13]. The run is terminated when termination criteria are met.

1.3.4 Deterministic hill climber

The deterministic hill climber (DHC) is incorporated to hBOA, GA and UMDA to improve their performance. DHC takes a candidate solution represented by an n-bit binary string on input. Then, it performs one-bit changes on the solution that lead to the maximum improvement of solution quality (maximum decrease in energy). DHC is terminated when no single-bit flip improves solution quality and the solution is thus locally optimal. Here, DHC is used to improve every solution in the population before the evaluation is performed. The hybrids created by incorporating DHC into hBOA, GA and UMDA are referred to as hBOA+DHC, GA+DHC, and UMDA+DHC, respectively.

DHC improves the performance of hBOA, GA and UMDA by a constant factor [27]. In Section 1.5 we describe cluster exact approximation (CEA), which is capable of making large updates of spin glass configurations, providing a much more effective search for low-energy configurations.

1.4 Initial Experiments

This section presents the results of initial experiments on various instances of the problem of finding ground states of Ising spin glasses. Three algorithms are compared: hBOA+DHC, GA+DHC and UMDA+DHC. The section starts by describing problem instances used in the tests. Next, experiments are described. Finally, the experimental results are presented and discussed.

1.4.1 Tested Spin Glass Instances

Three types of spin glass instances have been considered in the initial experiments:

1. Two-dimensional Ising spin glass with $\pm J$ couplings.
2. Two-dimensional Ising spin glass with Gaussian couplings.
3. Three-dimensional Ising spin glass with $\pm J$ couplings.

In all instances, periodic boundary conditions are used.

To analyze scalability, for two-dimensional spin glasses, instances of size 6×6 ($n = 36$ spins) to 20×20 ($n = 400$ spins) have been considered; 1000 random instances have been generated for each problem size. For three-dimensional spin glasses, instances of size $4 \times 4 \times 4$ ($n = 64$ spins) to $7 \times 7 \times 7$ ($n = 343$ spins) have been considered; in this case, only 8 random instances have been generated for each problem size because of the increased computational resources required to solve the three-dimensional instances.

1.4.2 Description of Experiments

All compared algorithms use binary tournament selection to select promising solutions. As a replacement strategy, RTR is used where the window size w is set to the number of bits in solution strings but it is always ensured to be at most 5% of the population size, $w = \min(n, N/20)$. In GA+DHC, one-point crossover is used and the probability of crossover is $p_c = 0.6$. GA+DHC uses also bit-flip mutation, where the probability of flipping a bit is set to $1/n$.

For each problem instance, bisection is run to determine the minimum population size to ensure convergence in 5 independent runs (out of 5 runs total). Each run is terminated either when the algorithm has found the optimum or when the algorithm has failed to find the optimum for a large number of iterations. The optimum for most 2D instances was verified with the branch-and-cut algorithm provided at the Spin Glass Ground State Server at the University of Köln [33]. The remaining 2D instances with ground states were obtained from S. Sabhapandit and S. N. Coppersmith from the University of Wisconsin who identified the ground states using flat-histogram Markov chain Monte Carlo simulations [7]. All 3D instances with their ground states were obtained from Alexander Hartmann [15].

The upper bound on the number of iterations (generations) is determined empirically so that the number of iterations is sufficiently large for all tests. In general, the bound on the number of iterations for GA+DHC is larger than that for hBOA+DHC and UMDA+DHC because of the slower mixing with one-point crossover [32].

The performance of the compared algorithms is measured by the number of evaluated spin glass configurations until the optimum has been found. Additionally, we show the number of DHC updates until the optimum has been found.

1.4.3 Results

Figure 1.2 compares the performance of hBOA+DHC, UMDA+DHC and GA+DHC on two-dimensional spin glasses with $\pm J$ couplings and periodic boundary conditions. The results indicate that the number of evaluations and the number of DHC flips for hBOA+DHC grows with a low-order polynomial of problem size, i.e., with $O(n^{1.63})$ and with $O(n^{2.46})$, respectively. Furthermore, the results show that hBOA significantly outperforms GA+DHC and UMDA+DHC in both the number of evaluations and the number of DHC flips. The worst performance is achieved by UMDA+DHC.

Figure 1.3 compares the performance of hBOA+DHC and GA+DHC on two-dimensional spin glasses with Gaussian couplings and periodic boundary conditions. The results indicate similar behavior as for $\pm J$ couplings in two dimensions; hBOA+DHC outperforms both GA+DHC and UMDA+DHC; UMDA+DHC performs worst. The results also show that Gaussian couplings are harder for all tested algorithms. The reason for this is that in Gaussian spin

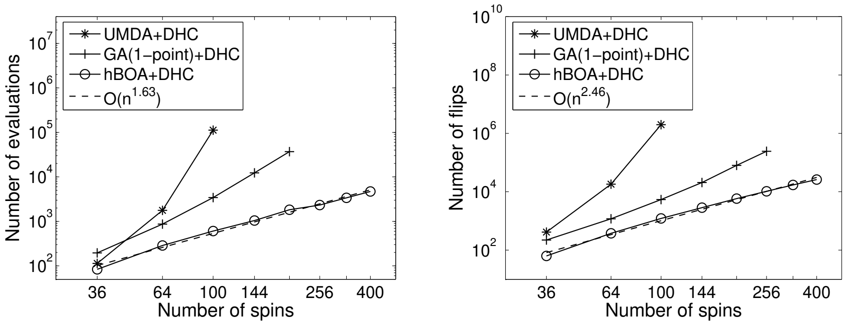

(a) Number of evaluations. (b) Number of flips.

Fig. 1.2. Performance of hBOA+DHC, UMDA+DHC, and GA+DHC on random 2D $\pm J$ Ising spin glasses.

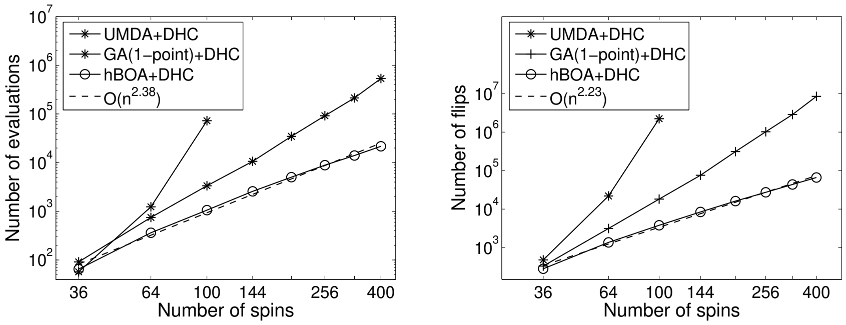

(a) Number of evaluations. (b) Number of flips.

Fig. 1.3. Performance of hBOA+DHC and GA+DHC on random 2D Gaussian Ising spin glasses.

glasses the constraints are scaled nonuniformly and it is known that most selectorecombinative evolutionary algorithms perform better on uniformly scaled problems than on non-uniformly scaled ones [36, 24].

For two-dimensional Ising spin glasses, a polynomial algorithm [10, 11] with complexity $O(n^{3.5})$ exists that computes the number of states at each energy level, including the ground state. It was shown [30] that on two-dimensional $\pm J$ Ising spin glasses, hBOA+DHC achieves asymptotic performance of the polynomial algorithm without any prior knowledge about spin

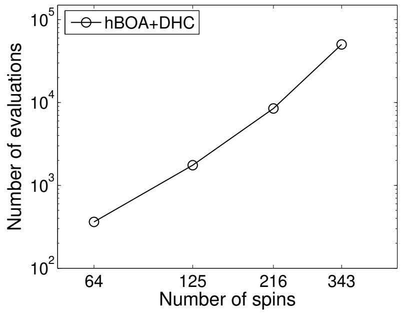

(a) Number of evaluations.

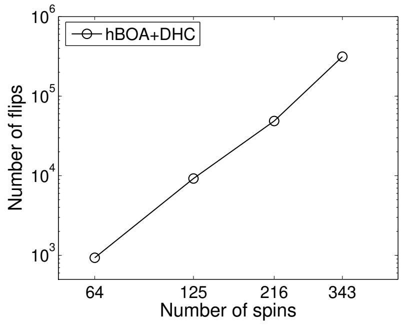

(b) Number of flips.

Fig. 1.4. Performance of hBOA+DHC on random 3D $\pm J$ Ising spin glasses.

glasses. The performance becomes slightly worse for the Gaussian spin glasses, although hBOA+DHC still retains low-order polynomial complexity.

Figure 1.4 shows the performance of hBOA+DHC on three-dimensional $\pm J$ spin glasses. Since both GA+DHC and UMDA+DHC have not been capable of solving most three-dimensional instances even with enormous computational resources, we only include the results for hBOA+DHC. The results show that the performance of hBOA+DHC appears to grow exponentially. This behavior is expected because the problem of finding ground states of three-dimensional spin glasses is NP-complete [1] and it is thus unlikely that there exists an algorithm that can solve this class of problems in polynomial time. However, we see that hBOA+DHC is still capable of solving instances of several hundreds spins, which are intractable with most standard optimization algorithms, such as genetic algorithms and simulated annealing.

1.5 Cluster Exact Approximation (CEA)

Due to the complex structure of the energy landscape of spin glasses, many local minima exist, which have energies very close to the ground-state energy. Usually these minima differ from the true ground states by flips of large domains. Hence, as already mentioned, the minima are surrounded by high energy barriers from the viewpoint of single-spin-flip dynamics. This leads to a very small efficiency of algorithms which apply single-bit (spin) changes as DHC. For this reason, we also consider an HC, which changes many spins at the same time, the so called *cluster-exact approximation* (CEA) [14].

This method constructs iteratively and randomly a non-frustrated cluster of spins. During the construction of the cluster a local gauge-transformation of the spin variables $s'_i = h_i s_i$ ($h_i = \pm 1$) is applied so that all interactions

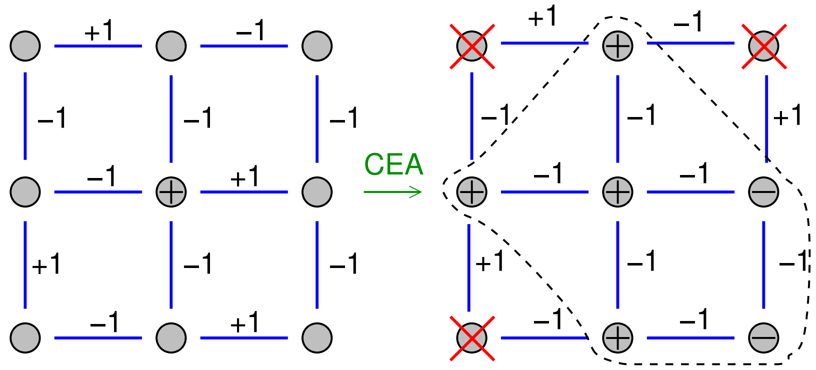

Fig. 1.5. Example of the cluster-exact approximation method. A part of a spin glass is shown, ferromagnetic bonds are denoted by -1, while antiferromagnetic bonds are denoted by $+1$, see Eq. (1.1). The construction starts with the spin i_0 at the center, the gauge field is set to $h_{i_0} = +1$. The right part displays the final situation. The spins which belong to the cluster carry a plus or minus sign which indicates how each spin is gauge-transformed, i.e. the gauge-field h_i, such that only ferromagnetic interactions remain inside the cluster. All other spins cannot be added to the cluster because it is not possible to multiply them by ± 1 to turn all adjacent bonds ferromagnetic. Please note that many other combinations of spins can be found to build a cluster without frustration.

between cluster spins become ferromagnetic ($J'_{i,j} = h_i J_{i,j} h_j < 0$), which is always possible for non-frustrated systems by definition. For each spin, which has not been considered so far, it is tested whether a value $h_i = \pm 1$ can be found, which makes all adjacent bonds ferromagnetic. If this is possible, the spin is added to the cluster, if not, it becomes a non-cluster spin. Fig. 1.5 shows an example of how the construction of the cluster works for a small spin-glass system. The order in which the spins are considered can be random or governed by some heuristics, here pure random order was used. Note that the clusters constructed in this way are typically very large, e.g. For 3d $\pm J$ spin glasses, each cluster contains typically 55 percent of all spins.

The non-cluster spins remain fixed during the following calculation; they act like local magnetic fields on the cluster spins. Consequently, the ground state of the gauge-transformed cluster is not trivial, although all interactions inside the cluster are now ferromagnetic. Since the cluster exhibits no bond-frustration, an energetic minimum state for its spins can be calculated in polynomial time by an algorithm which has been applied successfully for ferromagnetic random-field Ising models [16]. This algorithm is based on graph-theoretical methods [6, 34]: an equivalent network is constructed [31], the maximum flow is calculated [37, 35] and the spins of the cluster are set to orientations leading to a minimum in energy. Globally seen, either the energy is decreased or it remains the same.

1.5.1 Combining evolutionary algorithms and CEA

It is fairly straightforward to incorporate CEA into GA, hBOA or any other evolutionary algorithm. In this work, we repeat the CEA update step until the update fails to decrease the energy for a predefined number of iterations; more specifically, the bound on the number of failures is $\sqrt{n}$ for 2D spin glasses and it is $\sqrt[3]{n}$ for 3D spin glasses, where n is the number of spins. The hybrids created by incorporating CEA into hBOA and GA are referred to as hBOA+CEA and GA+CEA, respectively.

There are two important differences between DHC and CEA. DHC makes only single-bit updates and that is why it is able to reach the local optimum (with respect to one-bit flips) quickly but it is often trapped in a not-so-good local minimum. CEA makes much bigger updates, leading to a much more significant decrease in the energy, but also requiring more computational resources for each update.

The next section shows that CEA significantly improves the performance of both hBOA and GA, allowing these algorithms to solve much bigger problem instances efficiently.

1.6 Experiments

This section tests hBOA+CEA and GA+CEA on a broad range of two-dimensional and three-dimensional $\pm J$ Ising spin glass instances. Since the initial experiments indicated that UMDA can be expected to perform significantly worse than both hBOA and GA with DHC, we do not include UMDA in the comparison.

1.6.1 Tested Spin Glass Instances

Two types of spin glass instances have been considered in this set of experiments:

1. Two-dimensional Ising spin glass with $\pm J$ couplings.
2. Three-dimensional Ising spin glass with $\pm J$ couplings.

In all instances, periodic boundary conditions are used.

In the two-dimensional case, instances of size 15×15 ($n = 225$ spins) to 50×50 ($n = 2500$ spins) have been considered; 1000 random instances have been generated for each problem size. In the three-dimensional case, instances of size $6 \times 6 \times 6$ ($n = 216$ spins), $8 \times 8 \times 8$ ($n = 512$ spins), and $10 \times 10 \times 10$ ($n = 1000$ spins) have been considered; 1000 random instances have been generated for each problem size.

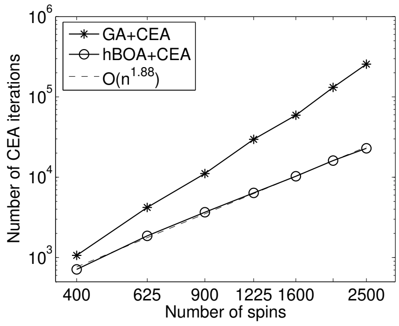

(a) 2D $\pm J$ Ising spin glass.

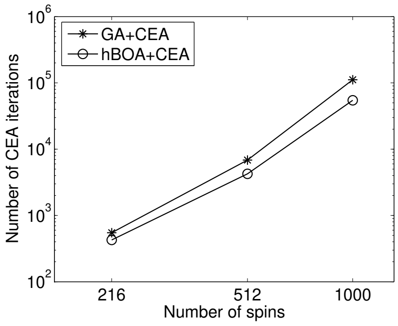

(b) 3D $\pm J$ Ising spin glass.

Fig. 1.6. Performance of hBOA+CEA and GA+CEA on $\pm J$ Ising spin glasses.

1.6.2 Description of Experiments

All parameters have been set like in the initial set of experiments in section 1.4 of this chapter. The only difference from the set-up of the initial experiments is that here CEA is used to improve all candidate solutions instead of DHC.

Since one update step of CEA is usually more computationally expensive than the entire evaluation [21], the time complexity of the compared algorithms is measured by the number of iterations of CEA as opposed to the number of evaluations.

1.6.3 Results

Figure 1.6a shows the performance of hBOA+CEA and GA+CEA on two-dimensional Ising spin glasses with $\pm J$ couplings. The results indicate that hBOA+CEA significantly outperforms GA+CEA and thus hBOA retains superior performance even with CEA. The results also show that incorporating CEA leads to a somewhat faster asymptotic growth of time complexity with problem size; on the other hand, the use of CEA provides a significant decrease of running time for the tested range of problems and, consequently, much larger problem sizes can be treated currently as compared to hBOA+DHC. Nonetheless, based on these results, it can be hypothesized that hBOA+DHC will become faster than hBOA+CEA for much larger two-dimensional spin glasses.

Figure 1.6b shows the performance of hBOA+CEA and GA+CEA on three-dimensional Ising spin glasses with $\pm J$ couplings. The results indicate that the performance grows faster than polynomially even with the use of CEA as is expected from the NP-completeness of this problem. However, CEA improves the performance of GA significantly and makes the difficult three-dimensional instances tractable even with GA. Nonetheless, hBOA+CEA still retains superior performance.

1.7 Summary and Conclusions

This paper tested the hierarchical Bayesian optimization algorithm (hBOA) on a large number of instances of the problem of finding ground states of Ising spin glasses with random couplings in two and three dimensions. The performance of hBOA was compared to that of the simple genetic algorithm (GA) and the univariate marginal distribution algorithm (UMDA). All algorithms were hybridized by using either a simple deterministic hill climber (DHC) or the cluster exact approximation (CEA). The results showed that hBOA significantly outperforms all other compared methods in all cases and that in some classes of Ising spin glasses, it achieves the performance of best polynomial-time analytical methods.

The problem of finding ground states of Ising spin glasses can be formulated as a constraint satisfaction problem and it has many features that can be found in other difficult constraint satisfaction problems, such as MAXSAT or graph coloring. Furthermore, this problem shares many challenges with other important classes of problems, such as protein folding. The most important source of problem difficulty is an enormous number of local optima that grows exponentially with problem size or faster. Additionally, to ensure effective exploration of the space of low-energy configurations, large-order interactions between spins must be considered. The results presented in this paper thus confirm that using hierarchical decomposition for solving difficult optimization problems with little problem-specific knowledge holds a big promise and that this line of computational-optimization research is likely to advance a number of areas that crucially depend on solving difficult optimization problems scalably.

Acknowledgments

The authors would like to thank Fabien Alet, Jiri Ocenasek, Kumara Sastry, Simon Trebst and Matthias Troyer for comments and help with the design, presentation and analysis of experiments.

This work was partially supported by the Research Award, the Research Board at the University of Missouri and by the *VolkswagenStiftung* (Germany) within the program “Nachwuchsgruppen an Universitäten”. Most experiments were done using the hBOA software developed by Martin Pelikan and David E. Goldberg at the University of Illinois at Urbana-Champaign. The experiments presented in this paper were completed at the Beowulf cluster at the University of Missouri at St. Louis and the Asgard cluster at the Swiss Federal Institute of Technology (ETH) Zürich.

References

[1] Barahona, F. (1982). On the computational complexity of ising spin glass models. *Journal of Physics A: Mathematical, Nuclear and General*, 15(10):3241–3253.

[2] Barahona, F., Maynard, R., Rammal, R., and Uhry, J. (1982). Morphology of ground states of a two dimensional frustration model. *J. Phys. A*, 15:673.

[3] Berg, B. A. and Neuhaus, T. (1992). Multicanonical ensemble - a new approach to simulate first order phase-transition. *Physical Review Letters*, 68(9).

[4] Bieche, I., Maynard, R., Rammal, R., and Uhry, J. (1980). On the ground states of the frustration model of a spin glass by a matching method of graph theory. *J. Phys. A*, 13:2553.

[5] Binder, K. and Young, A. (1986). Spin-glasses: Experimental facts, theoretical concepts and open questions. *Rev. Mod. Phys.*, 58:801.

[6] Claiborne, J. (1990). *Mathematical Preliminaries for Computer Networking*. Wiley, New York.

[7] Dayal, P., Trebst, S., Wessel, S., ürtz, D., Troyer, M., Sabhapandit, S., and Coppersmith, S. (2004). Performance limitations of flat histogram methods and optimality of Wang-Langdau sampling. *Physical Review Letters*, 92(9):097201.

[8] Fischer, K. and Hertz, J. (1991). *Spin Glasses*. Cambridge University Press, Cambridge.

[9] Fischer, S. and Wegener, I. (2004). The ising model on the ring: Mutation versus recombination. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2004)*, pages 1113–1124.

[10] Galluccio, A. and Loebl, M. (1999a). A theory of Pfaffian orientations. I. Perfect matchings and permanents. *Electronic Journal of Combinatorics*, 6(1). Research Paper 6.

[11] Galluccio, A. and Loebl, M. (1999b). A theory of Pfaffian orientations. II. T-joins, k-cuts, and duality of enumeration. *Electronic Journal of Combinatorics*, 6(1). Research Paper 7.

[12] Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley, Reading, MA.

[13] Harik, G. R. (1995). Finding multimodal solutions using restricted tournament selection. *Proceedings of the International Conference on Genetic Algorithms (ICGA-95)*, pages 24–31.

[14] Hartmann, A. K. (1996). Cluster-exact approximation of spin glass ground states. *Physica A*, 224:480.

[15] Hartmann, A. K. (2001). Ground-state clusters of two, three and four-dimensional +/-J Ising spin glasses. *Phys. Rev. E*, 63:016106.

[16] Hartmann, A. K. and Rieger, H. (2001). *Optimization Algorithms in Physics*. Wiley-VCH, Weinheim.

[17] Hartmann, A. K. and Rieger, H., editors (2004). *New Optimization Algorithms in Physics*. Wiley-VCH, Weinheim.

[18] Hartmann, A. K. and Weigt, M. (2005). *Phase Transitions in Combinatorial Optimization Problems*. Wiley-VCH, Weinheim.

[19] Holland, J. H. (1975). *Adaptation in natural and artificial systems*. University of Michigan Press, Ann Arbor, MI.

[20] Mezard, M., Parisi, G., and Virasoro, M. (1987). *Spin glass theory and beyond*. World Scientific, Singapore.

[21] Middleton, A. and Fisher, D. S. (2002). The three-dimensional random field Ising magnet: Interfaces, scaling, and the nature of states. *Phys. Rev. B*, 65:134411.

[22] Mühlenbein, H. and Mahnig, T. (1999). Convergence theory and applications of the factorized distribution algorithm. *Journal of Computing and Information Technology*, 7(1):19–32.

[23] Mühlenbein, H. and Paaß, G. (1996). From recombination of genes to the estimation of distributions I. Binary parameters. *Parallel Problem Solving from Nature*, pages 178–187.

[24] Mühlenbein, H. and Schlierkamp-Voosen, D. (1993). Predictive models for the breeder genetic algorithm: I. Continuous parameter optimization. *Evolutionary Computation*, 1(1):25–49.

[25] Naudts, B. and Naudts, J. (1998). The effect of spin-flip symmetry on the performance of the simple GA. *Parallel Problem Solving from Nature*, pages 67–76.

[26] Newman, C. and Stein, D. (2003). Finite-dimensional spin glasses: states, excitations and interfaces. preprint cond-mat/0301022.

[27] Pelikan, M. (2005). *Hierarchical Bayesian optimization algorithm: Toward a new generation of evolutionary algorithms*. Springer-Verlag.

[28] Pelikan, M. and Goldberg, D. E. (2001). Escaping hierarchical traps with competent genetic algorithms. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 511–518.

[29] Pelikan, M. and Goldberg, D. E. (2003). Hierarchical BOA solves Ising spin glasses and MAXSAT. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2003)*, II:1275–1286.

[30] Pelikan, M., Ocenasek, J., Trebst, S., Troyer, M., and Alet, F. (2004). Computational complexity and simulation of rare events of ising spin glasses. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2004)*, 2:36–47.

[31] Picard, J.-C. and Ratliff, H. (1975). Minimum cuts and related problems. *Networks*, 5:357.

[32] Sastry, K. and Goldberg, D. E. (2002). Analysis of mixing in genetic algorithms: A survey. IlliGAL Report No. 2002012, University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory, Urbana, IL.

[33] Spin Glass Ground State Server (2004). `http://www.informatik.uni-koeln.de/ls_juenger/research/sgs/sgs.html`. University of Köln, Germany.

[34] Swamy, M. and Thulasiraman, K. (1991). *Graphs, Networks and Algorithms*. Wiley, New York.

[35] Tarjan, R. (1983). *Data Structures and Network Algorithms.* Society for industrial and applied mathematics, Philadelphia.

[36] Thierens, D., Goldberg, D. E., and Pereira, A. G. (1998). Domino convergence, drift, and the temporal-salience structure of problems. *Proceedings of the International Conference on Evolutionary Computation (ICEC-98)*, pages 535–540.

[37] Träff, J. (1996). A heuristic for blocking flow algorithms. *Eur. J. Oper. Res.*, 89:564.

[38] Van Hoyweghen, C. (2001). Detecting spin-flip symmetry in optimization problems. In Kallel, L. et al., editors, *Theoretical Aspects of Evolutionary Computing*, pages 423–437. Springer, Berlin.

[39] Wang, F. and Landau, D. P. (2001). Efficient, multiple-range random walk algorithm to calculate the density of states. *Physical Review Letters*, 86(10):2050–2053.

[40] Young, A., editor (1998). *Spin glasses and random fields.* World Scientific, Singapore.