

Order or Not: Does Parallelization of Model Building in hBOA Affect Its Scalability?

Martin Pelikan

Missouri Estimation of Distribution Algorithms
Laboratory (MEDAL)
Dept. of Mathematics and Computer Science
University of Missouri–St. Louis
One University Blvd., St. Louis MO 63121
pelikan@cs.umsl.edu

James D. Laury, Jr.

Missouri Estimation of Distribution Algorithms
Laboratory (MEDAL)
Dept. of Mathematics and Computer Science
University of Missouri–St. Louis
One University Blvd., St. Louis MO 63121
jlaury01@arch.umsl.edu

ABSTRACT

It has been shown that model building in the hierarchical Bayesian optimization algorithm (hBOA) can be efficiently parallelized by randomly generating an ancestral ordering of the nodes of the network prior to learning the network structure and allowing only dependencies consistent with the generated ordering. However, it has not been thoroughly shown that this approach to restricting probabilistic models does not affect scalability of hBOA on important classes of problems. This paper demonstrates that although the use of a random ancestral ordering restricts the structure of considered models to allow efficient parallelization of model building, its effects on hBOA performance and scalability are negligible.

Categories and Subject Descriptors

I.2.6 [Artificial Intelligence]: Learning; I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search; G.1.6 [Numerical Analysis]: Optimization

General Terms

Algorithms, Performance

Keywords

Hierarchical Bayesian optimization algorithm, evolutionary algorithms, parallelization, efficiency enhancement, decomposable problems, probabilistic models, Bayesian networks.

1. INTRODUCTION

The hierarchical Bayesian optimization algorithm (hBOA) [26, 25, 28] can solve nearly decomposable and hierarchical problems in a quadratic or subquadratic number of function evaluations. Nonetheless, low-order polynomial scalability may still be insufficient to solve

problems with a large number of decision variables or large order of interactions. That is why a number of efficiency enhancement techniques have been proposed that address the most important computational bottlenecks of hBOA: fitness evaluation and model building [21, 38, 16, 17, 33, 35, 39].

One of the most promising approaches to enhancing the efficiency of model building in hBOA is to parallelize the learning of model structure [21, 19, 17, 18], which is the most expensive part of model building. However, efficient parallelization of model building represents a challenging problem because straightforward approaches to parallelizing model building necessitate heavy communication between the computational nodes [21, 19, 20]. Nonetheless, if one restricts the class of models by generating an ancestral ordering of the variables and then considering only models consistent with the generated ordering, model building can be parallelized with nearly no communication between the computational nodes and nearly linear speedups of model building can be obtained even for large numbers of processors [21, 19].

While it has been shown that the above approach enables efficient utilization of computational resources [21, 19, 22], it has not been investigated thoroughly how it affects scalability of hBOA on important classes of problems. The purpose of this paper is to test hBOA scalability with and without the modification that restricts models according to a random topological ordering of variables generated prior to model construction. As test problems, we consider standard nearly decomposable and hierarchical problems, including concatenated traps, hierarchical traps, and Ising spin glasses.

The paper is organized as follows. Section 2 outlines hBOA and its basic components. Section 3 describes the considered method for parallelizing model construction in hBOA. Section 4 provides and discusses experimental results. Finally, section 5 summarizes and concludes the paper.

2. HIERARCHICAL BAYESIAN OPTIMIZATION ALGORITHM (HBOA)

This section outlines the hierarchical Bayesian optimization algorithm and its components. Additionally, the section briefly discusses efficiency enhancement of hBOA and other estimation of distribution algorithms (EDAs) [16, 30, 34].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '07, July 7–11, 2007, London, England, United Kingdom.
Copyright 2007 ACM 978-1-59593-697-4/07/0007 ...\$5.00.

2.1 Basic hBOA Algorithm

The hierarchical Bayesian optimization algorithm (hBOA) [26, 25] evolves a population of candidate solutions represented by fixed-length strings over a finite alphabet. Here we assume that candidate solutions are represented by n -bit binary strings. 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. We use binary tournament selection without replacement. New solutions are generated by building a Bayesian network with decision trees [3, 8] 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) [11]. The run is terminated when termination criteria are met.

2.2 Bayesian Networks with Decision Trees

To model promising solutions and sample new candidate solutions, hBOA uses Bayesian networks with decision trees [3, 8]. A Bayesian network [15, 23] consists of two components: structure and parameters. The structure consists of a directed acyclic graph where the nodes correspond to variables (string positions) and the edges represent direct conditional dependencies. A Bayesian network represents the joint probability distribution

$$p(X) = \prod_{i=0}^{n-1} p(X_i | \Pi_i), \quad (1)$$

where $X = (X_0, \dots, X_{n-1})$ is a vector of all the variables in the problem; Π_i is the set of parents of X_i in the network (the set of nodes from which there exists an edge to X_i); and $p(X_i | \Pi_i)$ is the conditional probability of X_i given its parents Π_i .

To represent conditional probabilities efficiently, hBOA uses decision trees. For each variable X_i , there is a special decision tree T_i that encodes conditional probabilities $p(X_i | \Pi_i)$; for n variables, there are n decision trees. Each internal node of the decision tree T_i is labeled by a variable X_j where $j \neq i$. Children of a node labeled by X_j correspond to disjoint subsets of the potential values of X_j ; for each value of X_j , there is one child corresponding to this value. Each traversal of a decision tree T_i for X_i thus corresponds to a constraint on the values of some other variables. Each leaf node of T_i then stores the probabilities of X_i given the constraint defined by the traversal of T_i that ends in this leaf.

For the binary alphabet, there are two children of any internal node (one child corresponds to a 0, whereas the other one corresponds to a 1) and only one probability must be stored in each leaf because $p(X_i = 0 | \Pi_i = \pi_i) + p(X_i = 1 | \Pi_i = \pi_i) = 1$ for any instance π_i of Π_i .

2.3 Learning and Sampling Bayesian Networks

Learning a Bayesian network with decision trees consists of two steps [13]: (1) learn the structure, and (2) learn the parameters (conditional probabilities).

To estimate the parameters of a Bayesian network with decision trees, hBOA uses the maximum likelihood estimate

of the probabilities in the leaves of all decision trees [13]. Consider a decision tree T_i for X_i and a leaf in this decision tree that specifies a condition C (based on the traversal). Then, the maximum likelihood estimate of $p(X_i = x_i | C)$ where x_i is a potential value of X_i , is given by

$$p(X_i = x_i | C) = \frac{m(X_i = x_i, C)}{m(C)}, \quad (2)$$

where $m(X_i = x_i, C)$ denotes the number of instances with $X_i = x_i$ that satisfy C , and $m(C)$ denotes the number of instances that satisfy C .

To learn the structure of a Bayesian network with decision trees, a simple greedy algorithm is used [13, 3]. The greedy algorithm starts with an empty network, which is represented by single-node decision trees. Each iteration splits one leaf of any decision tree. The split is selected in order to maximize the improvement of network quality. The algorithm terminates when no more improvement is possible. There are several approaches to measuring quality of network structures, for example, one can use the Bayesian-Dirichlet metric with likelihood equivalence (BDe) [4, 13, 3]. In this work, we use the BDe metric with a model-complexity penalty as described in [25, 28, 8].

Splitting a leaf may increase the number of parents of the corresponding variable, depending on whether a split on this variable has already been made in this tree. The corresponding Bayesian network structure is implicitly defined by the set of decision trees. For more details on learning BNs with decision trees, see [3, 8, 25].

The sampling of a Bayesian network with decision trees can be done using probabilistic logic sampling [14], where the ancestral ordering of the variables is first determined, in which each variable is preceded by its parents. Then, the values of all variables are generated according to the ancestral ordering.

The asymptotic time complexity of building the network structure dominates the overall complexity of the variation operator that consists of building and sampling a BN [25]. A similar behavior can be observed in other EDAs that use complex probabilistic models, for example, in the extended compact genetic algorithm (ECGA) [10, 37]. If model building is the computational bottleneck, building the network structure is thus the most important component to tackle with efficiency enhancement techniques.

2.4 Efficiency Enhancement Techniques for hBOA and Other EDAs

To solve large and complex problems, it is necessary to use scalable optimization techniques. Nonetheless, even when the time complexity of an algorithm grows only quadratically with the number of decision variables, the computation may become impractical for extremely large problems with thousands of decision variables and problems for which the evaluation of solution quality is computationally expensive. That is why it is important to develop efficiency enhancement techniques, which provide additional mechanisms for speeding up hBOA and other evolutionary algorithms [9, 36, 17, 25, 35, 39]. From the practitioner's point of view, while scalability of optimizers addresses the transition *from intractability to tractability*, efficiency enhancement techniques address the transition *from tractability to practicality* [9].

There are two potential computational bottlenecks of hBOA and other EDAs: (1) fitness evaluation and (2) model

building. Efficiency enhancement techniques that address the above bottlenecks in hBOA can be classified into the following categories [25]:

1. Parallelization.
2. Hybridization.
3. Time continuation.
4. Fitness evaluation relaxation.
5. Prior knowledge utilization.
6. Incremental and sporadic model building.
7. Learning from experience.

This paper addresses parallelization; specifically, we analyze how parallelization of model building in hBOA affects hBOA performance and scalability.

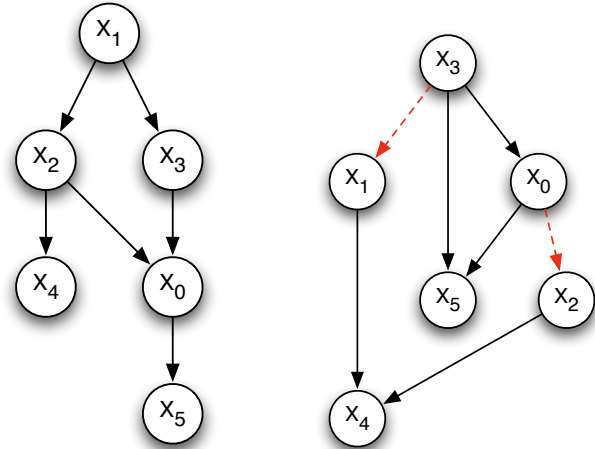
3. PARALLELIZATION OF MODEL BUILDING IN HBOA

There are several potential approaches to parallelization of the greedy algorithm for learning the Bayesian network structure [19, 17]. Probably the most efficient approach to parallelization is to distribute the nodes of the Bayesian network and let each processor identify parents of the nodes assigned to this processor. However, since the network must remain acyclic at all times, a straightforward approach to implementing this method will necessitate communication between all processors after adding a new parent to any node of the network, which will result in a highly inefficient use of parallel computer architectures.

To tackle the problem of heavy interprocessor communication, Ocenasek and Schwarz [21, 22, 20] proposed to generate a random ancestral ordering of the variables in the network before learning the network structure and to allow only network structures that are consistent with the generated ordering so that the parents of each variable are restricted to the variables that precede this variable in the generated ordering (see Figure 1). A new ancestral ordering is generated in each generation of hBOA to ensure that the bias on model structures changes over time. In this manner, the set of parents of each variable can be determined independently of the set of parents of any other variable and the different processors thus do not need to exchange any information until the model is finalized.

In addition to minimizing the communication between processors, it is important to ensure that the computational load on each processor is approximately equal. Since the computational complexity of processing each variable grows approximately linearly with the number of potential parents of the variable [21], it is best to assign nodes to processors so that each processor contains nodes with the same total number of potential parents. For example, two nodes can be assigned to each processor, where i th processor will be given the i th node from the beginning of the ancestral ordering and the i th node from the end of the ordering [21].

While restricting the models to ensure their consistency with the generated ancestral ordering ensures an efficient use of parallel computers [21, 19, 22], it also restricts the class of considered models analogically to the K2 learning algorithm [4, 13]. The models can thus be expected to represent



(a) Consistent network. (b) Inconsistent network.

Figure 1: Examples of consistent and inconsistent networks for the ancestral ordering ($X_1, X_3, X_2, X_4, X_0, X_5$). The dependencies inconsistent with the ancestral ordering are shown with dashed lines.

the actual distribution of promising solutions less accurately and the performance of hBOA may be affected, especially for problems with complex structure. The next section verifies the effects of the aforementioned restriction of model structure on hBOA performance on several important classes of nearly decomposable and hierarchical problems.

4. EXPERIMENTS

We have performed experiments on several common nearly decomposable and hierarchical problems, including concatenated deceptive problems of order 3 and 5, hierarchical traps, and the problem of finding ground states of $2D \pm J$ Ising spin glasses. The primary goal of experiments is to analyze the effects of modifying the model building algorithm according to the last section on hBOA performance; specifically, we study *scalability*, that is, how computational complexity of hBOA measured by the number of evaluations grows with problem size.

4.1 Test Problems

A brief description of the test problems follows:

- (1) *Dec-3: Concatenated deceptive of order 3.* In dec-3 [7], the input string is first partitioned into independent groups of 3 bits each. This partitioning is unknown to the algorithm, and it does not change during the run. A 3-bit deceptive function is applied to each group of 3 bits and the contributions of all deceptive functions are added together to form the fitness. Each 3-bit deceptive function is defined as follows:

$$dec(u) = \begin{cases} 1 & \text{if } u = 3 \\ 0 & \text{if } u = 2 \\ 0.8 & \text{if } u = 1 \\ 0.9 & \text{if } u = 0 \end{cases}, \quad (3)$$

where u is the number of ones in the input string of 3

bits. An n -bit dec-3 function has one global optimum in the string of all ones and $2^{n/3} - 1$ other local optima. To solve dec-3, it is necessary to consider interactions among the positions in each partition because when each bit is considered independently, the optimization is misled away from the optimum [40, 2, 29].

- (2) *Trap-5: Concatenated trap of order 5.* Trap-5 can be defined analogically to dec-3, but instead of 3-bit groups, 5-bit groups are considered. The contribution of each group of 5 bits is computed as [1, 6]

$$\text{trap}_5(u) = \begin{cases} 5 & \text{if } u = 5 \\ 4 - u & \text{otherwise} \end{cases}, \quad (4)$$

where u is the number of ones in the input string of 5 bits. An n -bit trap-5 has one global optimum in the string of all ones and $2^{n/5} - 1$ other local optima. Trap-5 also necessitates that all bits in each group are treated together, because statistics of lower order are misleading [6].

- (3) *hTrap: Hierarchical trap.* Hierarchical traps (hTraps) [24] cannot be tractably solved using single-level decomposition, but can be efficiently solved using a competent hierarchical optimizer. hTraps are created by combining trap functions of order 3 over multiple levels of difficulty. For hTraps, the string length should be an integer power of 3, that is, $n = 3^l$ where l is the number of levels.

In the variant of hTrap used in this work, on the lowest level, groups of 3 bits contribute to the overall fitness using a generalized 3-bit trap

$$\text{trap}_3(u) = \begin{cases} f_{high} & \text{if } u = 3 \\ f_{low} - u \frac{f_{low}}{2} & \text{otherwise} \end{cases}, \quad (5)$$

where $f_{high} = 1$ and $f_{low} = 1 + 0.1/l$. Note that for these values of f_{high} and f_{low} , the optimum of the trap is 000.

Each group of 3 bits corresponding to one of the traps is then mapped to a single symbol on the next level; a 000 is mapped to a 0, a 111 is mapped to a 1, and everything else is mapped to the null symbol '·'. The bits on the next level again contribute to the overall fitness using 3-bit traps defined above (see eq. 5), and the groups are mapped to an even higher level. This continues until the top level is evaluated that contains 3 bits total. However, on the top level, a trap with $f_{high} = 1$ and $f_{low} = 0.9$ is applied. As a result, the optimum of hTrap is in the string of all ones despite the superior performance of blocks of zeros on any level except for the top one. Any group of bits containing the null symbol does not contribute to the overall fitness.

To make the overall contribution at each level of the same magnitude, the contributions of traps on i th level from the bottom are multiplied by 3^i .

hTraps have many local optima, but only one global optimum in the string of all ones. Nonetheless, any single-level decomposition into subproblems of bounded order will lead away from the global optimum. That is why hTraps necessitate an optimizer that can build solutions hierarchically by juxtaposing good partial solutions over multiple levels of difficulty until the global optimum is found.

- (4) *Ising spin glass.* A 2D Ising spin glass is arranged on a regular 2D grid where each node i corresponds to a spin $s_i \in \{+1, -1\}$ and each edge $\langle i, j \rangle$ corresponds to a coupling between two spins s_i and s_j . 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. A specific set of coupling constants define a spin-glass instance. Each possible setting of all spins is called a spin configuration.

Given a set of coupling constants $\{J_{i,j}\}$, and a configuration of spins $C = \{s_i\}$ ($i = 1, \dots, n$), the energy can be computed as

$$E(C) = \sum_{\langle i,j \rangle} s_i J_{i,j} s_j, \quad (6)$$

where the sum runs over all couplings $\langle i, j \rangle$.

Here the task is to find ground states given a set of coupling constants, where a ground state is defined as a spin configuration with the minimum possible energy. Finding ground states is a challenging problem because of the rough energy landscape and a large order of interactions [12, 27, 5, 25, 32].

For each problem size, we consider 1000 random spin glass instances where each coupling is set randomly to either +1 or -1 with equal probabilities. All 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 [5].

4.2 Description of Experiments

To study scalability, we consider a range of problem sizes for each test problem. The results are then used to investigate the growth of the number of function evaluations until successful convergence to the global optimum with respect to the problem size with and without ordering the variables prior to model building.

For all problems and problem sizes except for spin glasses, bisection [36] is ran to determine the minimum population size to ensure reliable convergence to the global optimum in 30 out of 30 independent runs [30]. For spin glasses, only 5 successful runs out of 5 runs are required. Binary tournament selection is used to select promising solutions. The population of new solutions has the same size as the original population and RTR with window size $w = \min\{n, N/20\}$ is used to incorporate new solutions into the original population. To construct the model, we use the BDe metric for decision trees [3] modified by subtracting a complexity penalty term [31, 25]. A greedy network construction algorithm is used to construct network structure.

For spin glasses, we also use the deterministic hill climber (DHC) to improve all candidate solutions in the population [25, 32]. In each iteration, DHC performs a single-bit change on the solution that leads 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.

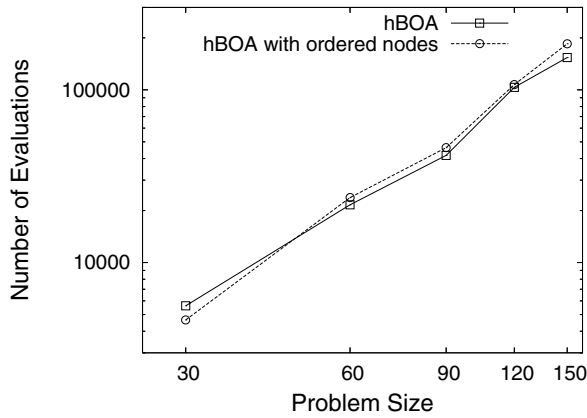


Figure 2: Results on dec-3.

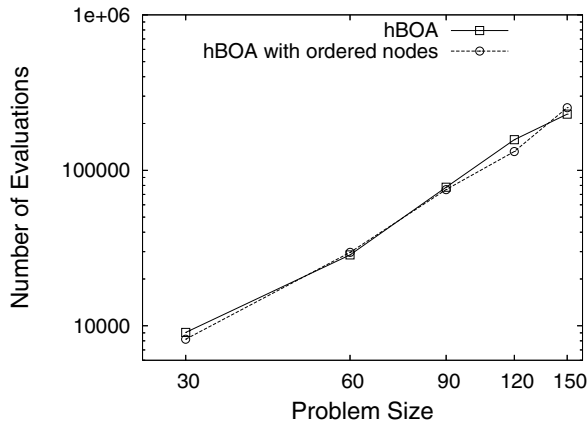


Figure 3: Results on trap-5.

4.3 Results

Figures 2 and 3 show the growth of the number of evaluations with problem size for the two separable problems of bounded difficulty: dec-3 and trap-5. The growth of the number of evaluations on dec-3 and trap-5 is also summarized in the following two tables:

Number of bits	dec-3	dec-3, ordered nodes
30	5616.00	4653.67
60	21594.00	23766.67
90	41750.00	46250.00
120	103400.00	107158.33
150	153766.67	184933.33

Number of bits	trap-5	trap-5, ordered nodes
30	9062.40	8195.83
60	28562.50	29625.00
90	77866.67	75250.00
120	157516.67	132166.67
150	229866.67	253016.67

The results on dec-3 and trap-5 indicate that for separable problems of bounded difficulty, restricting the Bayesian network according to a randomly generated ancestral ordering of variables does not affect hBOA performance significantly.

Figures 4 and 5 show the growth of the number of evaluations with problem size for test problems that cannot be

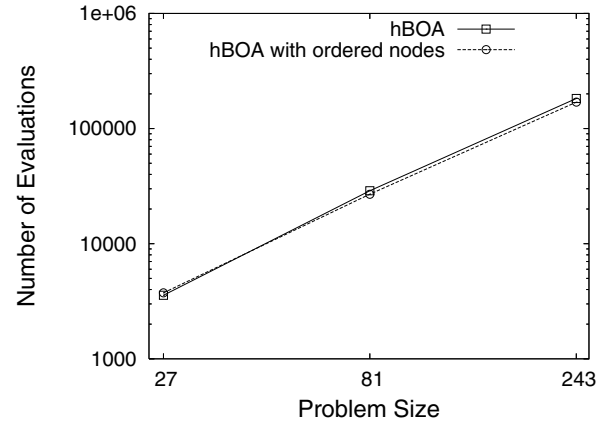


Figure 4: Results on hierarchical traps.

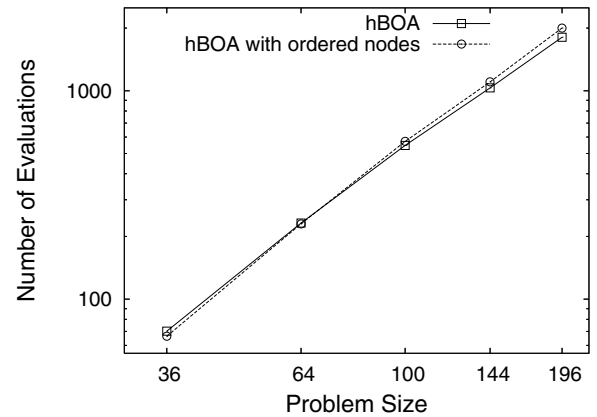


Figure 5: Results on 2D $\pm J$ Ising spin glasses.

decomposed into independent subproblems of bounded order: hTrap and 2D Ising spin glass. The growth of the number of evaluations on hTrap and 2D Ising spin glass is shown also in the following two tables:

Number of bits	hTrap	hTrap, ordered nodes
27	3567.20	3737.30
81	28833.33	26923.13
243	182466.67	169416.67

Number of bits	spin glass	spin glass, ordered nodes
$6 \times 6 = 36$	70.19	66.40
$8 \times 8 = 64$	231.86	229.95
$10 \times 10 = 100$	546.88	571.62
$12 \times 12 = 144$	1032.54	1103.11
$14 \times 14 = 196$	1806.87	2001.57

Analogously to the results on dec-3 and trap-5, restricting the probabilistic models does not affect hBOA performance on hTrap and 2D Ising spin glasses significantly.

Although the unpaired t-test revealed statistical significance of the differences obtained with and without ordering the variables prior to model building, the most important observation is that (1) the differences are relatively small with respect to the absolute values and (2) the order of the polynomial approximating the growth of the number of evaluations is affected only a little or not at all in all

test cases. We can thus conclude that restricting the probabilistic models in hBOA to models that are consistent with a randomly generated ancestral ordering does not significantly affect hBOA performance. Consequently, using this modification will ensure efficient parallelization of model building when model building becomes the computational bottleneck.

5. SUMMARY AND CONCLUSIONS

To efficiently parallelize model building in hBOA, it is useful to generate an ancestral ordering of nodes in the network prior to building the network structure, and to restrict models to those consistent with the generated ancestral ordering. While modifying the model building procedure allows its efficient parallel implementation, it also restricts the considered class of network structures, potentially affecting performance and scalability of hBOA. This paper provides empirical evidence that the modification of the model building algorithm does not significantly affect hBOA performance and scalability.

Acknowledgments

This work was supported by the National Science Foundation under CAREER grant ECS-0547013, by the Air Force Office of Scientific Research, the Air Force Materiel Command, USAF, under grant FA9550-06-1-0096, and by the University of Missouri in St. Louis through the High Performance Computing Collaboratory sponsored by Information Technology Services, and the Research Award and Research Board programs. The experiments were done using the hBOA software developed by Martin Pelikan and David E. Goldberg at the University of Illinois at Urbana-Champaign.

The U.S. Government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation thereon. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation, the Air Force Office of Scientific Research, or the U.S. Government.

6. REFERENCES

- [1] D. H. Ackley. An empirical study of bit vector function optimization. *Genetic Algorithms and Simulated Annealing*, pages 170–204, 1987.
- [2] P. A. N. Bosman and D. Thierens. Linkage information processing in distribution estimation algorithms. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-99)*, I:60–67, 1999.
- [3] D. M. Chickering, D. Heckerman, and C. Meek. A Bayesian approach to learning Bayesian networks with local structure. Technical Report MSR-TR-97-07, Microsoft Research, Redmond, WA, 1997.
- [4] G. F. Cooper and E. H. Herskovits. A Bayesian method for the induction of probabilistic networks from data. *Machine Learning*, 9:309–347, 1992.
- [5] P. Dayal, S. Trebst, S. Wessel, D. Würtz, M. Troyer, S. Sabhapandit, and S. Coppersmith. Performance limitations of flat histogram methods and optimality of Wang-Langdau sampling. *Physical Review Letters*, 92(9):097201, 2004.
- [6] K. Deb and D. E. Goldberg. Analyzing deception in trap functions. IlliGAL Report No. 91009, University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory, Urbana, IL, 1991.
- [7] K. Deb and D. E. Goldberg. Sufficient conditions for deceptive and easy binary functions. *Annals of Mathematics and Artificial Intelligence*, 10:385–408, 1994.
- [8] N. Friedman and M. Goldszmidt. Learning Bayesian networks with local structure. In M. I. Jordan, editor, *Graphical models*, pages 421–459. MIT Press, Cambridge, MA, 1999.
- [9] D. E. Goldberg. *The design of innovation: Lessons from and for competent genetic algorithms*, volume 7 of *Genetic Algorithms and Evolutionary Computation*. Kluwer Academic Publishers, 2002.
- [10] G. Harik. Linkage learning via probabilistic modeling in the ECGA. IlliGAL Report No. 99010, University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory, Urbana, IL, 1999.
- [11] G. R. Harik. Finding multimodal solutions using restricted tournament selection. *Proceedings of the International Conference on Genetic Algorithms (ICGA-95)*, pages 24–31, 1995.
- [12] A. K. Hartmann. Ground-state clusters of two, three and four-dimensional +/-J Ising spin glasses. *Phys. Rev. E*, 63:016106, 2001.
- [13] D. Heckerman, D. Geiger, and D. M. Chickering. Learning Bayesian networks: The combination of knowledge and statistical data. Technical Report MSR-TR-94-09, Microsoft Research, Redmond, WA, 1994.
- [14] M. Henrion. Propagating uncertainty in Bayesian networks by probabilistic logic sampling. In J. F. Lemmer and L. N. Kanal, editors, *Uncertainty in Artificial Intelligence*, pages 149–163. Elsevier, Amsterdam, London, New York, 1988.
- [15] R. A. Howard and J. E. Matheson. Influence diagrams. In R. A. Howard and J. E. Matheson, editors, *Readings on the principles and applications of decision analysis*, volume II, pages 721–762. Strategic Decisions Group, Menlo Park, CA, 1981.
- [16] P. Larrañaga and J. A. Lozano, editors. *Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation*. Kluwer, Boston, MA, 2002.
- [17] A. Mendiburu, J. A. Lozano, and J. Miguel-Alonso. Parallel estimation of distribution algorithms: New approaches. Technical Report Technical Report EHU-KAT-IK-1-3, Department of Computer Architecture and Technology, The University of the Basque Country, 2003.
- [18] A. Mendiburu, J. Miguel-Alonso, J. A. Lozano, M. Ostra, and C. Ubide. Parallel edas to create multivariate calibration models for quantitative chemical applications. *Journal of Parallel and Distributed Computing*, 66(8):1002–1013, 2006.
- [19] J. Ocenasek. *Parallel Estimation of Distribution Algorithms*. PhD thesis, Faculty of Information Technology, Brno University of Technology, Brno, 2002.
- [20] J. Ocenasek, E. Cantú-Paz, M. Pelikan, and J. Schwarz. Design of parallel estimation of

- distribution algorithms. In M. Pelikan, K. Sastry, and E. Cantú-Paz, editors, *Scalable Optimization via Probabilistic Modeling: From Algorithms to Applications*. Springer, 2006.
- [21] J. Ocenasek and J. Schwarz. The parallel Bayesian optimization algorithm. In *Proceedings of the European Symposium on Computational Intelligence*, pages 61–67. Physica-Verlag, 2000.
- [22] J. Ocenasek, J. Schwarz, and M. Pelikan. Design of multithreaded estimation of distribution algorithms. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2003)*, pages 1247–1258, 2003.
- [23] J. Pearl. *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Morgan Kaufmann, San Mateo, CA, 1988.
- [24] M. Pelikan. *Bayesian optimization algorithm: From single level to hierarchy*. PhD thesis, University of Illinois at Urbana-Champaign, Urbana, IL, 2002. Also IlliGAL Report No. 2002023.
- [25] M. Pelikan. *Hierarchical Bayesian optimization algorithm: Toward a new generation of evolutionary algorithms*. Springer-Verlag, 2005.
- [26] M. Pelikan and D. E. Goldberg. Escaping hierarchical traps with competent genetic algorithms. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 511–518, 2001. Also IlliGAL Report No. 2000020.
- [27] M. Pelikan and D. E. Goldberg. Hierarchical BOA solves Ising spin glasses and MAXSAT. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2003)*, II:1275–1286, 2003. Also IlliGAL Report No. 2003001.
- [28] M. Pelikan and D. E. Goldberg. Hierarchical bayesian optimization algorithm. In E. Cantú-Paz, M. Pelikan, and K. Sastry, editors, *Scalable optimization via probabilistic modeling: From algorithms to applications*, pages 63–90. Springer, 2006.
- [29] M. Pelikan, D. E. Goldberg, and E. Cantú-Paz. BOA: The Bayesian optimization algorithm. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-99)*, I:525–532, 1999. Also IlliGAL Report No. 99003.
- [30] M. Pelikan, D. E. Goldberg, and F. Lobo. A survey of optimization by building and using probabilistic models. *Computational Optimization and Applications*, 21(1):5–20, 2002. Also IlliGAL Report No. 99018.
- [31] M. Pelikan, D. E. Goldberg, and K. Sastry. Bayesian optimization algorithm, decision graphs, and Occam’s razor. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 519–526, 2001. Also IlliGAL Report No. 2000020.
- [32] M. Pelikan and A. K. Hartmann. Searching for ground states of Ising spin glasses with hierarchical BOA and cluster exact approximation. In E. Cantú-Paz, M. Pelikan, and K. Sastry, editors, *Scalable optimization via probabilistic modeling: From algorithms to applications*, pages 333–349. Springer, 2006.
- [33] M. Pelikan and K. Sastry. Fitness inheritance in the Bayesian optimization algorithm. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2004)*, 2:48–59, 2004.
- [34] M. Pelikan, K. Sastry, and E. Cantú-Paz, editors. *Scalable optimization via probabilistic modeling: From algorithms to applications*. Springer-Verlag, 2006.
- [35] M. Pelikan, K. Sastry, and D. E. Goldberg. Sporadic model building for efficiency enhancement of hierarchical boa. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2006)*, pages 405–412, 2006.
- [36] K. Sastry. Evaluation-relaxation schemes for genetic and evolutionary algorithms. Master’s thesis, University of Illinois at Urbana-Champaign, Department of General Engineering, Urbana, IL, 2001. Also IlliGAL Report No. 2002004.
- [37] K. Sastry and D. E. Goldberg. On extended compact genetic algorithm. IlliGAL Report No. 2000026, University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory, Urbana, IL, 2000.
- [38] K. Sastry, D. E. Goldberg, and M. Pelikan. Don’t evaluate, inherit. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2001)*, pages 551–558, 2001. Also IlliGAL Report No. 2001013.
- [39] K. Sastry, M. Pelikan, and D. E. Goldberg. Efficiency enhancement of estimation of distribution algorithms. In M. Pelikan, K. Sastry, and E. Cantú-Paz, editors, *Scalable Optimization via Probabilistic Modeling: From Algorithms to Applications*, pages 161–185. Springer, 2006.
- [40] D. Thierens. *Analysis and design of genetic algorithms*. PhD thesis, Katholieke Universiteit Leuven, Leuven, Belgium, 1995.