

Analysis of Estimation of Distribution Algorithms and Genetic Algorithms on NK Landscapes

Martin Pelikan

Missouri Estimation of Distribution Algorithms Laboratory (MEDAL)
Dept. of Mathematics and Computer Science, 320 CCB
Univ. of Missouri in St. Louis
One University Blvd., St. Louis, MO 63121
pelikan@cs.umsl.edu

ABSTRACT

This study analyzes performance of several genetic and evolutionary algorithms on randomly generated NK fitness landscapes with various values of n and k . A large number of NK problem instances are first generated for each n and k , and the global optimum of each instance is obtained using the branch-and-bound algorithm. Next, the hierarchical Bayesian optimization algorithm (hBOA), the univariate marginal distribution algorithm (UMDA), and the simple genetic algorithm (GA) with uniform and two-point crossover operators are applied to all generated instances. Performance of all algorithms is then analyzed and compared, and the results are discussed.

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

1. INTRODUCTION

NK fitness landscapes [13, 14] were introduced by Kauffman as tunable models of rugged fitness landscape. An NK landscape is a function defined on binary strings of fixed length and is characterized by two parameters: (1) n for the overall number of bits and (2) k for the neighborhood size. For each bit, k neighbors are specified and a function is given that determines the fitness contribution of the bit and its neighbors. Usually, both the neighbors as well as the subfunctions are initialized randomly in some way. NK landscapes are NP-complete for $k > 1$, although some variants of NK landscapes are polynomially solvable and there exist approximation algorithms for others [23, 8, 6]. Nonetheless, NK landscapes remain a challenge for any optimization

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'08, July 12–16, 2008, Atlanta, Georgia, USA.
Copyright 2008 ACM 978-1-60558-130-9/08/07...\$5.00.

algorithm and they are also interesting from the perspective of complexity theory and computational biology; that is why since their inception NK landscapes have attracted researchers in all these areas [14, 2, 23, 8, 1, 6].

This paper presents an in-depth empirical performance analysis of various genetic and evolutionary algorithms on NK landscapes with varying n and k . For each value of n and k , a large number of problem instances are first generated. Then, the branch-and-bound algorithm is applied to each of these instances to provide a guaranteed global optimum of this instance. Several genetic and evolutionary algorithms are then applied to all generated problem instances and their performance is analyzed and compared. More specifically, we consider the hierarchical Bayesian optimization algorithm (hBOA), the univariate marginal distribution algorithm (UMDA), and the simple genetic algorithm (GA) with bit-flip mutation, and uniform or two-point crossover operator. Additionally, GA without any crossover is considered. The results provide insight into the difficulty of NK landscapes with respect to the parameters n and k and performance differences between all compared algorithms.

The paper starts by describing NK landscapes and the branch-and-bound algorithm used to verify the global optima of generated NK landscapes in section 2. Section 3 outlines compared algorithms. Section 4 presents experimental results. Section 5 discusses future work. Finally, section 6 summarizes and concludes the paper.

2. NK LANDSCAPES

This section describes NK landscapes and the branch-and-bound algorithm, which was used to obtain global optima of all NK problem instances considered in this paper.

2.1 Problem Definition

An NK fitness landscape [13, 14] is fully defined by the following components: (1) The number of bits, n . (2) The number of neighbors per bit, k . (3) A set of k neighbors $\Pi(X_i)$ for the i -th bit, X_i , for every $i \in \{0, \dots, n-1\}$. (4) A subfunction f_i defining a real value for each combination of values of X_i and $\Pi(X_i)$ for every $i \in \{0, \dots, n-1\}$. Typically, each subfunction is defined as a lookup table with 2^{k+1} values. The task is to maximize the objective function

$$f_{nk}(X_0, X_1, \dots, X_{n-1}) = \sum_{i=0}^{n-1} f_i(X_i, \Pi(X_i)).$$

The difficulty of optimizing NK landscapes depends on all

of the four components defining an NK problem instance. One useful approach to analyzing complexity of NK landscapes is to focus on the influence of k on problem complexity. For $k = 0$, NK landscapes are simple unimodal functions similar to onemax or binint, which can be solved in linear time and should be easy for practically any genetic and evolutionary algorithm. The global optimum of NK landscapes can be obtained in polynomial time [23] even for $k = 1$; on the other hand, for $k > 1$, the problem of finding the global optimum of unrestricted NK landscapes is NP-complete [23]. The problem becomes polynomially solvable with dynamic programming even for $k > 1$ if the neighbors are restricted to only adjacent string positions (using circular strings) [23] or if the subfunctions are generated according to some distributions [8]. For unrestricted NK landscapes with $k > 1$, a polynomial-time approximation algorithm exists with the approximation threshold $1 - 1/2^{k+1}$ [23].

2.2 Generating Random NK Instances

Typically, both the neighbors as well as the lookup tables defining the subfunctions are generated randomly. In this paper, for each string position X_i , we first generate a random set of k neighbors where each string position except for X_i is selected with equal probability. Then, the lookup table defining f_i is generated using the uniform distribution over $\{0, 1\}$.

Consequently, the studied class of NK landscapes is NP-complete for any $k > 1$. Since the case for $k = 1$ is extremely simple to solve, we only considered $k > 1$; specifically, we considered $k = 2$ to 6 with step 1. To study scalability of various evolutionary algorithms, for each k , we considered a range of values of n with the minimum value of $n = 20$ and the maximum value bounded mainly by the available computational resources and the scope of the empirical analysis.

2.3 Branch and Bound

The basic idea of branch and bound is to recursively explore all possible binary strings of n bits using a recursion tree where each level corresponds to one of the bits and the subtrees below each level correspond to the different values of the bit corresponding to this level. To make the algorithm more efficient, some subtrees are cut if they can be proven to not lead to any solution that is better than the best-so-far solution found.

Before running the branch-and-bound algorithm, we first use a simple hill climber based on bit-flip mutation with several random restarts to locate high-quality local optima. The best of the discovered optima is then used as the best-so-far solution when the branch-and-bound algorithm is started. In the branch-and-bound approach used in this paper, the bits are assigned sequentially from X_0 to X_n , although reordering the bits might improve performance under some conditions.

When processing a node at level i , the best value we can obtain by setting the remaining $n - i$ bits is given by $\max_{x_i, \dots, x_n \in \{0, 1\}^{n-i}} f_{nk}(X_1 = x_1, \dots, X_n = x_n)$ where bits x_0 to x_{i-1} are assumed to be fixed to the values defined by the path from the root of the recursion tree to the current node. If a solution has been found already that has a higher fitness than this maximum possible value, the processing below the current node does not have to continue.

The aforescribed branch-and-bound algorithm is complete and it is thus guaranteed to find the global optimum

of any problem instance. Nonetheless, the complexity of branch and bound can be expected to grow exponentially fast and solving large NK instances becomes intractable with this algorithm. While the evolutionary algorithms presented in the next section should be capable of reliably solving larger instances, their convergence to the global optimum cannot be guaranteed (unless we know the global optimum a priori).

3. COMPARED ALGORITHMS

This section outlines the optimization algorithms discussed in this paper: (1) the hierarchical Bayesian optimization algorithm (hBOA) [18, 17], (2) the univariate marginal distribution algorithm (UMDA) [16], and (3) the genetic algorithm (GA) [12, 9]. Additionally, the section describes the deterministic hill climber (DHC) [19], which is incorporated into all compared algorithms to improve their performance. In all compared algorithms, candidate solutions are represented by n -bit binary strings and restricted tournament replacement (RTR) [10] is used for effective diversity maintenance.

3.1 Genetic Algorithm

The genetic algorithm (GA) [12, 9] evolves a population of candidate solutions typically represented by fixed-length binary strings. The first population is generated at random. Each iteration starts by selecting promising solutions from the current population. We use binary tournament selection. 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 uniform or two-point crossover, and bit-flip mutation [9]. The new candidate solutions are incorporated into the original population using restricted tournament replacement (RTR) [10]. The run is terminated when termination criteria are met.

3.2 UMDA

The basic procedure of the univariate marginal distribution algorithm (UMDA) [16] is similar to that of the GA. However, instead of using crossover and mutation, UMDA creates new candidate solutions by building and sampling a probabilistic model in the form of the probability vector. After selection, the probability vector is first learned that stores the proportion of 1s in each position of the selected population. Then, 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.

UMDA is an estimation of distribution algorithm (EDA) [3, 16, 15, 20]. EDAs—also called probabilistic model-building genetic algorithms (PMBGAs) [20] and iterated density estimation algorithms (IDEAs) [4]—replace standard variation operators of genetic algorithms such as crossover and mutation by building a probabilistic model of promising solutions and sampling the built model to generate new candidate solutions. The only difference between the GA and UMDA is in the way the selected solutions are processed to generate new solutions.

3.3 Hierarchical BOA (hBOA)

The hierarchical Bayesian optimization algorithm (hBOA) [18, 17] is also an EDA and the basic procedure of hBOA is similar to that of the UMDA variant described earlier. However, to model promising solutions and sample new solutions, Bayesian networks with local structures [5, 7] are used instead of the simple probability vector of UMDA. Similarly as in the considered GA and UMDA variants, the new candidate solutions are incorporated into the original population using RTR and the run is terminated when termination criteria are met.

3.4 Deterministic Hill Climber (DHC)

The deterministic hill climber (DHC) is incorporated into GA, UMDA and hBOA 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. 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.

4. EXPERIMENTS

This section describes experiments and presents experimental results. First, problem instances and experimental setup are discussed. Next, the analysis of hBOA, UMDA and several GA variants is presented. Finally, all algorithms are compared and the results are discussed.

4.1 Problem Instances

NK instances for $k = 2$ to $k = 6$ with step 1 were studied. The only restriction on problem size was the efficiency of the branch-and-bound algorithm, the complexity of which grew very fast with n . For $k = 2$, we considered $n = 20$ to $n = 52$ with step 2; for $k = 3$, we considered $n = 20$ to $n = 48$ with step 2; for $k = 4$, we considered $n = 20$ to $n = 40$ with step 2; for $k = 5$, we considered $n = 20$ to $n = 38$ with step 2; for $k = 6$, we considered $n = 20$ to $n = 32$ with step 2.

For each combination of n and k , we generated 10,000 random problem instances and for each instance we used the branch-and-bound algorithm to locate the global optimum. Then, we applied hBOA, UMDA and several GA variants to each of these instances and collected empirical results, which were subsequently analyzed. That means that overall 600,000 unique problem instances were generated and all of them were tested with every algorithm included in this study.

4.2 Experimental Setup

To select promising solutions, binary tournament selection is used. New solutions (offspring) are incorporated into the old population using RTR with window size $w = \min\{n, N/5\}$ as suggested in [17]. In hBOA, Bayesian networks with decision trees [5, 7, 17] are used and the models are evaluated using the Bayesian-Dirichlet metric with likelihood equivalence [11, 5] and a penalty for model complexity [7, 17]. All GA variants use bit-flip mutation with the probability of flipping each bit $p_m = 1/n$. Two common crossover operators are considered in a GA: two-point and uniform crossover. For both crossover operators, the probability of applying crossover is set to 0.6. To emphasize the importance of using crossover, the results for GA

without any crossover are also included, where only bit-flip mutation is used. A stochastic hill climber with bit-flip mutation has also been considered in the initial stage, but the performance of this algorithm was far inferior compared to any other algorithm included in the comparison and it was intractable to solve most problem instances included in the comparison; that is why the results for this algorithm are omitted.

For each problem instance and each algorithm, an adequate population size is approximated with the bisection method [22, 17], which estimates the minimum population size required for reliable convergence to the optimum. Here, the bisection method finds an adequate population size for the algorithms to find the optimum in 10 out of 10 independent runs. Each run is terminated when the global optimum has been found. The results for each problem instance comprise of the following statistics: (1) the population size, (2) the number of iterations (generations), (3) the number of evaluations, and (4) the number of flips of DHC. However, due to space limitations, we only present two of these statistics—the number of evaluations and the number of flips. For each value of n and k , all observed statistics were averaged over the 10,000 random instances. Since for each instance, 10 successful runs were performed, for each n and k and each algorithm the results are averaged over 100,000 successful runs. Overall, for each algorithm, the results correspond to 6,000,000 successful runs on a total of 600,000 unique problem instances.

4.3 Performance Analysis

Figure 1 shows the average number of fitness evaluations and the average number of DHC flips for all compared algorithms on NK problem instances for $k = 2$ to $k = 6$. Due to space limitations, we omit the results of GA with no crossover, which was outperformed by all other algorithms. As expected, performance of all algorithms gets worse with increasing k . More specifically, the population size, the number of iterations, the number of evaluations, and the number of DHC flips appear all to grow exponentially with k (the figure shows only graphs for the latter two statistics). For a fixed k , the time complexity appears to grow with n slightly faster than polynomially regardless of whether it is measured by the number of evaluations or the number of flips.

4.4 Comparison of All Algorithms

To compare performance of algorithms A and B , for each problem instance, we can compute the ratio of the number of evaluations required by A and the number of evaluations required by B ; analogically, we can compute the ratio of the number of flips required by A and B . Then, the ratios can be averaged over all instances with specific n and k . If A performs better than B , the computed ratios should be smaller than 1; if A performs the same as B , the ratios should be about 1; finally, if the A performs worse than B , the ratios should be greater than 1.

A comparison based on the aforementioned ratio was computed for each pair of algorithms studied in this work. To make the results easier to read, the superior algorithm was typically used as the second algorithm in the comparison (in the denominator of the ratios), so that the ratios should be expected to be greater than 1.

Figure 2 compares performance of hBOA with that of UMDA, GA with two-point crossover, and GA with uniform

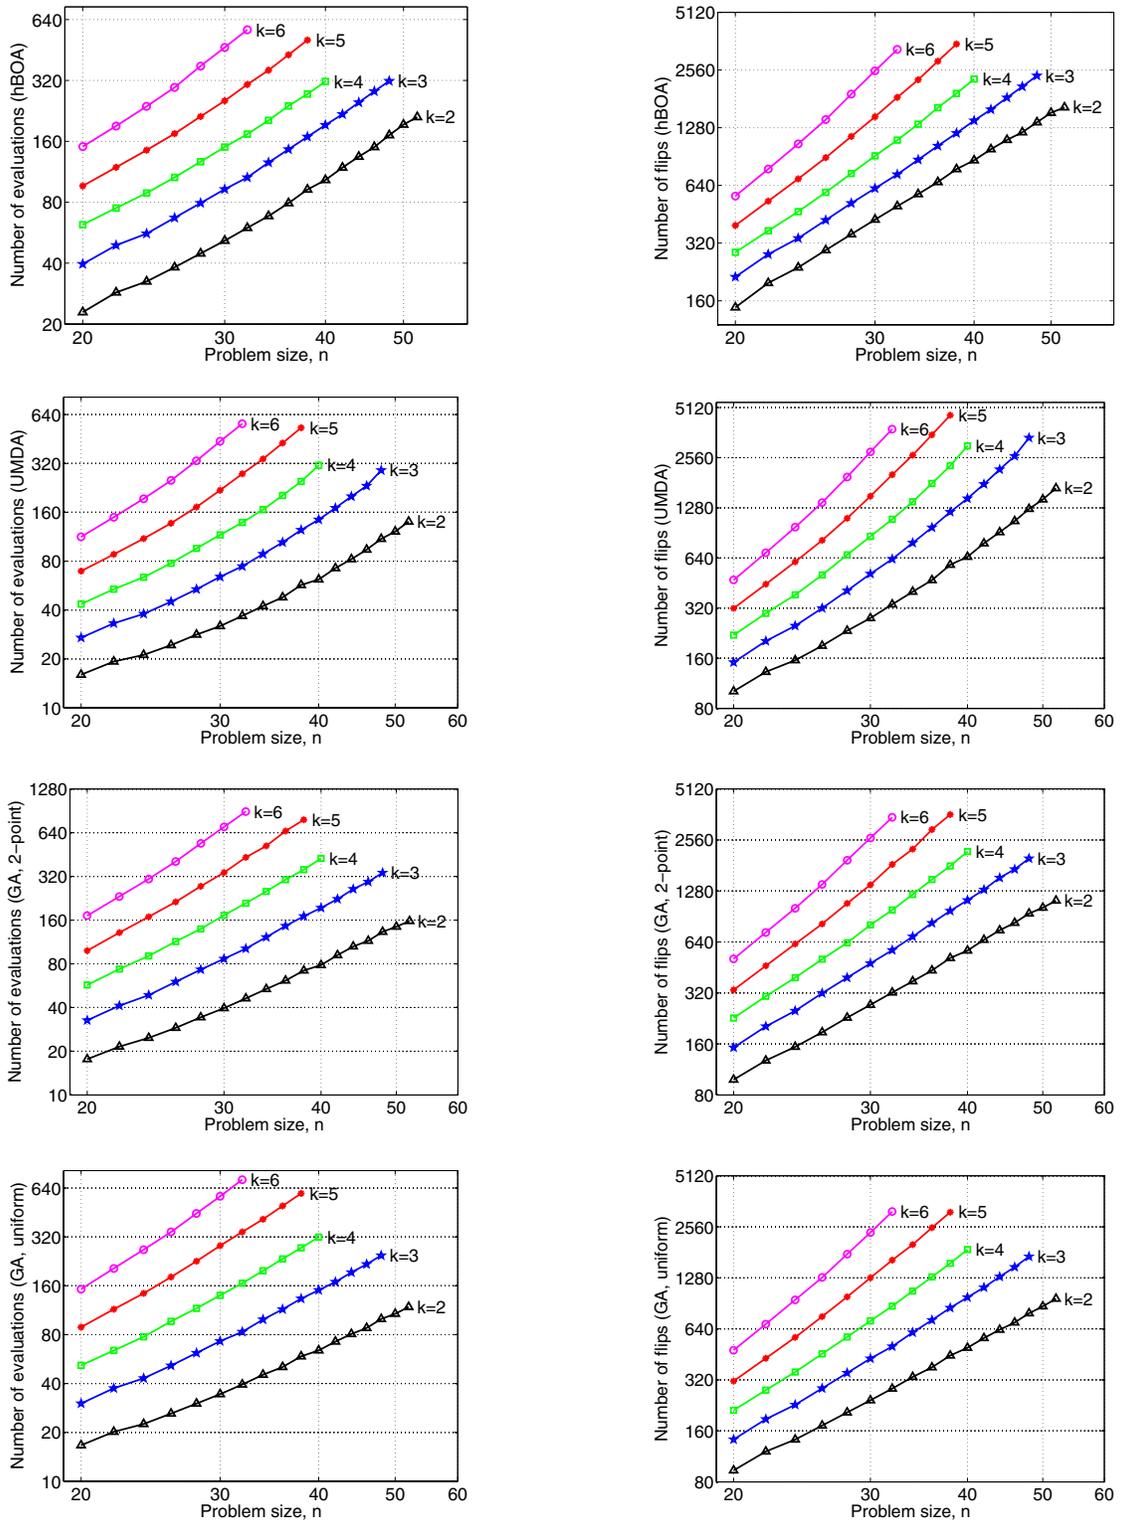


Figure 1: Performance of all compared algorithms except for GA with no crossover on NK landscapes with $k = 2$ to 6 (log-log scale).

crossover. Figure 4 compares performance of GA with two-point crossover and that of GA with uniform crossover. Figures 5 and 6 compare performance of GAs with and without crossover. Finally, figure 3 compares performance of UMDA and GA with two-point crossover. Some results of pairwise comparisons were omitted due to space limitations; nonetheless, the omitted results followed from those included here.

One of the important trends to observe in the results of the comparisons is the change in the two ratios with problem size. In most cases, when one algorithm outperforms another one, the differences become more significant as the problem size increases. In some cases, although one algorithm outperforms another one on small problems, because of the observed dynamics with problem size, we can expect the situation to reverse for large problems.

The comparisons based on the number of evaluations and the number of flips can be summarized as follows:

hBOA. While for small values of k , hBOA is outperformed by other algorithms included in the comparison, as k increases, the situation changes rapidly. More specifically, for larger k , hBOA outperforms all other algorithms and its relative performance with respect to other algorithms improves with increasing problem size. The larger the k , the more favorably hBOA compares to other algorithms.

GA with uniform crossover. GA with uniform crossover performs better than GA with two-point crossover and UMDA regardless of k and its relative performance with respect to these algorithms improves with problem size. However, as mentioned above, for larger values of k , GA with uniform crossover is outperformed by hBOA and the factor by which hBOA outperforms GA with uniform crossover grows with problem size.

GA with two-point crossover. GA with two-point crossover performs worse than hBOA and GA with uniform crossover for larger values of k , but it still outperforms UMDA with respect to the number of flips, which is the most important performance measure.

UMDA. UMDA performs worst of all recombination-based algorithms included in the comparison except for a few cases with small values of k .

Crossover versus mutation. Crossover has proven to outperform pure mutation, which is clear from all the results. First of all, for the most difficult instances, hBOA—which is a pure selectorecombinative evolutionary algorithm with no explicit mutation—outperforms other algorithms with increasing n . Second, eliminating crossover from GA significantly decreases its efficiency and the mutation-based approaches perform worst of all compared algorithms. Specifically, GA with no crossover is outperformed by all other variants of GA, and the stochastic hill climbing is not even capable of solving many problem instances in practical time.

5. FUTURE WORK

There are several interesting ways of extending the work presented in this paper. First of all, the problem instances

generated in this work can be used for analyzing performance of other optimization algorithms and comparing different optimization algorithms on a broad class of problems with tunable difficulty. Second, the class of problems considered in this study can be extended substantially using genetic and evolutionary algorithms with adequate settings for solving instances unsolvable with branch and bound. Although the global optimum would no longer be guaranteed, methods can be devised that still guarantee that the global optimum is found reliably as was done for example in ref. [21] for Sherrington-Kirkpatrick spin glasses. Finally, other probability distributions for generating NK problem instances can be considered to provide further insights into the difficulty of various classes of NK landscapes and the benefits and costs of using alternative optimization strategies in each of these classes.

6. SUMMARY AND CONCLUSIONS

This paper presented an in-depth empirical performance study of several genetic and evolutionary algorithms on NK landscapes with various values of n and k . Specifically, the algorithms considered in this work included hBOA, UMDA, and several GA variants. For each value of n and k , a large number of NK instances were generated and solved with the branch-and-bound algorithm, which is a complete algorithm that is guaranteed to find the global optimum. Performance of all algorithms was analyzed and compared.

The main contributions of this work are summarized in what follows. First of all, NK landscapes represent an important class of test problems and despite that there has been practically no work on using advanced estimation of distribution algorithms (EDAs) on NK landscapes. This work provides many experimental results on one advanced and one simple EDA, and it shows that advanced EDAs can significantly outperform other genetic and evolutionary algorithms on NK landscapes for larger values of k . Second, most studies concerned with NK landscapes do not verify the global optimum of the considered problem instances and it is thus often difficult to interpret the results and evaluate their importance. In this study, the global optimum of each instance is verified with the complete branch-and-bound algorithm. Third, while the difficulty of NK landscapes can be expected to vary substantially from instance to instance, most studies presented in the past used only a limited sample of problem instances; here we provide an in-depth study where about 600,000 unique problem instances are considered. Finally, the results in this paper are not based on only one evolutionary algorithm; instead, we consider several qualitatively different evolutionary algorithms, providing insight into the comparison of genetic algorithms and EDAs, as well as into the comparison of the mutation-based and recombination-based evolutionary algorithms.

The NK problem generator, the branch-and-bound algorithm and some of the NK problem instances used in this paper can be downloaded from <http://medal.cs.ums1.edu/>.

Acknowledgments

This project was sponsored by the National Science Foundation under CAREER grant ECS-0547013, by the Air Force Office of Scientific Research, Air Force Materiel Command, USAF, under grant FA9550-06-1-0096, and by the Univ. of Missouri in St. Louis through the High Performance Computing Collaboratory sponsored by Information Technology Services, and the Research

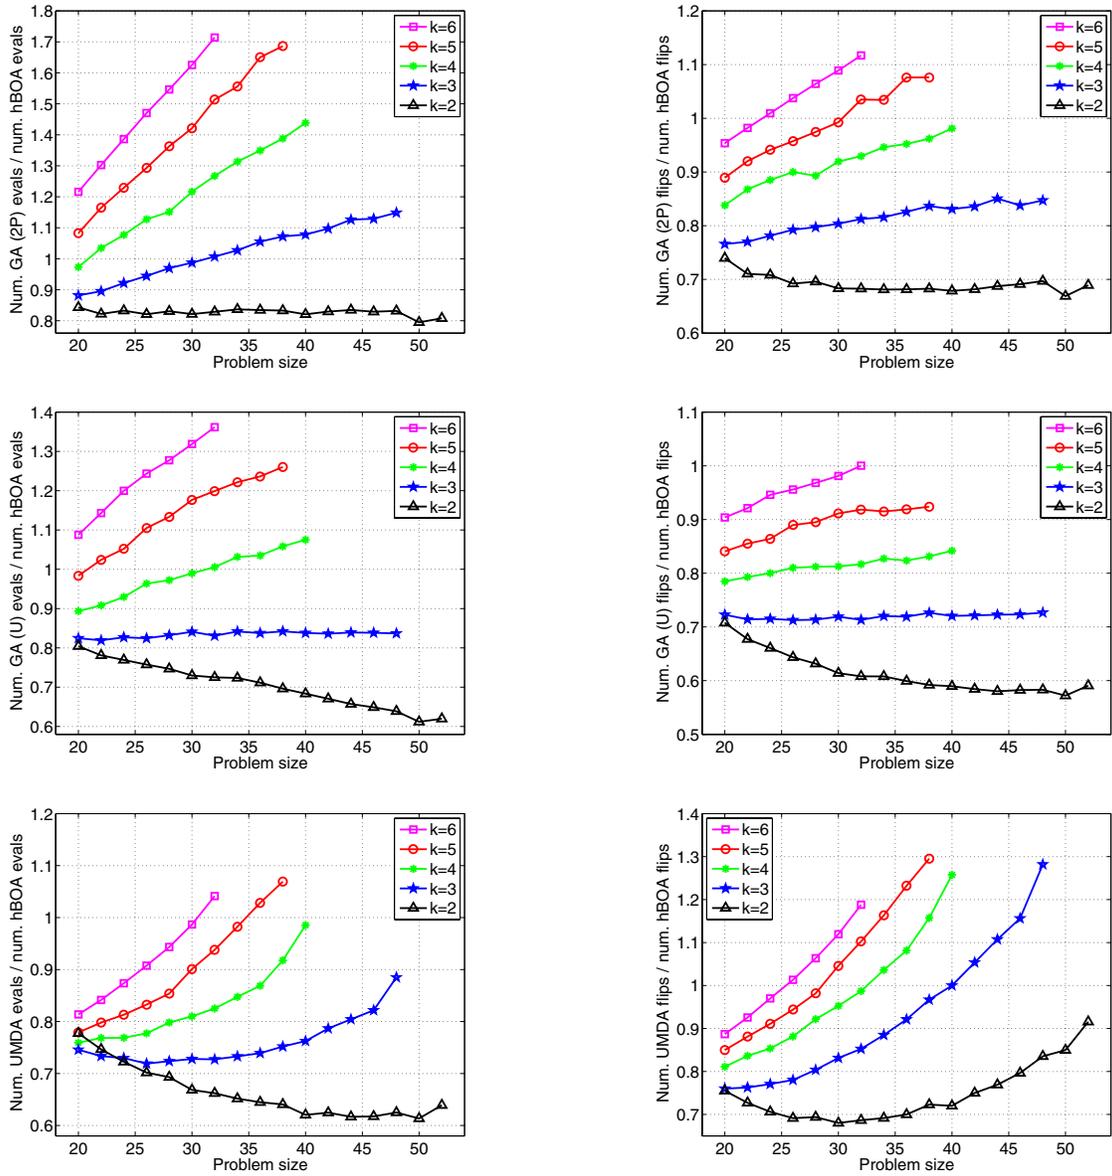


Figure 2: Comparison of hBOA with UMDA and two GA variants (uniform and two-point crossover) with respect to the number of evaluations and the number of flips.

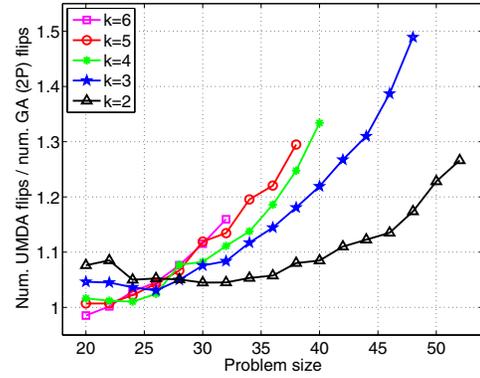
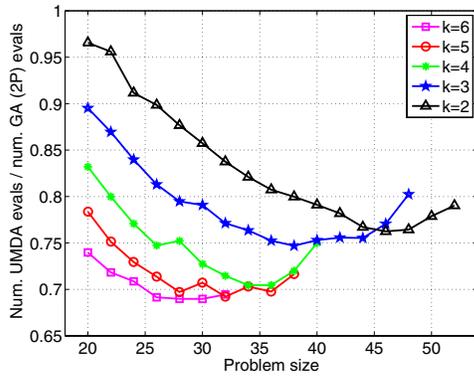


Figure 3: Comparison of GA with two-point crossover and UMDA with respect to the number of evaluations and the number of flips.

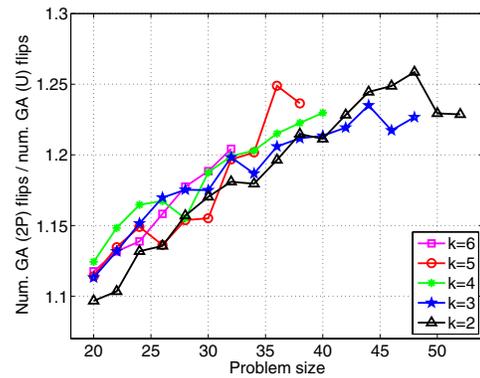
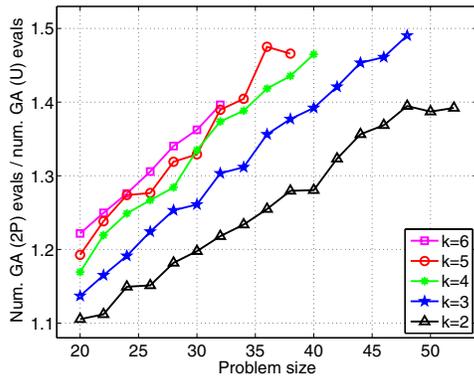


Figure 4: Comparison of GA with uniform and two-point crossover with respect to the number of evaluations and the number of flips.

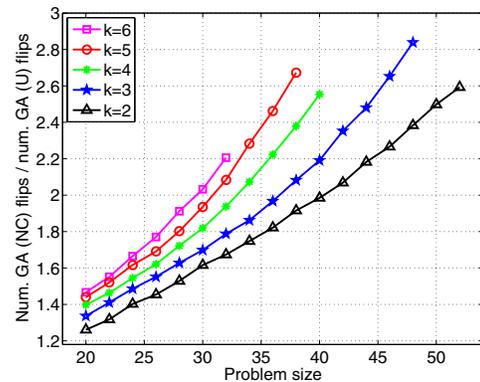
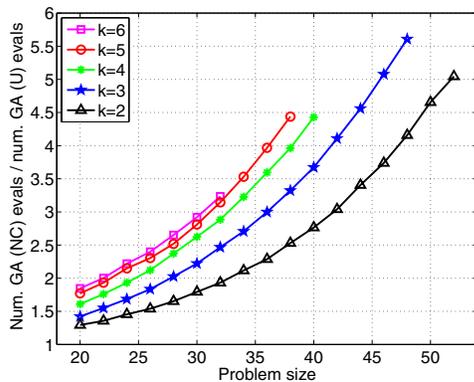


Figure 5: Comparison of GA with uniform crossover and GA with no crossover (mutation only) with respect to the number of evaluations and the number of flips.

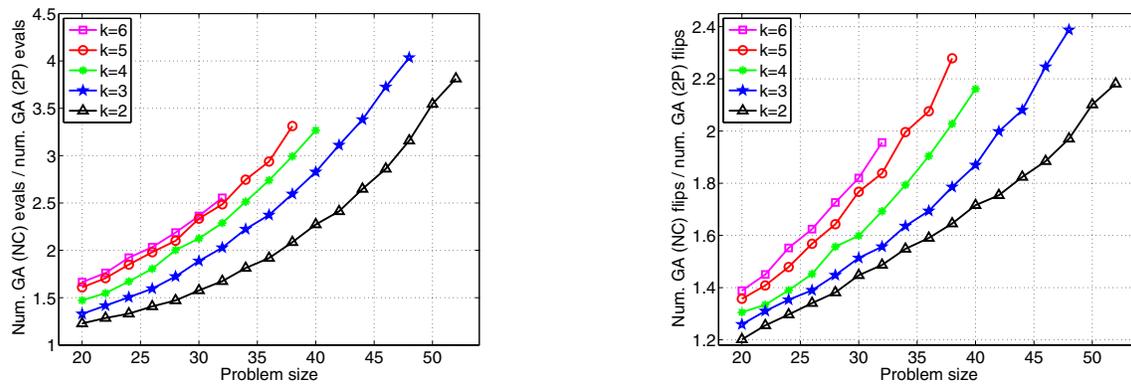


Figure 6: Comparison of GA with two-point crossover and GA with no crossover (mutation only) with respect to the number of evaluations and the number of flips.

Award and Research Board programs. 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. Some experiments were done using the hBOA software developed by M. Pelikan and D. E. Goldberg at the Univ. of Illinois at Urbana-Champaign and most experiments were performed on the Beowulf cluster maintained by ITS at the Univ. of Missouri in St. Louis.

7. REFERENCES

- [1] H. E. Aguirre and K. Tanaka. Genetic algorithms on nk-landscapes: Effects of selection, drift, mutation, and recombination. In G. R. Raidl et al., editors, *Applications of Evolutionary Computing: EvoWorkshops 2003*, pages 131–142, 2003.
- [2] L. Altenberg. NK landscapes. In T. Bäck, D. B. Fogel, and Z. Michalewicz, editors, *Handbook of Evolutionary Computation*, pages B2.7:5–10. Institute of Physics Publishing and Oxford University Press, Bristol, New York, 1997.
- [3] S. Baluja. Population-based incremental learning: A method for integrating genetic search based function optimization and competitive learning. Tech. Rep. No. CMU-CS-94-163, Carnegie Mellon University, Pittsburgh, PA, 1994.
- [4] P. A. N. Bosman and D. Thierens. Continuous iterated density estimation evolutionary algorithms within the IDEA framework. *Workshop Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2000)*, pages 197–200, 2000.
- [5] 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.
- [6] S.-S. Choi, K. Jung, and J. H. Kim. Phase transition in a random NK landscape model. pages 1241–1248, 2005.
- [7] 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.
- [8] Y. Gao and J. C. Culberson. An analysis of phase transition in NK landscapes. *Journal of Artificial Intelligence Research (JAIR)*, 17:309–332, 2002.
- [9] D. E. Goldberg. *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley, Reading, MA, 1989.
- [10] G. R. Harik. Finding multimodal solutions using restricted tournament selection. *Proceedings of the International Conference on Genetic Algorithms (ICGA-95)*, pages 24–31, 1995.
- [11] 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.
- [12] J. H. Holland. *Adaptation in natural and artificial systems*. University of Michigan Press, Ann Arbor, MI, 1975.
- [13] S. Kauffman. Adaptation on rugged fitness landscapes. In D. L. Stein, editor, *Lecture Notes in the Sciences of Complexity*, pages 527–618. Addison Wesley, 1989.
- [14] S. Kauffman. *The origins of Order: Self-organization and selection in evolution*. Oxford University Press, 1993.
- [15] P. Larrañaga and J. A. Lozano, editors. *Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation*. Kluwer, Boston, MA, 2002.
- [16] H. Mühlenbein and G. Paaß. From recombination of genes to the estimation of distributions I. Binary parameters. *Parallel Problem Solving from Nature*, pages 178–187, 1996.
- [17] M. Pelikan. *Hierarchical Bayesian optimization algorithm: Toward a new generation of evolutionary algorithms*. Springer, 2005.
- [18] 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.
- [19] 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.
- [20] 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.
- [21] M. Pelikan, H. G. Katzgraber, and S. Kobe. Finding ground states of Sherrington-Kirkpatrick spin glasses with hierarchical BOA and genetic algorithms. MEDAL Report No. 2008004, Missouri Estimation of Distribution Algorithms Laboratory, University of Missouri–St. Louis, St. Louis, MO, 2008.
- [22] 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.
- [23] A. H. Wright, R. K. Thompson, and J. Zhang. The computational complexity of N-K fitness functions. *IEEE Transactions Evolutionary Computation*, 4(4):373–379, 2000.