Extremal Optimization

Principal Investigator:

Madhav V. Marathe, CIC–3
7–8010, marathe@lanl.gov

Co-Investigator:

Allon G. Percus, CIC–3/T–CNLS
5–7435, percus@lanl.gov

Technical Category: Computer Science and Software Engineering (CSSE)

FY 2000 Funding Requested: $140K

Abstract:

Some of the most widely used heuristic tools for computationally hard problems have been inspired by equilibrium statistical mechanics. These methods — simulated annealing is an example — have been particularly successful because of their generality: given a new optimization problem about which one knows virtually nothing, one can often apply such heuristics with minimal effort, obtaining good near-optimal solutions. Equilibrium approaches, however, typically fail as one approaches *criticality* (another loan-word from statistical physics). The critical point is the transitional point that separates instances into distinct regimes of parameter space; computer scientists have observed that the instances of highest complexity in a hard problem lie close to this point.

Curiously, we have found that heuristic algorithms inspired by *non-equilibrium* physical processes do not seem to experience difficulties at criticality. A general-purpose method that we have been developing, called *extremal optimization*, shows no signs of diminished performance near the critical point, when tested on graph partitioning problems. We propose to analyze such approaches. The development of non-equilibrium optimization methods is likely, we believe, to lead to the next generation of general-purpose algorithms — intended, like simulated annealing, for broad application. We expect that this analysis will lead us to new insights into the role criticality plays in combinatorial optimization, as well as to a deeper (and more applied) understanding of computational complexity.

Scientific and Technical Impact

The main promise of research into non-equilibrium optimization methods lies in the development of general-purpose heuristics that can be applied to a wide range of computationally hard problems, particularly those where no good specialized algorithm exists. In the past decade, a very popular approach for attacking such problems has been *simulated annealing* [1], an algorithm that is easily described, easily implemented, and often leads to near-optimal solutions that could not otherwise be found without enormous conceptual and computational effort. We believe that our work on non-equilibrium methods has a good chance of providing the world of science with the "next" simulated annealing. The *extremal optimization* approach that we are developing appears, under preliminary testing, to succeed frequently where simulated annealing breaks down: at or near *critical points* [2]. These are transitional points, on the border of parameter regimes, that have been found to contain the computationally "hardest" instances [3]. Understanding the role of critical points in combinatorial optimization problems has been a goal of computer scientists for years now; we expect that an analysis of the effectiveness of non-equilibrium methods near criticality will contribute substantially to this goal. The result, ultimately, will be new insights into algorithmic performance on computationally hard problems in general.

Background

The last fifteen years have seen a considerable amount of cross-fertilization between the fields of operations research and statistical mechanics. One of the main benefits that combinatorial optimization has derived from this relationship is the development of powerful and quite general heuristic methods for finding near-optimal solutions. The canonical example is simulated annealing [1], which is inspired by the behavior of physical systems in thermal equilibrium. Simulated annealing has become, often, the method of choice for finding near-optimal solutions to *NP-hard* problems, *i.e.*, those for which no algorithm can with certainty find the *exact* optimum in polynomial time.

For some NP-hard problems, such as the traveling salesman problem, years of research have produced high-quality heuristics specifically crafted for that particular problem [4]. Where such algorithms have been developed, they often vastly outperform their more general-purpose cousins such as simulated annealing. However, this phenomenon is restricted to only a few well-studied cases. In fact, even some fairly classic combinatorial optimization problems, such as graph partitioning, still suffer from relatively poor special-purpose heuristics. And of course, the great majority of real-life optimization problems are not well-studied at all. Often, we know absolutely nothing about the general problem, except that we are able to write down an objective cost as a function of a finite number of components of the system. For such cases, a general-purpose optimization tool is desired — one that can practically be applied "off the shelf" and gives a good first approximation to an optimal solution.

Thus the appeal of a method such as simulated annealing. Like other methods based on "local search", simulated annealing works by taking an initial state of the system, and then attempting to make small, incremental changes in order to improve upon it. These changes, while taking the system sometimes to more and sometimes to less optimal states, are all governed by the laws of equilibrium statistical physics. Local search methods work satisfactorily when the system can hop easily from one local optimum to another. However, if the local optima are surrounded by high barriers in the "energy landscape", the search can get trapped at a single one of them — more likely than not, at a poor local optimum. This means that even if the procedure is run from multiple starting states, and even given very long runtimes, it does not find good near-optimal solutions.

In an empirical study, Cheeseman, Kanefsky and Taylor [3] found that this phenomenon tends to be correlated with instances near to *critical points*. In statistical mechanics, critical points are where *phase transitions* take place, singularities in the behavior of systems that are observed when an order parameter is varied. (A simple example of this is the melting point in water, where the phase transition from solid to liquid takes place.) In combinatorial problems, phase transitions represent an abrupt change in parameter space from a region of underconstrained instances, with many local optima, to a region of overconstrained instances, with very few local optima. The abruptness of this transition becomes increasingly pronounced with increasing system size. But for all finite sizes, close to the critical point where the transition occurs, there is a small region of instances "where the *really* hard problems are" [3], containing local optima that are both sparse and very inaccessible. These are the instances of high complexity, where local search methods get stuck, and in fact where exact algorithms have the highest running times as well.

Phase transitions in combinatorial optimization was first studied almost 30 years ago by Erdős and Rényi, in the context of random graph theory [5]. They found that for a wide range of properties, the probability of a random graph displaying that property exhibits a sharp increase at some critical value of the edge probability. The observations of Cheeseman *et al.* subsequently sparked the interest of the theoretical computer science and AI communities, with the result that the subject of criticality in NP-hard problems has developed into a rapidly growing theory. This has led to an improved understanding of the reasons for computational intractability. Phase transitions have been identified in practical contexts such as planning, scheduling and register allocation in compilers, and have led to renewed interest and research in the area of creating benchmark suites.

We have recently begun to realize that while equilibrium heuristics suffer from these complexity issues — as do exact algorithms — heuristics that take the system *far from equilibrium* may not. We have been developing just such a method, called *extremal optimization* [2], inspired by self-organizing processes often found in nature. Typically, these processes work by successively eliminating extremely undesirable components of sub-optimal solutions, and generating their replacements stochastically. Evolution is a prime example of how this happens in nature: the few most poorly-adapted species are selected *against*, without well-adapted species being expressly constructed — these simply come about through random mutations. Extremal optimization follows this paradigm. A "fitness" is assigned (according to some rule) to each component of the system, whether it is a vertex in a graph, a city in a traveling salesman problem, etc. The algorithm then successively selects the least "fit" component for an adaptive change, assigning to it a new random value. The optimal solution is defined as the best one seen so far in a given run. Note that no attempt is ever made at enforcing any equilibrium conditions.

The great surprise is that this simple extremal optimization algorithm (modulo a few finer details) routinely outperforms simulated annealing when applied to the problem of mincut graph bisection, for a wide choice of graphs [2]. The effect is most pronounced on instances with graph connectivities close to the percolation threshold, where a phase transition takes place. (Similar results have been obtained on random distance-matrix instances of the traveling salesman problem, though in that case it is far less clear where criticality is to be expected.) We consider that this merits further analysis. If the extremal optimization approach can be generalized to a wide class of optimization problems, and if it maintains its success at optimizing near phase transitions, this is a great improvement over current general-purpose methods. And even if its success is limited to a few problems, it already complicates substantially the relationship between criticality and optimization. Why is a non-equilibrium optimization method working so well here? And what does this tell us about our current notion of computational complexity? A host of new questions are raised.

Although the field of science that we propose here is not well covered in the literature, two useful background sources might be: the special issue [6] of *Artificial Intelligence*, for discussions of phase transitions and complexity, and the book [7] by Per Bak, for discussion and analysis of self-organization models in nature.

As far as applications of our proposal to ongoing LANL activities are concerned, our research is of great relevance to the Laboratory's core mission and its Stockpile Stewardship program goals, specifically in high-performance computing. Graph partitioning problems, on which extremal optimization has been found highly effective in early trials, are routinely applied in VLSI design (partitioning of gates between integrated circuits) as well as in load balancing. Progress of this sort is necessary for improved parallel computing capabilities [8] — for better use of current resources, as well as development of more rapid computing architecture.

Finally, to some extent the present proposal ties in with LDRD/DR grant #99517, *Novel Fundamentals in Strategic Computing* (PI: Olaf Lubeck). It should be stressed, however, that the two are complementary in that the DR project is primarily directed towards algorithmic theory, whereas ours concentrates more on analysis of complex systems with a view to improving heuristic performance.

Research Objectives and Goals

The main goals of our research are as follows:

- Introduce heuristic optimization methods inspired by *non-equilibrium* statistical physics.
- Develop our *extremal optimization* approach as the next generation of general-purpose methods, superseding simulated annealing, and widely applicable to new problems.
- Understand the behavior of non-equilibrium methods near critical points, and why they do not necessarily suffer there from the same complexity issues as so exact algorithms and equilibrium-based heuristics.
- Consequently, gain deeper insight into computational complexity, particularly as it affects average-case performance.

R&D Approach

We plan on elaborating on our extremal optimization testbed, so as to include many more combinatorial optimization problems, such as graph coloring, satisfiability, etc. The aim is to perform a comprehensive test of quality issues. This will involves both average performance (how it does compared to other algorithms, averaged over an ensemble of instances), and best performance (whether there are many instances where it finds near-optimal solutions that other methods simply never find, even given very long runtimes). The goal, particularly for the latter question, is to determine just how general its strong performance near criticality is — and whether this strength translates into an ability to outperform other algorithms at equal runtimes... or simply an ability to find better solutions, period.

We will also consider the issue of *ergodicity*. We have noticed, empirically, that extremal optimization responds favorably to a relaxation that involves selecting out components with poorly-ranked fitness from a probability distribution, rather than simply selecting out the least fit component with probability 1. Interestingly, the best performance seems to occur when this distribution is just barely broad enough that the system samples a substantial fraction of possible configurations, *i.e.*, just where ergodicity sets in. The link between ideal performance and the "edge of ergodicity" is a fascinating one, and will be the subject of both an empirical and a theoretical study.

The question of *why* extremal optimization works so well near critical regions, finally, will be extensively studied. We are privileged in that Gabriel Istrate, whose PhD thesis contains groundbreaking theoretical work on critical behavior in satisfiability problems, will be joining us as a postdoc this summer. Dr. Istrate has obtained rigorous models of critical phenomena [9] that he is interested in applying to a wide range of combinatorial optimization problems. These models will help us understand the algorithmic performance of non-equilibrium methods, and could translate the empirical success of such methods into a better understanding of complexity in general.

We are fortunate in that the other members of the present research team have strong credentials in the field of optimization as well. Madhav Marathe has made extensive contributions to the areas of computational complexity theory, performing pioneering work on the notion of approximation algorithms for multi-criteria optimization problems. Allon Percus has been working for some years on combinatorial optimization strategies motivated by statistical physics; in his research on the stochastic traveling salesman problem, he produced the most precise estimate to date for large N asymptotic tour lengths. The concept of extremal optimization has come as the result of collaborative work with Stefan Boettcher, who has a background in theoretical physics and has published extensively on far-from-equilibrium dynamics and self-organized criticality in evolutionary models.

Expected Scientific and Technical Results

We plan on obtaining the following results:

- *A theory of extremal optimization.* This includes completing the development of our algorithm, producing a highly efficient and general-purpose optimization method based on non-equilibrium dynamics. The algorithm is expected to outperform current leading general-purpose methods on many combinatorial optimization problems. By general-purpose, we mean that the algorithm will be applicable with a minimum of effort, and particularly useful on problems for which no finely-tuned heuristics yet exist.

- *A testbed documenting the performance of our algorithm on well-known problems.* This is necessary for establishing a consensus on state-of-the-art, something we have found to be lacking for many common optimization problems.

- *A theoretical framework for understanding the performance of non-equilibrium heuristics near criticality.* This involves an understanding of both the physical inspiration for these heuristics, and the algorithmic theory governing them. We plan on developing models that improve upon current notions of computational complexity, providing a more adequate explanation of the algorithmic performance (worst-case and average-case) of non-equilibrium methods.

Funding Breakout

Funding at or near the FY2000 level is requested for 3 years. The breakdown for the first year is expected to be as follows:

$65K = 1/2 TSM or postdoc (Allon Percus)
$65K = postdoc (Gabriel Istrate)
$10K = visitor funds (Stefan Boettcher) and miscellaneous

Key LANL staff participating in the project are expected to be:

Gabriel Istrate (CIC–3, to arrive in Summer '99)
Madhav V. Marathe (CIC–3)
Allon G. Percus (CIC–3/T–CNLS)

Dr. Stefan Boettcher of Emory University, Atlanta, GA, will also be one of the main collaborators on this project.

Possible Specialist Reviewers

Internal reviewers:

- Eli Ben-Naim, T–13
  7–9471, ebn@lanl.gov
- Olaf M. Lubeck, CIC–19
  7–6017, oml@lanl.gov

External reviewers:

- Per Bak, Niels Bohr Institute, Copenhagen, Denmark
  E-mail: bak@nbi.dk
- Paolo Sibani, Physics Department, Odense University, Denmark
  E-mail: paolo@planck.fys.ou.dk

References

[1] S. Kirkpatrick, C. D. Gelatt and M. P. Vecchi, "Optimization by simulated annealing", *Science* **220** (1983) 671–680; V. Černy, "Thermodynamical approach to the traveling salesman problem: an efficient simulation algorithm", *J. Optimization Theory Appl.* **45** (1985) 41–51.

[2] S. Boettcher and A. G. Percus, "Extremal optimization: methods derived from co-evolution", *Proc. of GECCO–99*, to appear. Available at `http://userwww.service.emory.edu/~sboettc`.

[3] P. Cheeseman, B. Kanefsky and W. M. Taylor, "Where the *really* hard problems are", *Proc. of IJCAI–91* 331–337.

[4] The best example for the Traveling Salesman Problem is the Lin-Kernighan heuristic. See: S. Lin and B. Kernighan, "An effective heuristic algorithm for the traveling salesman problem", *Operations Research* **21** (1973) 498–516.

[5] P. Erdős and A. Rényi, "On the evolution of random graphs", *Publication of the Mathematical Institute of the Hungarian Academy of Science* **5** (1960) 17–61.

[6] "Frontiers in problem solving: Phase transitions and complexity", Special issue of *Artificial Intelligence* **81**:1–2 (1996).

[7] P. Bak, *How Nature Works* (Springer, New York, 1996).

[8] For an overview, see *Los Alamos Science* **22** (1994), available at `http://lib-www.lanl.gov/pubs/la-sci.htm`.

[9] G. Istrate, "The phase transition in random horn satisfiability and its algorithmic implications", *Proc. of SODA '99* S925–S926; "Critical behavior in the satisfiability of random k-Horn formulae", *The EP-SRC/LMS Workshop on Phase Transitions in Combinatorial Problems*, Liverpool, England, January 1999.

MADHAV V. MARATHE

Research Interests:
- Modeling, Simulation and Analysis of Large Scale Systems, High Performance Computing, Design and Analysis (theoretical and experimental) of Algorithms, Complexity Theory, Combinatorial Optimization, Data Mining, Mobile Computing

Education:
- University at Albany, SUNY: Ph.D. August 1994, Computer Science
- Indian Institute of Technology, Madras, India: B.Tech. August 1989, Computer Science

Work Experience:
- June 1996 - Present: Technical Staff Member, LANL
- July 1998 - Present: Adjunct Assistant Professor, Department of Computer Science, University of New Mexico
- August 1994 - June 1996: Postdoctoral Research Associate, LANL

Reviewer:
- Associate Editor: Journal of Computing and Information
- Reviewer: SIAM J. on Computing, Theoretical Computer Science, Operations Research, International Journal on Foundations of Computer Science, INFORMS J. Computing, Networks, Information Processing Letters, ACM-SIAM Symposium on Discrete Algorithms, FST & TCS, Symposium on Parallel and Distributed Computing

Selected Publications:

1) "Loop Transformations for Performance and Message Latency Hiding in Parallel Object-Oriented Frameworks," F. Bassetti, K. Davis, M.V. Marathe and D. Quinlan, in: Proc. International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 98), Las Vegas Nevada, July 1998.
2) "Theory of Periodically Specified Problems: Complexity and Approximability," M.V. Marathe, H.B. Hunt III, D.J. Rosenkrantz and R.E. Stearns, in: Proc. 13th IEEE Conference on Computational Complexity, Buffalo, NY, June 1998.
3) "Map Labeling Problems," S. Doddi, M.V. Marathe, A. Mirzaian, B. Moret, and B. Zhu in: Proc. 8th ACM-SIAM Symposium on Discrete Algorithms (SODA), San Francisco, CA, pp. 148-157, January 1997.
4) "Complexity of Hierarchically and 1-Dimensional Periodically Specified Problems I: Hardness Results," M.V. Marathe, H.B. Hunt III, R.E. Stearns and V. Radhakrishnan, in: AMS-DIMACS Volume Series on Discrete Mathematics and Theoretical Computer Science: Workshop on Satisfiability Problem: Theory and Application, Vol 35, pp. 225-259, November 1996.
5) "Approximation Algorithms for the Minimum Satisfiability Problem," M.V. Marathe and S.S. Ravi, in: Information Processing Letters, (IPL), Vol. 58, No. 1, pp. 23-29, April 1996.
6) "Approximation Schemes for PSPACE-Complete Problems for Succinct Specifications," M.V. Marathe, H.B. Hunt III, R.E. Stearns and V. Radhakrishnan, in: Proc. 26th Annual ACM Symposium on the Theory of Computing (STOC), pp. 468-478, May 1994.

ALLON PERCUS

Contact Information: Tel: 5-7435, E-mail percus@lanl.gov
WWW: http://www.lanl.gov/home/percus

Present Position: *Postdoctoral Research Associate* at Los Alamos National Laboratory, jointly affiliated with the Computer Research and Applications group (CIC-3) and the Center for Nonlinear Studies (CNLS)

Education:
- *Université Paris-Sud*: PhD in statistical physics
  September 1997
- *Ecole Normale Supérieure*, Paris: DEA (MS equivalent) in theoretical physics
  September 1994
- *Harvard University*: BA cum laude (physics)
  June 1992

Selected Publications:
- S. Boettcher and A.G. Percus, "Extremal optimization: methods derived from co-evolution", *Proceedings of GECCO–99*, to appear.
- A.G. Percus and O.C. Martin, "The stochastic traveling salesman problem: Finite size scaling and the cavity prediction", *Journal of Statistical Physics* 94 (1999), to appear.
- A.G. Percus and D.C. Torney, "Greedy algorithms for optimized DNA sequencing", *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '99)*, S955–S956.
- A.G. Percus and O.C. Martin, "Finite size scaling universalities of kth-nearest neighbor distances on closed manifolds", *Advances in Applied Mathematics* 21 (1998) 424–436.
- N.J. Cerf, J. Boutet de Monvel, O. Bohigas, O.C. Martin and A.G. Percus, "The random link approximation for the Euclidean traveling salesman problem", *Journal de Physique I* 7 (1997) 117–136.
- A.G. Percus and O.C. Martin, "Finite size and dimensional dependence of the Euclidean traveling salesman problem", *Physical Review Letters* 76 (1996) 1188–1191.
- H.M. Lacker and A. Percus, "How do ovarian follicles interact? A many-body problem with unusual symmetry and symmetry-breaking properties", *Journal of Statistical Physics* 63 (1991) 1133–1161.

LANL Activities:
- Co-organizer of *Frontiers of Combinatorics* workshop, Summer 1998
- Member of *Reseach Library Advisory Board*

Invited Presentations:
- "Greedy algorithms for optimized DNA sequencing", invited talk at National Center of Biotechnology Information, National Library of Medicine, Bethesda, MD, 15 January 1999.
- "The stochastic traveling salesman problem and the random link approximation", Physics Dept. Colloquium, Emory University, Atlanta, 13 November 1998.

CURRICULUM VITAE

Personal Data:

Address:

Stefan Boettcher
Physics Department
Emory University
Atlanta, GA 30322-2430

E-mail: sboettc@emory.edu

Education:

Ph.D. in Physics, Washington University, St. Louis (1993)
"Diplom" degree in Physics, Kiel University, Kiel, F. R. Germany (1990)

Employment:

Since 1998: Faculty, Emory University
1996-1998: Postdoctoral Fellow, Los Alamos National Laboratory
1995-1996: Postdoctoral Fellow, The University of Oklahoma
1993-1995: Postdoctoral Fellow, Brookhaven National Laboratory

Selected Publications:

I have published about 30 peer-reviewed papers on statistical and mathematical physics. Among those, the most relevant to this project are:

1. "Nature's Way of Optimizing,"
S. Boettcher and A. G. Percus,
(submitted to *Artificial Intelligence*, cond-mat/9901351),

2. "Aging in a Model of Self-Organized Criticality,"
S. Boettcher and M. Paczuski,
Physical Review Letters **79**, 889 (1997).

3. "Punctuated Equilibria and Self-Organized Criticality in Evolution,"
P. Bak and S. Boettcher,
Physica D **107**, 143 (1997).

4. "Ultrametricity and Memory in a Solvable Model of Self-Organized Criticality,"
S. Boettcher and M. Paczuski,
Physical Review E **54**, 1082 (1996).

5. "Exact Results for Spatio-Temporal Correlations in a Self-Organized Critical Model of Punctuated Equilibrium,"
S. Boettcher and M. Paczuski,
Physical Review Letters **76**, 348 (1996).