Report on

Workshop on Computational Complexity and Statistical Mechanics

September 3-6, 2001, Santa Fe, New Mexico

Organizers: Gabriel Istrate, Allon Percus, Chris Moore

Gabriel Istrate

Los Alamos National Laboratories
Second Line Goes Here
Los Alamos, NM 2
E-mail: `@microsoft.com`

John Franco

ECECS, Computer Science
University of Cincinnati
Cincinnati, OH 45221-0030
E-mail: `franco@gauss.ececs.uc.edu`

September 12, 2001

This is a summary of activity at the first Workshop on Computational Complexity and Statistical Mechanics, held September 3-6 in Santa Fe, New Mexico. The workshop was organized by Gabriel Istrate and Allon Percus of Los Alamos National Laboratory and Chris Moore of the University of New Mexico and Santa Fe Institute and was attended by approximately 45 participants including university faculty, students, and researchers from industrial and government laboratories. Most attending are US residents but a few came from Europe.

The dominant theme of the workshop was thresholds or phase transitions associated with statistical properties of random instances of NP-complete problems. Secondary themes consider statistical properties of other random structures and the use of non-rigorous methods to gain intuition. Work on these topics spans several disciplines including Artificial Intelligence, Theoretical Computer Science, Physics and Mathematics, and seeks to reveal connections between physical phenomena and the nature of computation.

Since many computer scientists have not come to fully appreciate the use and meaning of stochastic processes and statistics in algorithmic complexity research we first provide a simple background related to thresholds. Although a true model for a typical physical system can involve so many components or sites as to become intractable, it is sometimes the case that an accurate mathematical representation can be obtained from an approximation of its average behavior as if it were an infinite system. Consider, for example, the percolation of a fluid in a large network of sites and pipes. Given an initial concentration of fluid at a particular site, the question is to determine the percentage of sites eventually receiving fluid. For infinitely large networks, regardless of detailed geometry, if the average number of neighbors for a site is below a critical value, then the probability that a site which is an arbitrarily large distance from the source receives fluid tends to 0, but if the average number of neighbors is greater than a critical value the probability of receiving fluid tends to 1. This transition behavior is characterized by rapidly increasing lengths of correlation paths between parts of the system as the transition is approached (that is, the behavior of correlation paths is not analytic for infinite systems at the transition point). Researchers are finding that this behavior is analogous to that of search algorithms for computer science problems such as the Satisfiability problem.

The Satisfiability problem (SAT) is the question of determining whether there exists an assignment of Boolean values to variables of a propositional expression in Conjunctive Normal Form (conjunction of disjunctions of variables or negations of variables, also called clauses) which causes the expression to evaluate to *true*. The problem is NP-complete. It is referred to as k-SAT if all clauses have k atomic elements. Random instances of k-SAT contain m clauses constructed uniformly and independently from n variables. It has been observed that search algorithms applied to random instances of k-SAT typically have poor performance around a critical threshold identified by $m/n = c(k)$, where $c(k)$ is not yet known precisely but is in the intersection of $\Omega(2^k/k)$ and $O(2^k)$. They are more efficient, usually, if m/n is far away from the critical threshold. The reason seems to be that with relatively few constraints (clauses) there are many opportunities to find a satisfying assignment because so many exist and with relatively many constraints there are many opportunities to combine constraints for a refutation. At the threshold (or transition) there are few satisfying assignments and, due to sparseness properties of random formulas, it is unlikely that the existing constraints can be combined in an efficient way to prove no satisfying assignment exists. In other words, correlation paths between clauses are very long resulting in long "backbones" of inferred values to variables which accompany solutions[1].

Curiously, the backbones of k-SAT near the threshold seem analogous to the well-studied backbones of molecular states that are observed during phase transitions in physics. As, say temperature, is reduced, a backbone of molecular magnetic "spins" develops during a phase transition. More than one particular backbone may develop, but once one forms, the energy needed to "backtrack" to another rises prohibitively. This is analogous to computational work needed to search for another assignment inferring a similar number of values and

[1]This is probably due to the fact that the smallest unsatisfiable subformulas are getting big near the threshold.

satisfying the same number of clauses (that is, energy state). If the analogy can be made concrete in some understandable way, we may be able to improve our understanding of the nature of "hard" problems and what it takes to make them easier. Thus, a number of people are studying k-SAT transitions, some with assistance from phase transition results. Among other things, they are looking at the shapes of the transition curves hoping to determine the relationship between sharp or first-order transitions with hard problems and the relationship of coarse or second-order transitions with easy problems. Some are looking for definitive ways that transition behavior implies algorithmic behavior of any kind.

SAT and a generalization known as Constraint Satisfaction Problems are important classes studied in this area and several talks at the workshop addressed these. Talks by Achlioptas and Kirousis presented past results on bounding the transition of 3-SAT from above and below and highlighted techniques for obtaining those bounds, including their limitations. Lower bounds have relied on analyzing Markovian migration of clauses through widths 3, 2 and 1 but such results have topped out at critical $m/n = 3.26$ which is far below the empirically obtained threshold of $m/n = 4.25$. Upper bounds have relied on first moment analyses of subsets of satisfying assignments with reduced variance. Kautz is interested in generators of satisfiable formulas with distributions ranging from ones that are everywhere tractable to ones that have a sharp hardness threshold. These are considered important for testing "one-sided" or incomplete solvers. Daudé presented transition bounds for k-XORSAT, instances of which can be solved in polynomial time by gaussian elimination. The aim is to help gather evidence for comparing transitions on "easy" classes with transitions on "hard" classes. Demopoulis presented experimental data with the aim of identifying easy-hard transitions for specific SAT solvers. Such results might contribute to improving SAT solvers of the future. Molloy discussed models for random constraint satisfaction problems. He is able to characterize those in which satisfiability transitions are sharp with respect to number of constraints.

Phase transitions on problems other than SAT have been studied and some results were reported at the workshop. Chayes considered the Number Partitioning Problem, which is closely related to the Subset Sum and other NP-complete Problems. It was shown that random instances of this problem, taken from a suitable distribution, exhibit a first-order phase transition. This is interesting in light of the fact that some polynomial time solvable problems such as 2-SAT exhibit a second-order phase transition. A similar result was reported by Mertens using statistical methods developed for the problem of identifying the minimum in a list of random numbers. However, Mertens also offered a calculation of the probability distributions of optimal and sub-optimal costs. Culberson experimentally considered the nature of "frozen" edges for Graph Coloring. Frozen edges in Graph Coloring are the counterpart of backbones in SAT. If random graphs are constructed by adding edges one at a time, the rate at which edges become frozen shows a rapid jump, indicating a sharp phase transition of the kind typically observed for NP-complete problems. Boettcher presented a heuristic for finding good solutions to hard physical and combinatorial optimization problems which is designed to be particularly successful near phase transitions. Duxbury considered phase transitions by constraint counting and presented results for a physical and combinatorial problem, namely Rigidity Percolation and Minimum Vertex Cover. Borgs considered phase transitions for random H-colorings of rectangular subsets of the hypercubic lattice.

Gibbs sampling has been applied successfully in some cases to combinatorial optimization problems and several speakers reported results with this technique. Braun investigated the behavior of typical solutions to optimization problems, whose parameters are set randomly, with respect to the empirical Gibbs measure of one training instance on a second test instance. His focus was the case where "typical" solutions for different sampled instances start to look very different below a certain temperature and their performance on the test instance decreases. He conjectured a connection exists between the robustness of a solution sampled from the Gibbs distribution at finite temperature and the hardness of approximating the solution to the same relative error as the "typical" instance. Sinclair presented a Markov chain Monte Carlo algorithm that samples perfect matchings in an arbitrary bipartite graph from the Gibbs distribution with arbitrary non-negative weights on the edges. He obtained a polynomial time algorithm that computes the partition function to any desired accuracy.

The following are highlights of the remaining presentations. Aldous considered bipartite matching problems with edge weights taken independently from an exponential distribution. The mean of the optimal solution had been conjectured to be $\zeta(2) = \pi^2/6$. Aldous verified this rigorously by cleverly identifying the limit in terms of a matching problem on a limit infinite tree. Selman discussed the "heavy-tailed" phenomena in combinatorial search. If distribution tails decay "slowly," search mechanisms may get "stuck" in fruitless portions of the search space. Restarting frequently can mitigate this problem. Hogg and van Dam discussed algorithms for quantum computing. Hogg explained how certain regularities related to phase transitions apply to quantum search algorithms and van Dam claimed results which indicate quantum adiabatic computing will probably not be useful for solving NP-complete problems in polynomial time. Moore showed that almost all graphs with degree 4.03 or less, including 4-regular graphs, are 3-colorable. Yukich found a law of large numbers and a central limit theorem for the problem of sequentially packing unit balls in a large cube (the existence of such had been predicted using non-rigorous techniques). Parkes presented some insights on interactions of phase transitions, constraint relaxation, and distributed and parallel search methods. Marathe described research aimed at proposing a "predictive complexity theory" which simultaneously characterizes complexity and approximability of combinatorial problems when instances are succinctly specified in a certain way. Reidys discussed combinatorial landscapes. Impagliazzo analyzed hill-climbing algorithms for planted bisection problems and compared the results to Metropolis algorithms.

Where do we go from here? An illusive goal has been to show that statistical properties imply some meaningful algorithmic complexity properties. Hoping for $P \neq NP$ may be asking too much. Some wonder whether we will be able to establish direct connections between statistical properties and complexity bounds of some specific problems. However, links between statistical properties and bounds for certain classes of algorithms have already been established, for example the link between sparseness and resolution in the case of unsatisfaible 3-SAT formulas. Statistical methods, such as Gibbs sampling, may lead to useful alternative algorithms for particular optimization problems such as those studied by Braun. Non-rigorous methods can reveal amazing insights. But how reliable are these insights? At the moment, we are gathering evidence by attempting considerably many comparisons between rigorous and non-rigorous results on the same problem. It will probably be some time before we know for sure.