

Critical Slowing Down in Polynomial Time Algorithms

A. Alan Middleton

Department of Physics, Syracuse University, Syracuse, New York 13244

(Received 10 April 2001; published 17 December 2001)

Combinatorial optimization algorithms that compute exact ground states for disordered magnets are seen to exhibit critical slowing down at zero temperature phase transitions. Using the physical features of the models, such as vanishing stiffness on one side of the transition and the ground state degeneracy, the number of operations needed in the push-relabel algorithm for the random field Ising model and in the algorithm for the 2D spin glass is estimated. These results strengthen the connections between algorithms and the physical picture and may be used to improve the speed of computations.

DOI: 10.1103/PhysRevLett.88.017202

PACS numbers: 75.10.Nr, 02.70.-c, 05.50.+q, 75.40.Mg

There exists a fundamental connection between the concepts of statistical physics and the algorithms used to simulate condensed matter systems: the mathematics of graphs. This connection is used in analytical computations of high temperature series and is used in high precision numerical calculations of percolation [1], using sophisticated connectivity algorithms from computer science. The Fortuin-Kasteleyn cluster representation [2] for magnets leads to the Swendsen-Wang algorithm [3]. By using non-local dynamics, these algorithms greatly reduce the simulation time near a phase transition. The study of disordered systems, such as spin glasses, pinned vortices in superconductors, and random field magnets, introduces graphs with nonuniform edges [4,5]. Early in the study of disordered systems, it was realized that the study of such graphs is directly related to issues of computational complexity. In some cases [6], computing the ground state of a disordered material is intractable, as the relevant graph problems are in the class of NP-hard problems [7]. Quite interestingly, some computationally intractable problems have been found to have phase transitions. For example, given an ensemble of logical expressions characterized by the number of Boolean variables N and clauses M , the fraction that are satisfiable can exhibit finite size scaling about a critical value of M/N [8].

In this Letter, I present results on the behavior of ground state algorithms near phase transitions in two models, the random field Ising magnet (RFIM) and the 2D spin glass (2DSG), which are always solvable in a time polynomial in the size of the graphs. The phase transitions lead to singularities in the mean solution time, however. For the RFIM, a connection is made between this critical slowing down, the correlation length, and the degeneracy of the ground state. Numerical evidence suggests that the dynamic critical exponent is $z \approx 1$ (RFIM); scaling arguments suggest $z \geq 1$. The algorithm for the 2DSG is difficult to analyze, but the observed uniform time per spin in the ferromagnetic (FM) phase is argued to be natural. These results suggest avenues both for improving the algorithms for disordered systems and confirm and provide insights into the physical picture of disordered systems.

The dynamics at low temperature T is exceedingly slow in disordered magnets, as seen in experiment and in local move Monte Carlo simulations [4]. The glassy dynamical behavior is due to the complex structure of the energy functional; free energy barriers at a length scale ℓ scale as ℓ^ψ , so that the equilibration times for a volume of size ℓ^d are expected to scale as $\sim \exp(\ell^\psi/T)$. However, for some random magnets there are algorithms to find an *exact* ground state that take time polynomial in ℓ^d . Ground states are found using “nonphysical” configurations or moves: A local search or simulated annealing that uses only physical configurations and local moves is hindered by the large energy barriers.

For local Monte Carlo moves in *uniform* magnets at finite T , the correlation time at continuous transitions scales as L^z , with $z_{\text{loc}} \geq \gamma/\nu$ [9]. Nonlocal cluster moves [3] can reduce the critical exponent z to $z_{\text{cl}} \geq \alpha/\nu$, with α and ν the heat capacity and correlation length exponents, respectively. As the ground state algorithms for disordered magnets were not *designed* to utilize clusters related to a critical point, it is less obvious how many operations are required. Polynomial bounds on the *worst case* behavior of these algorithms exist: For RFIM graphs with N spins the highest level version of the push-relabel (PR) algorithm uses $O(N^{5/2})$ time. Algorithms for general matching [10], applied here to the 2DSG, use from $O(N^{3/2})$ to $O(N^3)$ operations, up to logarithmic corrections. For typical disorder realizations in finite dimension d , these algorithms are faster, with the average running time for many problems scaling roughly as N^q , with q often in the range 1.1–1.3 [5,11]. For systems with a single $T = 0$ fixed point, such as the elastic medium in a random potential [11], this scaling would not be expected to vary with parameters.

The $T = 0$ RFIM has a phase transition that has been extensively studied [12] using a mapping [13] of the ground state to the optimization problem max-flow [14]. The Gaussian RFIM has Hamiltonian $H_R = -J \sum_{\langle ij \rangle} s_i s_j - \sum_i h_i s_i$, where the spins on the $N = L^d$ lattice sites i are $s_i = \pm 1$, the ferromagnetic couplings J are positive, and the random fields h_i are Gaussian with mean 0 and variance $\Delta^2 J^2$. In the paramagnetic phase

($\Delta > \Delta_c \approx 2.27$), the spins are pinned by the external fields h_i , and the net magnetization $m = N^{-1} \sum s_i$ vanishes as $L \rightarrow 0$. In the ferromagnetic phase ($\Delta < \Delta_c$), the ferromagnetic coupling J dominates and $|m|$ has a nonzero limit, with the sign of m determined roughly by the total of $\sum h_i$. The time to find the ground state in a 3D cubic lattice with periodic boundaries is plotted as a function of Δ and L in Fig. 1. Plots of the number of primitive operations executed are nearly identical in form. Near the phase transition, the time to find the ground state shows a critical slowing down. This notable effect reflects the deep connection between the dynamics of the ground-state algorithm and the physics of the RFIM.

The application of max-flow algorithms to the RFIM is well established [5], but, to connect scaling in the RFIM and algorithm timing, it is useful to review the algorithm. The network flow algorithm used to solve the RFIM, because of its speed, is the PR algorithm of Tarjan and Goldberg [15]. The algorithm described here uses a modification that removes the source and sink nodes [16], reducing memory usage and also clarifying the physical connection.

The modified PR algorithm starts by assigning variables x_i to each site and r_{ij} to each bond, with $x_i = h_i$ and $r_{ij} = J$, initially. Reflecting the motivation of a fluid flow problem through a pipe network, the x_i are “excesses” and the r_{ij} are the “residual capacities.” The *flow problem* is to maximize the total flow through the network between the sources ($x_i \geq 0$) and sinks ($x_i < 0$). Solving this problem can be shown to be equivalent to finding the ground state of the RFIM [5]. An auxiliary height variable u_i is initially assigned to each node via a global update step. In this step, the value of u_i at each site in the set $\mathcal{T} = \{j | x_j < 0\}$ of

negative excess sites is set to zero. Sites with $x_i \geq 0$ have u_i set to the length of the shortest path, via edges with positive capacity, from i to \mathcal{T} ; if there is no such path to a site i , $u_i = N$ and i is disconnected from \mathcal{T} .

The ground state is found by rearranging the excesses x_i , via “push” operations, and updating the heights u_i , via “relabel” operations. A push operation moves excess “fluid” from a site i to a lower height neighbor j , if possible: Whenever $x_i > 0$, $r_{ij} > 0$ and $u_j = u_i - 1$. In a push, the excesses and residual capacities are modified by the amount $\delta = \min(x_i, r_{ij})$ of fluid pushed through bond ij : $x_i \rightarrow x_i - \delta$, $x_j \rightarrow x_j + \delta$, $r_{ij} \rightarrow r_{ij} - \delta$, and $r_{ji} \rightarrow r_{ji} + \delta$. Push operations tend to move the positive excess towards sites in \mathcal{T} . When $x_i > 0$ and no push is possible, the site is relabeled (raised): $u_i \rightarrow 1 + \max_{\{j | r_{ij} > 0\}} u_j$. In addition, if a set of highest sites \mathcal{U} becomes isolated, with $u_i > u_j + 1$, for all $i \in \mathcal{U}$ and all $j \notin \mathcal{U}$, the height u_i for all $i \in \mathcal{U}$ is increased to its maximum value, N , as these sites will always be isolated from \mathcal{T} . When no more pushes or relabels are possible, a final global update determines the RFIM ground state: Those sites which are path connected by bonds with $r_{ij} > 0$ to \mathcal{T} have $s_i = -1$, while the sites which are disconnected from \mathcal{T} have $s_i = 1$. A proof of the correctness of the PR flow algorithm can be found in standard textbooks [14]. Periodic global updates are often crucial to the practical speed of the algorithm. The highest site heuristic is used here, which applies pushes and relabels where u_i is maximal and $x_i > 0$.

The PR algorithm is intuitively appealing: When $J \gg \Delta$, the excess can be rearranged freely, unconstrained by the r_{ij} , so that the positive and negative excesses cancel as much as possible. The final excess values have the sign of the original total excess. This corresponds to the ferromagnetic bonds being strong enough to favor complete alignments of the spins, with the spin direction given by the sign of the total h_i . If the initial capacities are not large enough (weaker J), the rearrangement of the excess is limited and regions align independently, according to the local field. The number of steps needed to move excess across the diameter of a region via push operations is bounded below by the linear size of the region.

I now argue that these correspondences can be used to bound the running time of the algorithm, using the physical properties of the RFIM, particularly the behavior of the ground state degeneracy in the thermodynamic limit. Recent work [12] has given strong numerical evidence of insensitivity of the interior spins in the ground state to boundary conditions, when $\Delta > \Delta_c$. (This is in contrast with the scenario for, e.g., mean field [4] or highly disordered spin glasses in $d > 8$ [17], where the entire solution is sensitive to “surface” changes.) The ground state solution is then determined by the h_i within a volume typically of the size of the correlation volume ξ^d . However, in the FM phase, the interior *is* sensitive to the boundaries and a finite fraction of the spins flip infinitely often as the sample

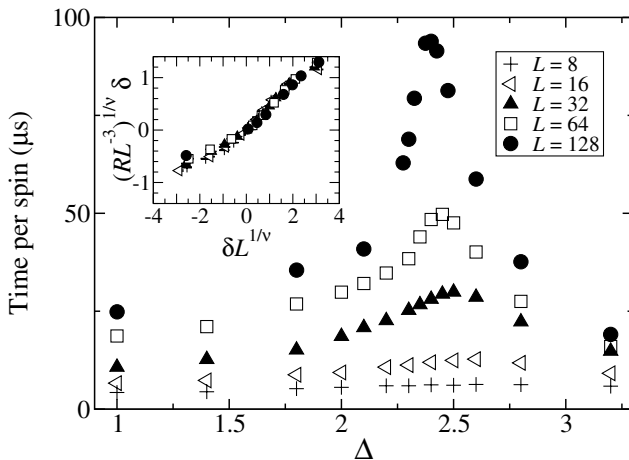


FIG. 1. The CPU time needed to find the ground state in the 3D RFIM, for a 766 MHz PIII. The inset plots $(RL^{-3})^{1/\nu} \delta$, the scaled number of relabel operations per site, with $\delta = (\Delta - \Delta_c)/\Delta_c$, vs the finite size scaling variable $\delta L^{1/\nu}$. The values $\Delta_c = 2.27$ and $\nu = 1.37$ are not fit parameters, but are derived independently. Statistical error bars are about 1/5 of the symbol size in both plots.

volume is increased, as the net magnetization is given by a coarse-grained global sum of the h_i .

First, consider the case $\Delta \gg 2d$, where, in the ground state, almost all spins satisfy $s_i = \text{sgn}(h_i)$. In the algorithm, as the total capacity of bonds leaving the site is only $2dJ$, where $h_i > 2dJ$, the x_i is always positive, so that push operations can reduce h_i only by $2dJ$. Where $h_i < -2dJ$, x_i is always negative. There will be some small rearrangement, but only locally, and the sign of the excess will change only at a very few sites during the execution of the algorithm, which terminates after a number of operations $\propto N$ (up to logarithmic corrections [18]).

A similar scenario holds, but at scales ξ , for $\Delta > \Delta_c$. The algorithm establishes the boundaries of correlation volumes by rearrangement of excess over distances of scale ξ . Further rearrangement is blocked by the effective decrease in residual capacity with scale (as the stiffness, corresponding to the scale dependent J , decreases rapidly on scales greater than ξ [12]). The question to be answered, then, is how long the algorithm takes, per spin, to redistribute excess on scale ξ . The number of push and relabel operations in a volume ξ^d is bounded below by ξ^{d+1} . The relative heights must differ by at least ξ , for excess to be pushed a distance ξ , so that at least $O(\xi^{d+1})$ relabels R must be performed for each correlation volume. This gives the estimate $R \sim L^d \xi$. This scaling is consistent with numerical results, up to logarithmic corrections, for $d = 1, 3$ [19]. The inset of Fig. 1 shows a scaling plot for R in three dimensions, for example, with $(RL^{-3})^{1/\nu} \delta$ plotted as a function of the finite size scaling variable $\delta L^{1/\nu}$, $\delta = (\Delta - \Delta_c)/\Delta_c$, with the values $\nu = 1.37$ and $\Delta_c = 2.27$ fixed parameters, determined independently [12]. A fit in $d = 1$, with $\nu = 2$ and $\Delta_c = 0$ also fixed by known results, describes the data well for $L \leq 5 \times 10^6$, with $R \sim L \xi \ln(L/\xi)$ (without the global relabeling heuristic).

One limit where the asymptotic time for the algorithm can be described more precisely is when $\Delta \ll [L \ln(L)]^{-1/2}$. Here, the capacities do not limit the rearrangement of excess: The final state is either $x_i \geq 0$ everywhere or $x_i \leq 0$ everywhere, corresponding to a uniform $s_i = \pm 1$ state, according to whether $\sum h_i \geq 0$ or $\sum h_i \leq 0$, respectively. The dynamics of the PR algorithm is set by the fluctuations in the total h_i in regions at each length scale. At any time scale of the computation, the positive excess will be pushed towards the nearest negative excess region, with the distribution of excess negative or positive over a particular volume with ℓ^d sites, as the rearrangement of excess will have been completed over shorter lengths at an earlier time scale. Generally, but especially given sufficiently frequent global updates, the sites with greatest u_i will be furthest from the set \mathcal{T} . As the highest sites are examined for pushing, the excess will be moved from these sites to the next lower level. These sites will then have their excess moved to the next lower level and the algorithm will “sweep” the excess




FIG. 2. Images of the heights u_i at intermediate times (indexed by the global update number g) of the PR algorithm, in the 2D RFIM for 100^2 samples, for $\Delta = 10^{-5}$ (top) and $\Delta = 1.0$ (bottom). In the former case, rearrangement of excess x_i on scales < 100 is coupled with the coalescence of small peaks into a single large peak in u_i . In the latter case, coarsening of the “terrain” given by u_i is cut off by the finite correlation length; the visible spin-up domain shows a complex “drainage” pattern which funnels the excess (not shown) to the edges of the domains. The shading is darkest at the maximum u_i and is white where $u_i = 0$.

through the volume of diameter $\approx 2\ell$. This will establish a distribution of excess that will have uniform sign over a region of size 2ℓ . In this fashion, the algorithm will solve for the sum of h_i recursively. The number of steps at each stage will be L^d and $O(\ln L)$ stages will be required, giving a total running time $\sim L^d \ln(L)$. This result is consistent with the numerical timings for very small Δ , with the data for RL^{-d} linear in $\ln(L)$ over sample dimensions $16 \leq L \leq 128$ to within the 1% numerical error for $d = 3$ and over $10^3 \leq L \leq 5 \times 10^6$ to within the same error for $d = 1$. The raising of the heights u_i and the rearrangement of the excesses x_i is a type of coarsening process, up to the scale of the correlation length. This coarsening is displayed in Fig. 2, which shows the heights as a function of time (L^d relabels per generation g).

To indicate the generality of critical slowing down and partial arguments for other systems, it is useful to compare the RFIM results with results for the 2DSG. The Hamiltonian is $H_S = \sum_{\langle ij \rangle} J_{ij} s_i s_j$, where again the s_i are Ising spins. The Gaussian distributed J_{ij} have mean J_0 and variance 1. The SG to FM transition [20] takes place as J_0 is increased through the critical value $J_c \approx 0.96$. The mapping from the 2DSG with free boundaries to a general matching problem is given in Ref. [6]: energy is minimized by selecting a state with a minimum total weight for frustrated bonds. Timing results for the 2DSG are shown in Fig. 3. One notable difference from the RFIM is the apparent convergence to *constant time per spin* in the FM phase. As the algorithm used for the 2DSG uses a bond representation, the algorithm does not need to




FIG. 3. The CPU time needed to find the ground state in the 2D spin glass, as a function of the ferromagnetic strength J_0 and system size L .

distinguish degenerate states related by global spin flips. If the frustration is low enough, the FM phase is obtained by *local* operations giving a solution with percolating unfrustrated bonds. For low J_0 , in the SG phase, in contrast, though locally the ground state is insensitive to the boundaries, the global ground state is sensitive to the disorder and the operations (augmenting paths) must therefore be carried out over all scales, so that the time per spin may diverge as a power of L (numerically, approximately as $L^{0.78 \pm 0.10}$ for $96 \leq L \leq 720$ and $J_0 = 0$). At the critical point, the time per spin scales more quickly, as $L^{1.25 \pm 0.15}$ [21].

In summary, the time for a polynomial time algorithm to find ground states is examined near zero temperature critical points for two models: the random field Ising model and the 2D spin glass. At the critical points, the combinatorial optimization algorithms used, PR flow and general matching, while exponentially faster than, say, simulated annealing at finding the ground state, slows down. This slowing down is argued to be closely related to important physical ideas, namely, the uniqueness or twofold degeneracy of the ground state in various phases and the divergence of the correlation length as the transition is approached. It will certainly be of interest to consider such slowing down and scaling at other fixed points for other polynomial time algorithms, to expand our understanding of the physics and algorithms for these systems. Such results might most directly be used for speeding up calculations: parallel computations for $\Delta > \Delta_c$ in the RFIM should be asymptotically scalable, as only operations up to length ξ are needed. The perspective provided by the algorithm timings provides independent support for the conjectured ground state structure and implies a useful dynamical viewpoint for understanding the static structure. This work complements results [8] on mean field models, where there can be many states and there is not a length scale.

I would like to thank J. Machta for useful discussions. This work was supported in part by the NSF (DMR-9702242 and DMR-0109164.)

- [1] See, e.g., M. E. J. Newman and R. Ziff, Phys. Rev. E **64**, 016706 (2001).
- [2] C. M. Fortuin and P. W. Kasteleyn, Physica (Utrecht) **57**, 536 (1972).
- [3] R. H. Swendsen and J.-S. Wang, Phys. Rev. Lett. **58**, 86 (1987); X.-J. Li and A. D. Sokal, Phys. Rev. Lett. **63**, 827 (1989).
- [4] *Spin Glasses and Random Fields*, edited by A. P. Young (World Scientific, Singapore, 1998); K. Binder and A. P. Young, Rev. Mod. Phys. **58**, 801 (1986).
- [5] M. J. Alava *et al.*, in *Phase Transitions and Critical Phenomena*, edited by C. Domb and J. L. Lebowitz (Academic, San Diego, 2001), Vol. 18.
- [6] For the case of the spin glass, see F. Barahona, J. Phys. A **15**, 3241 (1982).
- [7] C. H. Papadimitriou, *Computational Complexity* (Addison-Wesley, Reading, MA, 1994).
- [8] For the distinct behaviors of heuristics with *typical* polynomial and exponential times, see R. Monasson *et al.*, Nature (London) **400**, 133 (1999).
- [9] K. Kawasaki, in *Phase Transitions and Critical Phenomena*, edited by C. Domb and M. S. Green (Academic, New York, 1972), Vol. 2.
- [10] W. Cook and A. Rohe, INFORMS J. Comp. **11**, 138 (1999).
- [11] A. A. Middleton, Phys. Rev. E **59**, 2571 (1999).
- [12] A. T. Ogielski, Phys. Rev. Lett. **57**, 1251 (1986); A. K. Hartmann and U. Nowak, Eur. Phys. J. B **7**, 105 (1999); D. S. Fisher and A. A. Middleton, cond-mat/0107489 [Phys. Rev. B (to be published)].
- [13] J. C. Anglès d'Auriac, M. Preissmann, and R. Rammal, J. Phys. (Paris), Lett. **46**, L173 (1985).
- [14] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction To Algorithms* (MIT Press, Cambridge, MA, 1990).
- [15] A. V. Goldberg and R. E. Tarjan, J. Assoc. Comput. Mach. **35**, 921 (1988); B. V. Cherkassky and A. V. Goldberg, Algorithmica **19**, 390 (1997).
- [16] The use of negative excesses is also useful in min-cost flow algorithms.
- [17] C. M. Newman and D. S. Stein, Phys. Rev. Lett. **72**, 2286 (1994).
- [18] A. A. Middleton (unpublished).
- [19] In the case of $d = 2$, the lower critical dimension for the RFIM, the correlation length diverges exponentially at $h = 0$, likely making a quantitative investigation of the scaling quite difficult.
- [20] W. L. McMillan, Phys. Rev. B **29**, 4026 (1984).
- [21] This exponent increase most likely reflects the increase in the domain wall fractal dimension [18] at criticality and possibly the need to readjust the solution at many competing scales. Matching algorithms operate by finding paths of the domain wall dimension at successive scales, until the density of unmatched nodes goes to zero; a scaling hypothesis can be made for the time $T \sim \rho^{-d_t/d}$ to match a node at each density ρ , with d_t an algorithm-dependent dimension of a tree.