PHYSICAL REVIEW E **70**, 035103(R) (2004)

Percolation of unsatisfiability in finite dimensions

J. M. Schwarz and A. Alan Middleton
Department of Physics, Syracuse University, Syracuse, New York 13244, USA
(Received 14 September 2004; published 21 September 2004)

The optimization of two-dimensional Boolean formulas is studied using percolation theory, rare region arguments, and boundary effects. In contrast with mean-field results, there is no satisfiability transition as the constraint density is varied, although there is a logical connectivity transition. In the disconnected phase, there is a transition in the solution time. The thermodynamic ground state for this NP-hard optimization problem is unique; local solutions can be adjoined to find the global ground state. These results have implications for the computational study of disordered materials.

DOI: 10.1103/PhysRevE.70.035103 PACS number(s): 05.70.Fh, 89.20.Ff, 64.60.Ak, 75.10.Nr

Complex problems with many degrees of freedom are of interest to both physicists and theoretical computer scientists. The overlap is especially strong between the physics of disordered materials and optimization problems in the typical case. For example, there is a close correspondence between the ground states of Ising spin glasses, with up and down spins, and optimal assignments of Boolean variables, which can be true or false, in a random logical formula. This correspondence is more than formal as both systems exhibit phase transitions in the structure of minimal configurations and in the dynamics of the physical systems or optimization algorithms [1]. Combinatorial optimization algorithms from computer science are often employed to simulate disordered condensed matter systems [2]. Approaches from statistical physics, including techniques such as replica theory and concepts such as the thermodynamic limit and scaling, have proven useful in studying the running time of algorithms and the structure of solution space.

Motivated by work on mean-field Boolean formulas [3] and progress in understanding models of finite-dimensional disordered materials, we investigate ensembles of Boolean formulas whose graphs are two dimensional. These formulas are composed by conjunctively joining logical clauses, with each clause formed using nearest neighbor variables. The optimization problem is to assign truth values to the variables so as to satisfy the maximum number of clauses. This problem is analogous to minimizing the number of broken bonds in an Ising spin glass [4] and is NP hard [5], so that worst-case realizations are believed to require a time exponential in the problem size to solve. We first decompose the Boolean formulas into components that contain logical contradictions. These contradictory strongly connected components (CSCs) need not percolate even though the clauses themselves do percolate. Balancing the distribution of these clusters with the time to find their "ground state" leads to an exponential divergence in the running time at some density of clauses for a particular algorithm. We also find an exponentially rapid convergence to a unique ground state as the size of the problem increases. This suggests that the problem is easy in the typical case, though it is classified as difficult in the worst case sense. It may well be that many NP-hard problems derived from physical systems, such as finding the ground state configuration for the two-dimensional (2D) spin glass in a magnetic field [6], are typically solvable in polynomial time. Our results support this possibility. NP-hard problems with algorithms that typically take polynomial time on some problem sets are known [5], but have not been extensively and directly studied for physical problems in finite dimensions.

We consider finite-dimensional Boolean formulas Z of the form

$$Z = \wedge_{\ell=1}^{M}(\vee_{i=1}^{K} y_i^{\ell}), \tag{1}$$

where $\vee$ is the logical OR operation, $\wedge$ is the logical AND operation, and $\{y_i^{\ell}\}$ are *literals* chosen from a set $Y=\{x_1,\ldots,x_N,\overline{x_1},\ldots,\overline{x_N}\}$ of N Boolean *variables* and their negations. The variables are identified with the vertices of a two-dimensional lattice. The clauses are the terms $\vee_{i=1}^{K} y_i^{\ell}$; we specialize to $K=1,2$. A 2-clause contains two neighboring variables, each negated with probability $1/2$. The 1-clauses are single literals, with probability $1/2$ of negation. No two clauses contain the same set of variables. A sample formula is depicted in Fig. 1(a). The ensemble is defined by parameters α and γ, respectively the ratios of the number of 2-clauses to N and 1-clauses to N. Given a truth assignment $x_i \to \{T,F\}$ for all Boolean variables, a clause is satisfied if one of the literals in the clause is T. If all clauses are satisfied, the formula Z is satisfied. Determining the existence of a satisfying truth assignment is the problem of satisfiability (SAT).

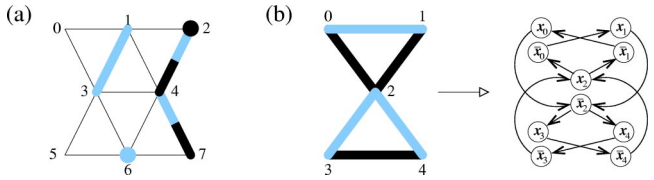

FIG. 1. (a) A finite-dimensional Boolean formula. Each 2-clause is represented by two segments on an edge. Circles represent 1-clauses. Black segments or circles indicate negated variables, while the lighter shaded segments or circles represent variables that are not negated. The formula depicted is $(x_1 \vee x_3) \wedge (x_2 \vee \overline{x_4}) \wedge (x_4 \vee \overline{x_7}) \wedge (\overline{x_2}) \wedge (x_6)$. (b) An unsatisfiable subgraph (left) and its digraph (right). The subgraph's formula is $(x_0 \vee x_1) \wedge (\overline{x_0} \vee \overline{x_2}) \wedge (\overline{x_1} \vee \overline{x_2}) \wedge (x_2 \vee x_3) \wedge (x_2 \vee x_4) \wedge (\overline{x_3} \vee \overline{x_4})$. A contradictory cycle (CC) is $x_2 \to \overline{x_0} \to x_1 \to \overline{x_2} \to x_3 \to \overline{x_4} \to x_2$.

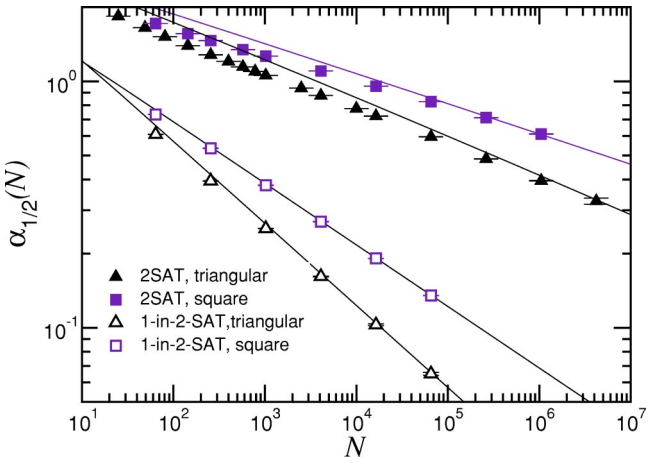

FIG. 2. Plot of $\alpha_{1/2}(N)$, the clause density at which 1/2 of the graphs are satisfiable, as a function of lattice size N. Symbols indicate numerical results for 2SAT and 1-in-2-SAT (Ising spin glass) on triangular and square lattices. Curves are analytic approximations found in a small subgraph expansion.

The optimization of the number of satisfied clauses in Z can be mapped to the problem of determining the ground state of a spin glass in a heterogeneous field. This mapping translates Boolean assignments $x_i=\{\mathrm{F},\mathrm{T}\}$ to spin variables $S_i=\{-1,+1\}$. A bond energy E_ℓ can be assigned to a clause $(y_0^\ell \vee y_1^\ell)$ connecting variables x_i and x_j via [3]

$$E_\ell = \frac{J}{4}[1-\Delta(y_0^\ell)S_i-\Delta(y_1^\ell)S_j+\Delta(y_0^\ell)\Delta(y_1^\ell)S_iS_j], \qquad (2)$$

where $\Delta(y_0^\ell)=1$ if $y_0^\ell=x_i$ and $\Delta(y_0^\ell)=-1$ if $y_0^\ell=\overline{x_i}$ and similarly for j (nearest neighbor to i) replacing 0 with 1. The total spin glass energy E is given by $E=\Sigma_{\ell=1}^{M} E_\ell$. Any clause that is not satisfied costs an energy J; the existence of an $E=0$ ground state is equivalent to satisfiability of the Boolean formula.

Resolution [5] is a method that can be used to quickly decide SAT for $K\leqslant 2$. This procedure is equivalent to mapping each 2-clause to a pair of logical implications and searching for "contradictory cycles" (CCs). For example, the clause $x_1\vee x_2$ is mapped to $\overline{x_1}\rightarrow x_2$ and $\overline{x_2}\rightarrow x_1$. Clauses with $K=1$ map to a single implication, e.g., $\overline{x_1}$ becomes $x_1\rightarrow\overline{x_1}$. The Boolean formula can be represented by an implication directed graph (digraph) $G=(Y,E)$ with $2N$ vertices and $(2\alpha+\gamma)N$ edges E. For a sample mapping, see Fig. 1(b). The formula Z cannot be satisfied if there is a CC, which is a path p in G that connects a variable to its negation and vice versa, i.e., $p=(x_i\rightarrow x_j\ldots\rightarrow\overline{x_i}\rightarrow\ldots\rightarrow x_i)$. The existence of a CC can be decided in time linear in the size of E [7].

We find that there are CCs for any $\alpha>0$ (taking $\gamma=0$), as $N\rightarrow\infty$. Defining $\alpha_{1/2}(N)$ as the value of α for which 1/2 of the N-variable formulas are satisfiable, $\alpha_{1/2}(N)\rightarrow 0$ as $N\rightarrow\infty$ (see Fig. 2). This crossover is coarse, in that the width of the crossover from low to high probability of satisfiability is proportional to $\alpha_{1/2}(N)$ for large N. This is to be contrasted with random mean-field $K=2$ formulas where, for $N\rightarrow\infty$, there is a sharp SAT to UNSAT phase transition (the probability that a formula is satisfiable is 1 for $\alpha<\alpha_c=1$ and 0 for larger α). This difference results from small CCs, which are exponentially rare in the mean-field case but appear with Poissonian statistics in finite dimensions at any α.

The location of the SAT/UNSAT crossover can be computed by an expansion in α. Some subgraphs are "forcing," i.e., in all satisfying assignments one of the variables has a fixed truth value. The smallest unsatisfiable graph is found by joining two contradictory forcing subgraphs. An example of this graph type is depicted in Fig. 1(b). On the triangular lattice, these subgraphs have density $\rho_\triangle(\alpha)=(\alpha^6/2^7 3^4)+O(\alpha^7)$. The density of the simplest unsatisfiable graphs on the square lattice is $\rho_\square(\alpha)=(\alpha^8/2^{16})+O(\alpha^9)$. In general, if the smallest unsatisfiable subgraph has r bonds and density $c_r\alpha^r$, the probability of satisfiability is $P_{\mathrm{SAT}}(N)=(1-c_r\alpha^r)^N$, to lowest order in α, giving $\alpha_{1/2}(N)\approx(c_r^{-1/r}\ln 2)N^{-1/r}$. We plot numerical results and analytic expansions for $\alpha_{1/2}(N)$ in Fig. 2, which includes the analytic corrections to next order in α [seven-edged subgraphs with density $(5^2/2^7 3^7)\alpha^7+O(\alpha^8)$ on the triangular lattice and nine-edged subgraphs with density $(\alpha^9/2^{16})+O(\alpha^{10})$ on the square lattice].

We also plot analytic estimates and numerical results for $\alpha_{1/2}(N)$ for the 1-in-2-SAT problem in Fig. 2. While a clause in 2SAT (i.e., $K=2$) is satisfied if either literal is true, a clause is satisfied in 1-in-2-SAT when exactly one literal in a clause is true. The 1-in-2-SAT problem maps both to an Ising spin glass in the absence of a magnetic field and to the two-color problem [8]. The smallest unsatisfiable graphs are frustrated cycles, giving $\alpha_{1/2}(N)\approx 3(N/\ln 2)^{-1/3}$ and $\alpha_{1/2}(N)\approx 2^{5/4}(N/\ln 2)^{-1/4}$, for the triangular and square lattices, respectively.

Given the lack of a sharp SAT/UNSAT transition, due to the existence of *small* unsatisfiable graphs, we have investigated the percolation of *large* unsatisfiable graphs as a phase transition. We study these graphs within the context of MAXSAT, which is the problem of minimizing the number of unsatisfied clauses. In two dimensions, the determination of the ground state for the Ising spin glass (or MAX-1-in-2-SAT) is in P [6], while determining the ground state for planar MAXSAT with $K=2$ is NP-hard. We studied the CSCs, sets of literals that contain a CC and for which there is a directed path from any literal to any other in the set. We find that the probability of having a spanning CSC has a transition that becomes sharper with increasing N, with a critical value for α of $\alpha_S=1.82454(5)$ on the triangular lattice [$\alpha_S=1.8128(6)$ on the square lattice]. The cluster size distribution $n(s,\alpha)$ at criticality behaves as $n(s,\alpha_S)\sim s^{-\tau}$, with $\tau=2.02(5)$. The scaling of the probability for a spanning CSC near α_S gives a correlation length exponent of $\nu=1.32(3)$. These values are consistent with standard percolation, where $\tau=187/91$ and $\nu=4/3$ [9].

In mean-field graphs, there is a connection between the percolation of directed half-cycles (paths from x to $\overline{x}$) and connectivity percolation [10]. This connection is made by trimming paths in the implication digraph of $2N$ literals and projecting them onto a connectivity graph of N variables. The edge probability p in the connectivity graph is related to the 2-clause probability $\tilde{p}$ via $\tilde{p}=2p-p^2$. The construction directly carries over to finite dimensions. However, in mean field, the half-cycle percolation coincides with the SAT/

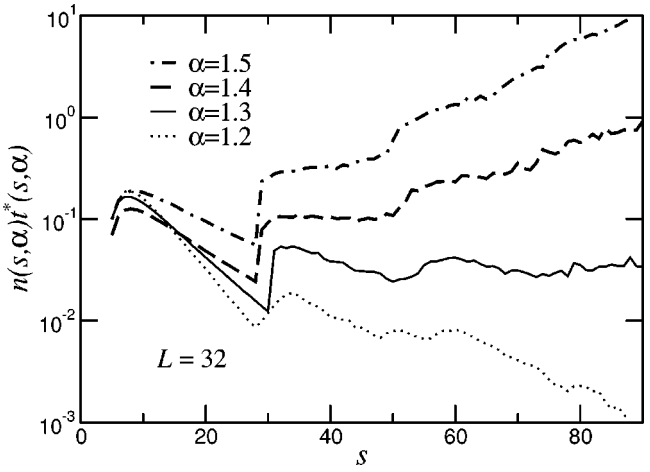

FIG. 3. Convolution of the median number of backtracks $t^*(s,\alpha)$ with the CSC cluster size distribution $n(s,\alpha)$, where s is the cluster mass. The clause density α is less than the percolation value $\alpha_S \approx 1.8$, for these curves.

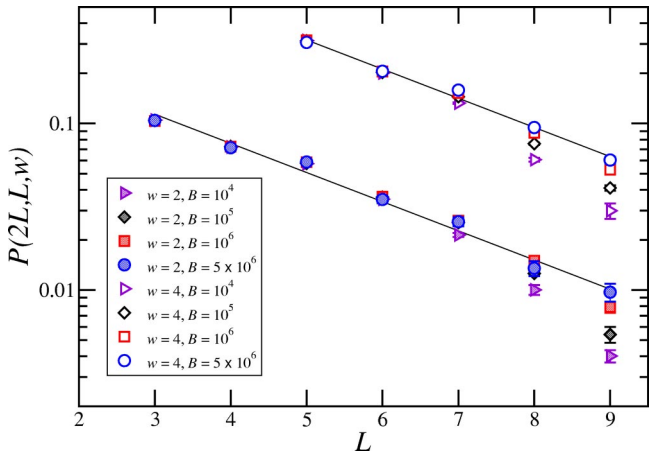

FIG. 4. Logarithm-linear plot of $P(2L,L,w)$ for $\alpha=1.7$ and $\gamma=0.2$ for weighted MAX2SAT. The bound on backtracks is B. The lines are exponential fits for the $w=2,4$, $B=5\times10^6$ data.

UNSAT transition, while in finite dimensions we are interested in CSC percolation. Certainly the half-cycle percolation threshold sets a lower bound on CSC percolation and it is not surprising that CSC percolation appears to be in the same universality class as connectivity percolation (also see [11]).

The decomposition of the graphs into CSCs speeds up exact search algorithms for MAXSAT. We apply this decomposition to estimate running times. We used a MAXSAT code [12] that first finds a heuristic bound to the solution and then applies an exact Davis-Putnam-Loveland-Logemann (DPLL) search. Each CSC cluster is loaded into the algorithm individually [13]. The running time measure t is the number of "backtracks" that are executed while partially exploring the tree of all possible assignments. The sum of the unsatisfied clauses from each cluster gives the minimal number of unsatisfied clauses for the entire formula. As is to be expected from percolation theory, when $\alpha<\alpha_S$, the distribution of sizes of the CSCs, is exponentially decaying in the cluster size, $n(s,\alpha)\sim e^{-s/s_\xi(\alpha)}$, with $s_\xi \propto \xi^d \propto (\alpha_S-\alpha)^{-d\nu}$. In addition, plotting our results for the median number of backtracks for each cluster, we find that the median running time of the DPLL-type algorithm grows exponentially with the cluster size, $t^*(s,\alpha)\sim e^{s/s_\tau(\alpha)}$. Balancing these two exponentials implies the existence of a transition in the behavior of the running time on N. When $s_\xi(\alpha)<s_\tau(\alpha)$, the median running time for a sample, $T^*(L,\alpha)$, is bounded by a multiple of the system volume, $T^*(L,\alpha)\sim L^2$. However, when $s_\xi>s_\tau$ an estimate for the largest cluster size in a *typical* finite sample gives $T^*(L,\alpha)\sim L^{2s_\tau/s_\xi}$. The *mean* running time, $\overline{T}(L,\alpha)$, diverges exponentially with L. The separation between linear and superlinear median time behaviors defines the transition location α_G, with $\alpha_G<\alpha_S$, via $s_\xi(\alpha_G)=s_\tau(\alpha_G)$. Figure 3 shows convolutions of the cluster size distribution $n(s,\alpha)$ and the median time $t^*(s,\alpha)$ as a function of size. The change from negative to positive slope on the semilogarithmic plot gives $\alpha_G\approx1.3$ for the DPLL code we use. This slowing down of the algorithmic dynamics is similar to that for the dynamics of random magnets [14] and is reminiscent of the change from the easy-SAT to hard-SAT phases in random graphs [1].

Despite the divergence of the running times for the exhaustive DPLL-type algorithms, we expect that the ground states can be found in time proportional to the system volume in the typical case, even above the CSC percolation transition. Assuming that the droplet picture describes finite-dimensional spin glasses, the presence of a random magnetic field destroys the spin glass phase [15]. So while the CSCs percolate, the effects of frustration remain localized over some length scale. The ground state is unique and insensitive to boundaries. If the convergence to a unique state is *exponential with size*, a heuristic algorithm of solving subsystems and patching together the subsolutions might be useful.

To test whether the ground state is unique, we study the effect of boundary conditions, similar to studies of the Ising spin glass [16]. By comparing ground states for a system of linear size L and an expanded system of linear size $L'>L$, one can determine if the ground state is unique. If the solutions in a common subsystem of linear size w become fixed as L and L' diverge, a unique ground state exists in the thermodynamic limit. Note that the ground state must be unique for $\alpha<\alpha_S$, as the logical structure of the graph does not percolate.

Since the $\pm J$ spin glass with magnetic field (equivalent to optimal assignments for MAX2SAT) has many degenerate ground states, we study the weighted MAX2SAT (WMAXSAT) question, where the degeneracy is broken, to be able to compare ground state solutions directly. We add 1-clauses, $\gamma=0.2$, to study sparser graphs $[\alpha_S(\gamma=0.2)<\alpha=1.7]$ and thus a larger range system of system sizes. Each clause has a weight, chosen uniformly in the interval $[0,1)$, and the optimization problem is now to minimize the sum of the weights of the unsatisfied clauses.

We estimate $P(2L,L,w)$, the probability that there is a change in the central region of area w^2 when the system size expands from L to $2L$ [16], by sampling from the WMAXSAT ensemble. To be able to complete the simulations, we impose an upper limit B on the number of back-

tracks in DPLL (some samples are not solved) and study P for increasing B. For both $w=2$ and $w=4$, P is well fit by exponentials in L with the same slope, in the limit of large B (see Fig. 4). (A power-law fit gives an exponent less than -2, which is inconsistent with a fractal domain wall picture for a model with two states [16].) For $\alpha=1.7$ and $\gamma=0.2$, we estimate a correlation length of $\xi=2.5\pm0.3$. The exponential approach to a unique state holds for all α and γ that we explored.

Given this result, we tested a heuristic WMAXSAT algorithm for samples of length L' in which subsamples of length L are solved exactly, with $\xi \ll L < L'$. For each subsample, the truth assignment within a window of length $w \ll L$ centered on the subsample is recorded. This process is iterated until the windows cover the sample. Given the exponential convergence to a unique state, we expect the mean error of this heuristic to scale no worse than $(L')^2 e^{-L/\xi}$. We tested this expectation directly by comparing the patched solution with the exact global solution. For $\alpha=1.7$, $\gamma=0.2$, $L'=16$, and $w=2$, the mean error n_{err} is well fit by an exponential in L, $n_{\mathrm{err}} \sim \exp(-L/\xi)$ with $\xi=2.2\pm0.1$. The ability of this algorithm to find exact global solutions, controlled by subsample size L, is to be contrasted with similar approaches to the traveling salesperson problem, where such a decomposition method provides an excellent approximation [17], but, with high probability, fails to give an exact result. We note that this method would also fail for the 2D Ising spin glass, where $P(2L,L,w)$ is a power law rather than an exponential.

To conclude, we have studied the problem of optimally satisfying Boolean formulas in finite dimensions. While there is no thermodynamic SAT to UNSAT transition, there is a percolation transition in the logical structure of the formulas as the clause density is increased, that is apparently in the class of standard percolation. Below this transition, we use rare region arguments to predict a transition in the mean running time of an optimization algorithm. We also find that the ground state is unique even in the high clause density regime; the exponential insensitivity to the boundaries suggests that patching algorithms can provide exact solutions with high probability.

This project was supported in part by the National Science Foundation through grant No. DMR-0109164.

[1] G. Semerjian and R. Monasson Phys. Rev. E **67**, 066103 (2003); W. Barthel, A. K. Hartmann, and M. Weigt, *ibid.* **67**, 066104 (2003).

[2] A. K. Hartmann and H. Rieger, *Optimization Problems in Physics* (Wiley-VCH, Berlin, 2002).

[3] R. Monasson *et al.*, Nature (London) **400**, 133 (1999); O. C. Martin *et al.*, Theor. Comput. Sci. **265**, 3 (2001); M. Mezard *et al.*, Science **297**, 812 (2002); A. K. Hartmann and M. Weigt, J. Phys. A **36**, 11069 (2003).

[4] K. Binder and A. P. Young, Rev. Mod. Phys. **58**, 801 (1986).

[5] See, e.g., C. H. Papadimitriou, *Computational Complexity* (Addison-Wesley, Reading, MA, 1994). Two classes in computational complexity theory are P and NP decision (yes/no) problems. Problems in P can be *decided* in time polynomial in the size of the problem description, while a proof of the answer for NP problems can be *checked* in polynomial time. NP-hard optimization problems are at least as hard as NP decision problems.

[6] F. Barahona, J. Phys. A **15**, 3241 (1982).

[7] T. H. Cormen, C. E. Lieserson, and R. L. Rivest, *Introduction to Algorithms* (MIT Press, Cambridge, MA, 1990).

[8] T. Walsh, in *Proceedings of AAAI 2002* (AAAI Press, Menlo Park, CA, 2002).

[9] D. Stauffer and A. Aharony, *Introduction to Percolation Theory* (Taylor and Francis, London, 1992).

[10] B. Bollobas *et al.*, Random Struct. Algorithms **18**, 201 (2001).

[11] F. O. Pfeiffer and H. Rieger, Phys. Rev. E **67**, 056113 (2003).

[12] B. Borchers and J. Furman, J. Comb. Optim. **2**, 299 (1999).

[13] As each CSC is strongly connected, assignment of truth values to variables outside of a given CSC does not affect satisfying assignments for that CSC. Note that weighted MAXSAT might lead to CSC-CSC interactions.

[14] M. Randeria, J. P. Sethna, and R. G. Palmer, Phys. Rev. Lett. **54**, 1321 (1985).

[15] D. S. Fisher and D. A. Huse, Phys. Rev. B **38**, 386 (1988).

[16] C. M. Newman and D. L. Stein, J. Phys.: Condens. Matter **15**, R1319 (2004); A. A. Middleton, Phys. Rev. Lett. **83**, 1672 (1999); M. Palassini and A. P. Young, Phys. Rev. B **60**, R9919 (1999).

[17] R. M. Karp, Math. Methods Oper. Res. **2**, 209 (1977); S. Arora, J. ACM **45**, 753 (1998).