Theory of Nonequilibrium Local Search on Random Satisfaction Problems

Erik Aurell,[1,*] Eduardo Domínguez,[2] David Machado,[2] and Roberto Mulet[2]
[1]Department of Computational Science and Technology, AlbaNova University Center, SE-106 91 Stockholm, Sweden
[2]Group of Complex Systems and Statistical Physics. Department of Theoretical Physics, Physics Faculty, University of Havana, CP 10400 La Habana. Cuba

(Received 8 March 2019; revised manuscript received 8 October 2019; published 4 December 2019)

We study local search algorithms to solve instances of the random k-satisfiability problem, equivalent to finding (if they exist) zero-energy ground states of statistical models with disorder on random hypergraphs. It is well known that the best such algorithms are akin to nonequilibrium processes in a high-dimensional space. In particular, algorithms known as focused, and which do not obey detailed balance, outperform simulated annealing and related methods in the task of finding the solution to a complex satisfiability problem, that is to find (exactly or approximately) the minimum in a complex energy landscape. A physical question of interest is if the dynamics of these processes can be well predicted by the well-developed theory of equilibrium Gibbs states. While it has been known empirically for some time that this is not the case, an alternative systematic theory that does so has been lacking. In this Letter we introduce such a theory based on the recently developed technique of cavity master equations and test it on the paradigmatic random 3-satisfiability problem. Our theory predicts the qualitative form of the phase boundary between the satisfiable (SAT) and unsatisfiable (UNSAT) region of the phase diagram where the numerics of a focused Metropolis search and cavity master equation cannot be distinguished.

DOI: 10.1103/PhysRevLett.123.230602

Introduction.—Nonequilibrium phenomena and the structure of models with disorder have been two major frontiers of statistical physics in the last half century. Famous results include the discovery of a host of equalities which generalize the second law [1–3], and the solution of (equilibrium) spin glass models, first with all-to-all couplings [4], and later with random-graph (dilute) couplings [5].

An equilibrium distribution can be generated by a dynamic rule where variables change state according to local change of energy. This is the core of the Monte Carlo procedure of scientific computing, and in particular of the method of simulated annealing [6]. The physical principle that underlies such methods is that a process obeying detailed balance will eventually converge towards a Gibbs distribution, although the time required may be exponential in problem size for models in a spin glass phase.

The combination of nonequilibrium processes and complicated high-dimensional energy functions has a natural application to combinatorial optimization problems. As it is now widely appreciated, these are equivalent to the physical problem of finding ground states in statistical mechanics models with disorder, an analogy which has generated a large amount of literature [5]. Constraint satisfaction is the subset where the energy function is non-negative, and the problem is to find a zero-energy ground state (if any exists).

Equivalently, methods to solve combinatorial optimization problems that rely on local rules can be understood as dynamical processes of high-dimensional spin systems. The dynamical rules can, as for simulated annealing, obey detailed balance, but do not have to. This Letter introduces a new theory to describe nonequilibrium algorithms which find solutions to the paradigmatic case random 3-satisfiability.

Random 3-satisfiability and the related random constraint satisfaction problem have been central benchmark quantum annealing and quantum optimization algorithms [7–12]. Comparisons have however so far exclusively been to processes which like simulated annealing work close to equilibrium, and which perform far less well than the out-of-equilibrium processes studied here. Our Letter therefore also injects into these discussions the best-performing classical local algorithms within the purview of physical analysis.

The random 3-satisfiability problem and local search.—Boolean satisfiability is the problem to determine if a set of logical clauses can be satisfied. Random K-satisfiability (K-SAT) is the problem where all clauses have cardinality K, and are picked at random. Although it is one of the basic examples of a worst-case hard computational problem [13], the empirical run time of solving a given instance of a random 3-satisfiability problem varies greatly [14]. Very underconstrained problems are in practice easy (SAT), the difficult region is for problems that are on the verge of being unsatisfiable (UNSAT) which for 3-satisfiability (3-SAT), means a ratio of clauses to variables (M/N) of about 4.27. Very overconstrained problems are again easy

for complete algorithms (i.e., algorithms that either find a solution or give a proof that there is no solution), but this aspect will not be discussed further here.

In the run-up from underconstrained to critical 3-SAT problems, different algorithms can be characterized where they fail to work. Working here means finding a solution with high probability, in time scaling as a small power in the system size, and failing means everything else. According to this criterion the best algorithm for random K-SAT is "survey propagation" [15] which in its most recent version is able to find solutions extremely close to the SAT-UNSAT threshold [16]. Survey propagation is however a quite complex algorithm tailored to random constraint satisfaction problems, and is not competitive on most real-world problems [17]. It is therefore of general interest to step back and consider other simpler and more general solution procedures, of which the first example is "simulated annealing" [6], a workhorse of scientific computing. The performance of simulated annealing at slow enough cooling rate can be analyzed by spin glass techniques [18] and is known to fail at some distance from the SAT-UNSAT threshold.

The best local search algorithms that have been invented for satisfiability are however not processes in detailed balance. They all rely on "focusing," meaning that only variables that participate in some unsatisfied clause are considered for update. A focused algorithm hence obeys the dictum "if it works, don't fix it." For constraint satisfiability, the most well-known algorithm in this class is "walksat" [19] which is also competitive on many real-world problems [17,20], and which works on random 3-satisfiability up to a clause density of about 4.2 [21]. Several other local search procedures have been shown to work up to a similar threshold [21–27].

Here we will consider a focused Metropolis search (FMS) [23,25]. This algorithm can be described very simply as first making a focusing step and then a standard Metropolis step, as in simulated annealing at one temperature. For the best choice of temperature a FMS has been empirically shown to work up to a clause density of about 4.23 [23]. With the notable exception of the average rate equations of Refs. [20,28] that only work for special cases, theory has been lacking for these important processes. That is what we will provide here.

Cavity master equation applied to a random FMS for 3-SAT.—The cavity master equation is a closure of the dynamic cavity equations. The dynamic cavity starts from the joint probability distribution of all histories of a set of dynamic variables interacting in a locally treelike (locally loop-free) graph. It is then possible to write a self-consistent equation for the probabilities of the history of single variables when the history of one of their neighboring variables is held fixed; one says that the first variable is in the cavity of the second variable. These dynamic cavity equations are formally belief propagation updates. As is, they are of little practical value since the variable (the history of one dynamic variable) is very high dimensional. For dynamics in discrete time with synchronous updates, closure assumptions have been explored for some time [29–31].

The cavity master equation is appropriate for dynamics of discrete variables in continuous time. In satisfiability problems these variables naturally take values 1 (true) or -1 (false), which we here call spins. The cavity master equation takes as input the jump rates r_i (for spin i) defining the dynamics, and is for spins, interacting in groups labeled by $a, b, c, \ldots$ (constraints, clauses) formulated in terms of quantities $p_{a\to i}(\sigma_{a\setminus i}|\sigma_i)$ where $\sigma_{a\setminus i}$ are the current values in group a except i [32]. These quantities should be considered closures imposed on the corresponding full cavity quantities $\mu_{a\to i}(X_{a\setminus i}|X_i)$ where $X_{a\setminus i}$ is the whole history of all the spins in group a except i in the cavity of i, and X_i the cavity history.

In practice, to describe a FMS on random K-SAT we then have to solve the following set of coupled differential equations:

$$\begin{aligned}\dot{p}(\sigma_{a\setminus i}|\sigma_i) &= -\sum_{j\in a\setminus i}\sum_{\substack{\{\sigma_{b\setminus j}\}\\ b\in\partial j\setminus a}} r_j(+)\prod_{b\in\partial j\setminus a} p(\sigma_{b\setminus j}|\sigma_j)p(\sigma_{a\setminus i}|\sigma_i)\\ &\quad+\sum_{j\in a\setminus i}\sum_{\substack{\{\sigma_{b\setminus j}\}\\ b\in\partial j\setminus a}} r_j(-)\prod_{b\in\partial j\setminus a} p(\sigma_{b\setminus j}|-\sigma_j)p(F_j[\sigma_{a\setminus i}]|\sigma_i). \qquad (1)\end{aligned}$$

F_j in equation above is the flip operator acting on spin j while the combination of several terms of the type $p(\sigma_{a\setminus i}|\sigma_i)$ is characteristic of the cavity master equation closure, and structurally analogous to the earlier described case of the (ferromagnetic) p-spin model [32]. The term $r_j(\pm)$ in Eq. (1) is the jump rate of spin j when it takes value $\pm\sigma_j$. This quantity depends on the instantaneous value of spin j and on the instantaneous values of all the spins interacting with j, through all the clauses in which spin j appears. To describe the dynamics of the FMS algorithms one takes

$$r_i = \frac{E_i(\sigma_i,\sigma_{\partial i})}{KE}\min\left[e^{-\beta\Delta E(\sigma_i,\sigma_{\partial i})},1\right] \qquad (2)$$

where the expression in brackets is the standard Metropolis factor and $E_i(\sigma_i,\sigma_{\partial i})$ is the number of unsatisfied constraints of which spin i is a member. Each of these can be written

$$E_a = \frac{1}{2^K}\prod_{i\in a}(1-l_i^a\sigma_i) \qquad (3)$$

where the l_i^a are quenched random variables taking values ±1. When σ_i equals l_i^a, one says that variable i satisfies

clause a, and the value of E_a is zero. If $\sigma_i \neq l_i^a$ for all $i \in \partial a$ the value of E_a is one, and one says that clause a is unsatisfied. K satisfiability is thus a mixture of p-spin problems, where p ranges from 1 to K. In the focusing step of a FMS all variables partaking in unsatisfied clauses are picked with probability proportional to $E_i(\sigma_i, \sigma_{\partial i})$, which explains this factor in Eq. (2).

To model the dynamics of a FMS in overall algorithmic time (wall-clock time), we have to further take into account that the number of unsatisfied clauses changes. When this becomes smaller the rate per unit time of a given unsatisfied clause to be picked goes up. This is reflected by the denominator KE in Eq. (2), where K is the number of variables per clause (3 for 3-SAT) and E is the total number of unsatisfied clauses. This globally defined time reparametrization kicks in strongly when there are only a few unsatisfied clauses left, and when the variables in these clauses are probed more often. By a more efficient coding one can bring down the number of sums in Eq. (1) from $2^{(K-1)c}$ to 2^c (c is the number of clauses per variable). This coding is described in the Supplemental Material [33] (see also references therein [34–36])

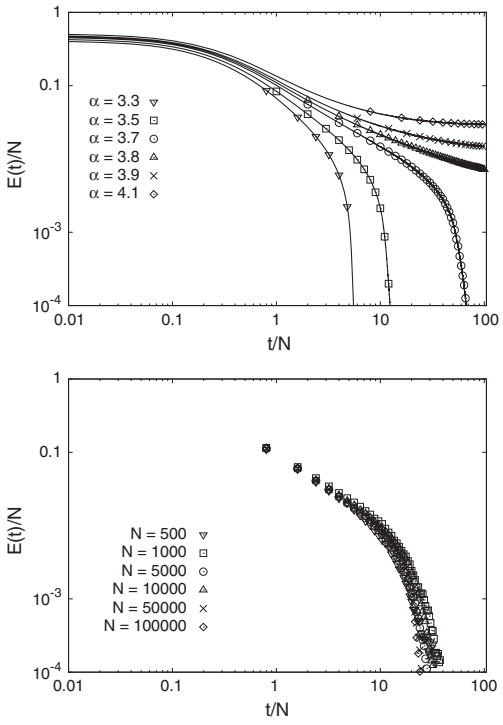

FIG. 1. FMS results on 3-SAT instances. Both panels show the number of unsatisfied clauses (system energy) as a function of time, in log-log scale. (Top) Dependency on α of FMS behavior. There is a transition between when a FMS works (finds a solution in integration time considered), and when it does not. These calculations were done with $\eta = 0.45$ and system size $N = 10^5$. Averages over 500 different histories were made for each α. (Bottom) The FMS's dependency on N for $\eta = 0.45$ and $\alpha = 3.6 < \alpha_c$. Lines are a guide to the eyes.

Results.—The problem is defined by the ratio between the number of clauses (M) and the number of literals (N) of some given instance of 3-SAT written as $\alpha = M/N$, and by $\eta = e^{-\beta}$ as the noise parameter that enters into the rates of Eq. (2). In order to understand the behavior of a FMS we need to study its dependence on these two parameters.

For a given noise η, a FMS has been empirically shown to have a zone, for α lower than some $\alpha_c(\eta)$, where it solves 3-SAT instances in times linear with system size N. For $\alpha \geq \alpha_c$ solutions are found in times that grow exponentially with N, or solutions do not exist. This is shown in Fig. 1. As can be seen in the top panel of this figure, for $\eta = 0.45$, a FMS is able to solve instances that have $\alpha \leq 3.7$, and seems to fail otherwise. In the bottom panel, size effects are represented. For $\alpha = 3.6$, FMS results seem to be almost independent of N.

Then, by numerically integrating Eqs. (1) one can obtain the behavior for the same values of the parameter η. Results can be seen in Fig. 2. Although the transition α is not identical to Fig. 1, the results of the CME are qualitatively

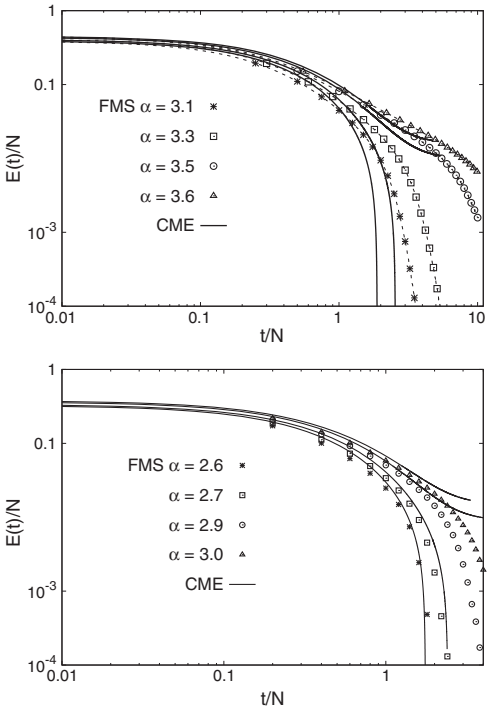

FIG. 2. CME results on 3-SAT instances. Both figures show the number of unsatisfied clauses (system energy) as a function of time, in log-log scale for different values of α. (Top) Comparison between the CME (lines) and FMS (points) with $\eta = 0.4$ and system size $N = 2000$. (Bottom) Comparison between the CME (lines) and FMS (points) for $\eta = 0.65$ and $N = 5000$. There is a transition between a phase in which solutions of the CME reaches zero energy in finite time, and a phase where they do not. In the region where frustration increases, i.e., for high α, the CME will not reach zero energy even when the FMS typically is able to find solutions. In this region, long range fluctuations in time and/or in the graph are important.

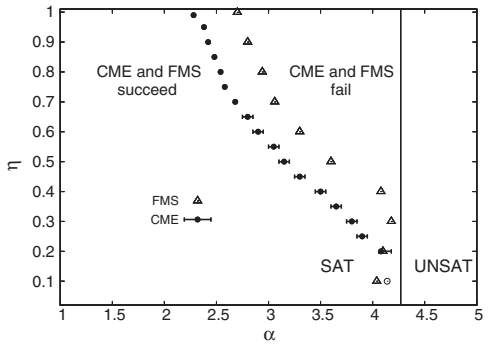

FIG. 3. Comparison between the phase diagrams of a focused Metropolis search (FMS) and cavity master equation (CME). The phase boundary of a FMS was obtained by running 100 instances of the problem at different α for a time $10^5 \cdot N$ and determining at which value of α half of them find a solution in the given time (filled-in circles). Convergence is slower in the lower "descending" branch of a FMS. The phase diagram of the CME was found by integrating the CME for 24 instances (thick-line circles) or at least 4 instances (thin-line circles) of the problem at different α. Results for each α were placed in log-log plots as the ones of Fig. 2 and it was determined for which value of α half of them did not converge to zero. We used $N = 5000$ for $\eta \geq 0.65$, $N = 2000$ for $0.4 \leq \eta \leq 0.65$, and $N = 500$ for $\eta < 0.4$.

very similar. The differences are that the CME, as is natural of the solution of a set of ordinary differential equations, either converges to zero fairly rapidly, or does not converge to zero. The zone where a FMS solves the problem by fluctuations is hence not well described by the CME. The predicted threshold of the CME (α_c for given η) is thus generally slightly smaller than the empirically determined threshold of a FMS.

As illustration of the quality of our approximation on Fig. 2 a comparison is made between the CME and FMS, for $\eta = 0.65$ and $\eta = 0.4$ near the transition, and several values of α, (see also the Supplemental Material [33] for more plots). Figure 3 shows the phase diagrams of a FMS and CME. For the FMS the phase boundary is obtained as the last value of α where in 50% of the runs the number of unsatisfied clauses goes to zero. We used 20 random graphs and for each one 10 000 histories. For the CME the phase boundary is similarly obtained for a given η, as the last value of α where in at least half of the graphs (only one run of course) the number of unsatisfied clauses goes to zero. Close to the phase boundary integration, times are thus in both cases long, and nontrivial correlations between nearby variables have time to build up. If a variable is correlated with nearby variables at one time it will be correlated with its own state at a later time. Since the CME builds on a closure assumption where temporal correlations are ignored, it is not surprising that the two phase boundaries do not coincide, i.e., that there is a gap. Far from the phase transition, however, the approximation is much better. This is more evident if it is compared with the results obtained assuming, that the spins are independent (see the Supplemental Material [33]). In this case, one can easily see that the local version of the master equation fails completely to describe the dynamical behavior of a FMS while the CME follows very close the results of the simulations.

The size and precise form of the gap is more difficult to predict. As the correlations which the CME ignores contain information that a FMS may explore, the phase boundary for the CME should be shifted to the left relative to the FMS. Below $\eta \leq 0.3$, in the region of optimal performance of the FMS, the gap however shrinks. We conjecture that this better agreement may be the result of two competing effects. One is the buildup of correlations, and the other is that in this region, once the FMS hits a local minima, it gets trapped, and its behavior is dominated by fluctuations. The CME instead explores the whole configuration space, neglects fluctuation effects, which would lead it to hit the absorbing boundary of no unsatisfied clauses relatively faster. In any case, it is important to note that neither the FMS or CME appears to be limited by or sensitive to the dynamical transition from spin glass theory which occurs at $\alpha_d = 3.9$.

Discussion.—The qualitative and quantitative description of the energy landscapes in combinatorial optimization problems is one of the most important results of statistical physics of disordered systems, with many applications in many areas of science [4,5]. The quantitative prediction of the exact threshold between a SAT and an UNSAT phase in random satisfiability problems by a one-step replica symmetry breaking (1RSB) technique was a breakthrough [15], which has been extended to many other paradigmatic problems in computer science such as, e.g., graph coloring [37] vertex covering [38], and the stochastic block model [39].

Yet, these advances *a priori* describe statics, and not dynamics. A long line of empirical investigations surveyed in the introduction have shown that the phase diagram of the nonequilibrium local search appears unrelated to bounds derived from the complexity of (equilibrium) free energy landscapes. A further and more recent discussion that nonequilibrium may be "unreasonably effective" was given recently in Ref. [40] and similarly in Ref. [41]. For combinatorial optimization nonequilibrium may make possible to achieve what has been posited to be impossible, from equilibrium considerations. A full realization (and exploitation) of these results has however been hampered by a lack of systematic theory. This is what we have furnished here, by adapting recent advances in the description of dynamics on locally treelike graphs to the understanding of the FMS.

Our theory for how the local search proceeds in time is very accurate, in the parameter regime near the optimal performance of the algorithm. We find that at large η's our algorithm described well the form of the phase boundary, although it is shifted a constant to a somewhat smaller value

of α (clause density). On the contrary, in the high density (large α) regime, fluctuations dominate the behavior of the FMS, and the CME keeps properly exploring the energy landscape of the system. So, the gap between both transitions vanishes. The last results suggest that there is also room to use the CME as the basis to develop optimized algorithms in the same spirit already tested in Ref. [32]. We note that also in the simpler case of synchronously updated spin systems (parallel updates) it was possible to further improve on an analogous Markov approximation presented in Ref. [31] by using the matrix product approximation of quantum many-body theory [42]. It will be interesting to try a similar approach for the continuous dynamics tested here.

We acknowledge support from the European Union Horizon 2020 research and innovation programme MSCA-RISE-2016 under Grant Agreement No. 734439 INFERNET and by an Erasmus + International Credit Mobility to KTH (EU).

*eaurell@kth.se

[1] E. Sevick, R. Prabhakar, S. R. Williams, and D. J. Searles, Annu. Rev. Phys. Chem. **59**, 603 (2008).
[2] C. Jarzynski, Annu. Rev. Condens. Matter Phys. **2**, 329 (2011).
[3] U. Seifert, Rep. Prog. Phys. **75**, 126001 (2012).
[4] M. Mézard, G. Parisi, and M. Virasoro, *Spin Glass Theory and Beyond* (World Scientific, Singapore, 1987).
[5] M. Mézard and A. Montanari, *Information, Physics and Computation* (Oxford University Press, Oxford, 2009).
[6] S. Kirkpatrick, C. Gelatt, and M. P. Vecchi, Science **220**, 671 (1983).
[7] H. G. Katzgraber, F. Hamze, and R. S. Andrist, Phys. Rev. X **4**, 021008 (2014).
[8] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer, Science **348**, 215 (2015).
[9] L. Hormozi, E. W. Brown, G. Carleo, and M. Troyer, Phys. Rev. B **95**, 184416 (2017).
[10] Y. Susa, Y. Yamashiro, M. Yamamoto, and H. Nishimori, J. Phys. Soc. Jpn. **87**, 023002 (2018).
[11] K. Kudo, Phys. Rev. A **98**, 022301 (2018).
[12] R. Harris *et al.*, Science **361**, 162 (2018).
[13] M. R. Garey and D. S. Johnson, *Computers and Intractability* (Freeman, New York, 1979).
[14] D. Mitchell, B. Selman, and H. Levesque, in *Proceedings of the Tenth National Conference on Artificial Intelligence* (AAAI Press, San Jose, 1992), pp. 459–465.
[15] M. Mézard, G. Parisi, and R. Zecchina, Science **297**, 812 (2002).
[16] R. Marino, G. Parisi, and F. Ricci-Tersenghi, Nat. Commun. **7**, 12996 (2016).
[17] H. Kautz and B. Selman, Discrete Appl. Math. **155**, 1514 (2007).
[18] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborová, Proc. Natl. Acad. Sci. U.S.A. **104**, 10318 (2007).
[19] B. Selman, H. Kautz, and B. Cohen, in *Coloring, and Satisfiability: Second DIMACS Implementation Challenge*, Vol. 26, edited by D. S. Johnson and M. A. Trick (American Mathematical Society, Rhode Island, 1996), pp. 521–532.
[20] W. Barthel, A. K. Hartmann, and M. Weigt, Phys. Rev. E **67**, 066104 (2003).
[21] E. Aurell, U. Gordon, and S. Kirkpatrick, in NIPS, Vol. **17**, p. 49 (2005).
[22] U. Schöning, in Foundations of Computer Science (IEEE Computer Society, Washington, 1999), Vol. **40**, pp. 410–414.
[23] S. Seitz, M. Alava, and P. Orponen, J. Stat. Mech. (2005) P06006.
[24] J. Ardelius and E. Aurell, Phys. Rev. E **74**, 037702 (2006).
[25] M. Alava, J. Ardelius, E. Aurell, P. Kaski, S. Krishnamurthy, P. Orponen, and S. Seitz, Proc. Natl. Acad. Sci. U.S.A. **105**, 15253 (2008).
[26] L. Kroc, A. Sabharwal, and B. Selman, in *Applications of Satisfiability Testing—SAT 2010*, Vol. 6175, edited by O. Strichman and S. Szeider (Springer Verlag, Berlin-Heidelberg, 2010), pp. 346–351.
[27] R. Lemoy, M. Alava, and E. Aurell, Phys. Rev. E **91**, 013305 (2015).
[28] G. Semerjian and M. Weigt, J. Phys. A **37**, 5525 (2004).
[29] I. Neri and D. Bollé, J. Stat. Mech. (2009) P08009.
[30] Y. Kanoria and A. Montanari *et al.*, Ann. Appl. Probab. **21**, 1694 (2011).
[31] G. Del Ferraro and E. Aurell, Phys. Rev. E **92**, 010102(R) (2015).
[32] E. Aurell, E. Domínguez, D. Machado, and R. Mulet, Phys. Rev. E **97**, 050103(R) (2018).
[33] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.123.230602 for a comparison of CME with FMS in the 2-SAT problem, a Finite Size Scaling analysis and with an Independent Spin Approximation.
[34] D. Gillespie, J. Comput. Phys. **22**, 403 (1976).
[35] B. Aspvall, M. Plass, and R. Tarjan, Inf. Proc. Lett. **8**, 121 (1979).
[36] S. Cook, *Proceedings of the 3rd ACM Symp. Theory Computing* (ACM New York, New York, 1971), p. 151.
[37] R. Mulet, A. Pagnani, M. Weigt, and R. Zecchina, Phys. Rev. Lett. **89**, 268701 (2002).
[38] J. Zhou and H. Zhou, Phys. Rev. E **79**, 020103(R) (2009).
[39] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, Phys. Rev. E **84**, 066106 (2011).
[40] C. Baldassi, C. Borgs, J. T. Chayes, A. Ingrosso, C. Lucibello, L. Saglietti, and R. Zecchina, Proc. Natl. Acad. Sci. U.S.A. **113**, E7655 (2016).
[41] L. Budzynski, F. Ricci-Tersenghi, and G. Semerjian, J. Stat. Mech. (2019) 023302.
[42] T. Barthel, C. De Bacco, and S. Franz, Phys. Rev. E **97**, 010104(R) (2018).