

A New Look at Survey Propagation and Its Generalizations

ELITZA MANEVA, ELCHANAN MOSSEL, AND MARTIN J. WAINWRIGHT

University of California—Berkeley, Berkeley, California

Abstract. This article provides a new conceptual perspective on *survey propagation*, which is an iterative algorithm recently introduced by the statistical physics community that is very effective in solving random k -SAT problems even with densities close to the satisfiability threshold. We first describe how any SAT formula can be associated with a novel family of Markov random fields (MRFs), parameterized by a real number $\rho \in [0, 1]$. We then show that applying belief propagation—a well-known “message-passing” technique for estimating marginal probabilities—to this family of MRFs recovers a known family of algorithms, ranging from pure survey propagation at one extreme ($\rho = 1$) to standard belief propagation on the uniform distribution over SAT assignments at the other extreme ($\rho = 0$). Configurations in these MRFs have a natural interpretation as partial satisfiability assignments, on which a partial order can be defined. We isolate *cores* as minimal elements in this partial ordering, which are also fixed points of survey propagation and the only assignments with positive probability in the MRF for $\rho = 1$. Our experimental results for $k = 3$ suggest that solutions of random formulas typically do not possess non-trivial cores. This makes it necessary to study the structure of the space of partial assignments for $\rho < 1$ and investigate the role of assignments that are very close to being cores. To that end, we investigate the associated lattice structure, and prove a weight-preserving identity that shows how any MRF with $\rho > 0$ can be viewed as a “smoothed” version of the uniform distribution over satisfying assignments ($\rho = 0$). Finally, we isolate properties of Gibbs sampling and message-passing algorithms that are typical for an ensemble of k -SAT problems.

Categories and Subject Descriptors: G. Mathematics of Computing [**Discrete Mathematics**]; Miscellaneous

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Satisfiability problems; k -SAT; survey propagation; belief propagation; sum-product; message-passing; factor graph; Markov random field; Gibbs sampling.

The work of E. Maneva was supported by NSF grants CCR-0121555 and DMS-0528488.

The work of E. Mossel was supported by a Miller Fellowship in Computer Science and Statistics, NSF grants DMS-0504245 and DMS-0528488, and a Sloan Fellowship in Mathematics.

The work of M. J. Wainwright was supported by a Sloan Fellowship in Computer Science and NSF grant DMS-0528488.

Authors' addresses: E. Maneva, IBM Almaden Research Center, San Jose, CA 95120, e-mail: elitza.maneva@gmail.com; M. J. Wainwright, Department of Statistics and Department of Electrical Engineering and Computer Science, UC Berkeley, Berkeley, CA 94720, e-mail: wainwrig@stat.berkeley.edu; E. Mossel, Department of Statistics, UC Berkeley, Berkeley, CA 94720, e-mail: mossel@stat.berkeley.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2007 ACM 0004-5411/2007/07-ART17 \$5.00 DOI 10.1145/1255443.1255445 <http://doi.acm.org/10.1145/1255443.1255445>

ACM Reference Format:

Maneva, E., Mossel, E., and Wainwright, M. J. 2007. A new look at survey propagation and its generalizations. *J. ACM* 54, 4, Article 17 (July 2007), 41 pages. DOI = 10.1145/1255443.1255445 <http://doi.acm.org/10.1145/1255443.1255445>

1. Introduction

Constraint satisfaction problems (CSPs) play an important role across a broad spectrum of computer science, including computational complexity theory [Cook 1971], coding theory [Gallager 1963; Richardson and Urbanke 2001], and artificial intelligence [Dechter 2003; Kumar 1992; Pearl 1988]. Important but challenging problems include devising efficient algorithms for finding satisfying assignments (when the problem is indeed satisfiable), or conversely providing a certificate of unsatisfiability. One of the best-known examples of a constraint satisfaction problem is the k -SAT problem, which is a classical NP complete problem [Cook 1971] for all $k \geq 3$. In trying to understand the origin of its hardness, a great deal of research has been devoted to the properties of formulas drawn from different probability distributions. One of the most natural models for random k -SAT problems is the following: for a fixed density parameter $\alpha > 0$, choose $m = \alpha n$ clauses uniformly and with replacement from the set of all k -clauses on n variables. Despite its simplicity, many essential properties of this model are yet to be understood: in particular, the hardness of deciding if a random formula is satisfiable and finding a satisfying assignment for a random formula are both major open problems [Levin 1986; Wang 1997; Feige 2002].

Given the significance of CSPs in various branches of computer science, there is now a considerable literature on algorithms for solving CSPs (see Kumar [1992], Dechter [2003], and Russell and Norvig [2003]). However, one of the most exciting recent developments in satisfiability problems has its origins *not* in computer science, but rather in statistical physics. More specifically, the ground-breaking contribution of Mézard et al. [2002], as described in an article published in “Science”, is the development of a new algorithm for solving k -SAT problems. A particularly dramatic feature of this method, known as *survey propagation*, is that it appears to remain effective at solving very large instances of random k -SAT problems—even with densities very close to the satisfiability threshold, a regime where other “local” algorithms (e.g., the WSAT method [Selman et al. 1996]) typically fail. Given this remarkable behavior, the survey propagation algorithm has generated a great deal of excitement and follow-up work in both the statistical physics and computer science communities (e.g., Braunstein et al. [2005, 2003], Braunstein and Zecchina [2004], Aurell et al. [2004], Achlioptas and Ricci-Tersenghi [2006], Mora et al. [2005], Parisi [2002], and Wainwright and Maneva [2005]). Nonetheless, despite the considerable progress to date, the reasons underlying the remarkable performance of survey propagation are not yet fully understood.

1.1. OUR CONTRIBUTIONS. This article provides a novel conceptual perspective on survey propagation—one that not only sheds light on the reasons underlying its success, but also places it within a broader framework of related “message-passing” algorithms that are widely used in different branches of computer science. More precisely, by introducing a new family of Markov random fields (MRFs) that are associated with any k -SAT problem, we show how a range of

algorithms—including survey propagation as a special case—can all be recovered as instances of the well-known belief propagation algorithm [Pearl 1988], as applied to suitably restricted MRFs within this family. This equivalence is important because belief propagation is a message-passing algorithm—widely used and studied in various areas, including coding theory [Gallager 1963; Richardson and Urbanke 2001; Kschischang 2001; Yedidia et al. 2005], computer vision [Freeman et al. 2000; Coughlan and Ferreira 2002] and artificial intelligence [Pearl 1988; Yedidia et al. 2003]—for computing approximations to marginal distributions in Markov random fields. Moreover, this equivalence motivates a deeper study of the combinatorial properties of the family of extended MRFs associated with survey propagation. Indeed, one of the main contributions of our work is to reveal the combinatorial structures underlying the survey propagation algorithm.

The configurations in our extended MRFs turn out to have a natural interpretation as particular types of *partial SAT assignments*, in which a subset of variables are assigned 0 or 1 variables in such a way that the remaining formula does not contain any empty or unit clauses. To provide some geometrical intuition for our results, it is convenient to picture these partial assignments as arranged in layers depending on the number of assigned variables, so that the top layer consists of fully assigned satisfying configurations. Figure 1 provides an idealized illustration of the space of partial assignments viewed in this manner. It is argued [Mézard et al. 2003; Mora et al. 2005; Achlioptas and Ricci-Tersenghi 2006] that for random formulas with high density of clauses, the set of fully assigned configurations are separated into disjoint clusters that cause local message-passing algorithms like belief propagation to break down (see Figure 2 for an illustration). Based on our results, the introduction of partial SAT assignments yields a *modified search space* that is far less fragmented, thereby permitting a local algorithm like belief propagation to find solutions.

We show that there is a natural partial ordering associated with this enlarged space, and we refer to minimal elements in this partial ordering as *cores*. We prove that any core is a fixed point of the pure form of survey propagation ($\rho = 1$). This fact indicates that each core represents a summary of one cluster of solutions. However, our experimental results for $k = 3$ indicate that the solutions of random formulas typically have trivial cores (i.e., the empty assignment). This observation motivates deeper study of the full family of Markov random fields for the range $0 \leq \rho \leq 1$, as well as the associated belief propagation algorithms, which we denote by $\text{SP}(\rho)$. Accordingly, we study the lattice structure of the partial assignments, and prove a combinatorial identity that reveals how the distribution for $\rho \in (0, 1)$ can be viewed as a “smoothed” version of the MRF with $\rho = 0$. Our experimental results on the $\text{SP}(\rho)$ algorithms indicate that they are most effective for values of ρ close to and not necessarily equal to 1. One intriguing possibility is that the effectiveness of pure survey propagation (i.e., $\text{SP}(1)$) may be a by-product of the fact that $\text{SP}(\rho)$ is most effective for values of ρ less than 1, but going to 1 as n goes to infinity. The near-core assignments which are the ones of maximum weight in this case, may correspond to quasi-solutions of the cavity equations, as defined by Parisi [2002]. In addition, we consider alternative sampling-based methods (e.g., Gibbs sampling) for computing marginals for the extended MRFs. We also study properties of both message-passing and Gibbs sampling that are typical over a random ensemble of k -SAT problems. We establish results that link the typical behavior of Gibbs sampling and message-passing algorithms under suitable initialization, and when applied to the extended family of MRFs with ρ sufficiently close to one.




FIG. 1. The set of fully assigned satisfying configurations occupy the top plane, and are arranged into clusters. Enlarging the space of partial assignments leads to a new space with better connectivity. Minimal elements in the partial ordering are known as cores. Each core corresponds to one or more clusters of solutions from the top plane. In this example, one of the clusters has as a core a non-trivial partial assignment, whereas the others are connected to the all-* assignment.




FIG. 2. The black dots represent satisfying assignments, and white dots unsatisfying assignments. Distance is to be interpreted as the Hamming distance between assignments. (a) For low densities the space of satisfying assignments is well connected. (b) As the density increases above α_d the space is believed to break up into an exponential number of clusters, each containing an exponential number of assignments. These clusters are separated by a “sea” of unsatisfying assignments. (c) Above α_c all assignments become unsatisfying.

The fact that the pure form of survey propagation (i.e., $\text{SP}(1)$ in our notation) is a form of belief propagation was first conjectured by Braunstein et al. [2005], and established independently of our work by Braunstein and Zecchina [2004]. In other independent work, Aurell et al. [2004] provided an alternative derivation of $\text{SP}(1)$ that established a link to belief propagation. However, both of these papers treat only the case $\rho = 1$, and do not provide a combinatorial interpretation based on an underlying Markov random field. The results established here provide the generalization to the full range of $\rho \in [0, 1]$. Moreover, the structures intrinsic to our Markov random fields—namely cores and lattices—highlight the importance of values $\rho \neq 1$, and place the survey propagation algorithm on a combinatorial ground. As we discuss later, this combinatorial perspective has already inspired subsequent work [Achlioptas and Ricci-Tersenghi 2006] on survey propagation for satisfiability problems. Looking forward, the methodology of partial assignments may also open the door to other problems where a complicated landscape prevents local search algorithms from finding good solutions. As a concrete example, a

subset of the current authors [Wainwright and Maneva 2005] have recently shown that related ideas can be leveraged to perform lossy data compression at near-optimal (Shannon limit) rates.

1.2. ORGANIZATION. The remainder of this article is organized in the following way:

- In Section 1.3, we provide further background on the k -SAT problem, as well as previous work on survey propagation.
- In Section 2, we introduce required notation and set up the problem more precisely.
- In Section 3, we define a family of Markov random fields (MRFs) over partial satisfiability assignments, and prove that survey propagation and related algorithms correspond to belief propagation on these MRFs.
- Section 4 is devoted to analysis of the combinatorial properties of this family of extended MRFs, as well as some experimental results on cores and Gibbs sampling.
- In Section 5, we consider properties of random ensembles of SAT formulae, and prove results that link the performance of survey propagation and Gibbs sampling to the choice of Markov random field.
- We conclude with a discussion in Section 6.

We note that many of the results reported here have been presented (without proofs or details) as an extended SODA abstract [Maneva et al. 2005].

1.3. PREVIOUS WORK ON k -SAT AND SURVEY PROPAGATION. As a classical NP complete problem [Cook 1971], the k -SAT problem for $k \geq 3$ has been extensively studied. One approach is to consider ensembles of random formulas; in particular, a commonly studied ensemble is based on choosing $m = \alpha n$ clauses uniformly and with replacement from the set of all k -clauses on n variables. Clearly, a formula drawn randomly from this ensemble becomes increasingly difficult to satisfy as the clause density $\alpha > 0$ increases. There is a large body of work [Friedgut 1999; Goerdts 1996; Chvatal and Reed 1992; de la Vega 1992; Dubois et al. 2000; Kaporis et al. 2000; Achlioptas and Peres 2003] devoted to the study of the *threshold* density where the formula becomes unsatisfiable; however, except for the case $k = 2$, the value of the threshold is currently unknown. However, non-rigorous techniques from statistical physics can be applied to yield estimates of the threshold; for instance, results from Monasson and Zecchina [1997] yield a threshold estimate of $\alpha_c \approx 4.267$ for $k = 3$.

The survey propagation (SP) algorithm, as introduced by Mézard et al. [2002], is an iterative message-passing technique that is able to find satisfying assignments for large instances of SAT problems at much higher densities than previous methods. The derivation of SP is based on the cavity method in conjunction with the 1-step replica symmetry breaking (1-RSB) ansatz of statistical physics. We do not go into these ideas in depth here, but refer the reader to the physics literature [Mézard and Zecchina 2002; Braunstein et al. 2005; Mézard et al. 2002] for further details. In brief, the main assumption is the existence of a critical value α_d for the density, smaller than the threshold density α_c , at which the structure of the space of solutions of a random formula changes. For densities below α_d the space of solutions is highly connected—in particular, it is possible to move from one solution to any other by

single variable flips,¹ staying at all times in a satisfying assignment. For densities above α_d , the space of solutions breaks up into clusters, so that moving from a SAT assignment within one cluster to some other assignment within another cluster requires flipping some constant fraction of the variables simultaneously. Figure 2 illustrates how the structure of the space of solutions evolves as the density of a random formula increases. The clustering phenomenon that is believed to occur in the second phase is known in the statistical physics literature as 1-step replica symmetry breaking [Mézard and Zecchina 2002], and the estimated value for α_d in the case $k = 3$ is $\alpha_d \approx 3.921$. Within each cluster, a distinction can be made between *frozen* variables—ones that do not change their value within the cluster—and *free* variables that do change their value in the cluster. A concise description of a cluster is an assignment of $\{0, 1, *\}$ to the variables with the frozen variables taking their frozen value, and the free variables taking the joker or wild card value $*$. The original argument for the clustering assumption was the analysis of simpler satisfiability problems, such as XOR-SAT, where the existence of clusters can be demonstrated by rigorous methods [Mézard et al. 2003]. In addition, if one assumes that there are no clusters, the cavity method calculation yields a value for $\alpha_c > 5$ (for $k = 3$), which is known to be wrong. More recently, Mora et al. [2005], as well as Achlioptas and Ricci-Tersenghi [2006] have demonstrated via rigorous methods that for $k \geq 8$ and some clause density below the unsatisfiability threshold, clusters of solutions do indeed exist.

The survey propagation (SP) algorithm is so-named, because like the belief propagation algorithm [Pearl 1988; Yedidia et al. 2003], it entails propagating statistical information in the form of messages between nodes in the graph. In the original derivation of the updates [Mézard et al. 2002; Braunstein et al. 2005], the messages are interpreted as “surveys” taken over the clusters in solution space, which provide information about the fraction of clusters in which a given variable is free or frozen. However, prior to the work presented here, it was not clear how to interpret the algorithm as an instantiation of belief propagation, and thus as a method for computing (approximations) to marginal distributions in a certain Markov random field (MRF). Moreover, as discussed above, our formulation of SP in this manner provides a broader view, in which SP is one of many possible message-passing algorithms that can be applied to smoothed MRF representations of SAT problems.

2. Background and Problem Set-Up

In this section, we begin with notation and terminology necessary to describe the k -SAT problem, and then provide a precise description of the survey propagation updates.

2.1. THE k -SAT PROBLEM AND FACTOR GRAPHS.

2.1.1. *Basic Notation.* Let C and V represent index sets for the clauses and variables, respectively, where $|V| = n$ and $|C| = m$. We denote elements of V using the letters i, j, k , etc., and members of C with the letters a, b, c , etc. We use

¹ There is no general agreement on whether assignments should be considered neighbors if they differ in only one variable, or any constant number of variables.




FIG. 3. Factor graph representation of a 3-SAT problem on $n = 5$ variables with $m = 4$ clauses, in which circular and square nodes correspond to variables and clauses respectively. Solid and dotted edges (a, i) correspond to the weightings $J_{a,i} = 0$ and $J_{a,i} = 1$ respectively. The clause a is defined by the neighborhood set $V(a) = \{1, 2, 3\}$ and weights $J_a = (0, 1, 1)$. In traditional notation, this corresponds to the formula $(x_1 \vee x_2 \vee x_3) \wedge (\bar{x}_1 \vee x_2 \vee x_3 \vee x_5) \wedge (\bar{x}_2 \vee x_4 \vee x_5)$.

x_5 to denote the subset of variables $\{x_i : i \in S\}$. In the k -SAT problem, the clause indexed by $a \in C$ is specified by the pair $(V(a), J_a)$, where $V(a) \subset V$ consists of k elements, and $J_a := (J_{a,i} : i \in V(a))$ is a k -tuple of $\{0, 1\}$ -valued weights. The clause indexed by a is *satisfied* by the assignment x if and only if $x_{V(a)} \neq J_a$. Equivalently, letting $\delta(y, z)$ denote an indicator function for the event $\{y = z\}$, if we define the function

$$\psi_a(x) := 1 - \prod_{i \in V(a)} \delta(J_{a,i}, x_i), \quad (1)$$

then the clause a is satisfied by x if and only if $\psi_a(x) = 1$. The overall formula consists of the AND of all the individual clauses, and is satisfied by x if and only if $\prod_{a \in C} \psi_a(x) = 1$.

2.1.2. Factor Graphs. A convenient graphical representation of any k -SAT problem is provided by the formalism of factor graphs (see Kschischang [2001] for further background). As illustrated in Figure 3, any instance of the k -SAT problem can be associated with a particular bipartite graph on the variables (denoted by circular nodes) and clauses (denoted by square nodes), where the edge (a, i) between the clause $a \in C$ and variable $i \in V$ is included in E if and only if $i \in V(a)$. Following Braunstein et al. [2005], it is convenient to introduce two labellings of any given edge—namely, solid or dotted, corresponding to whether $J_{a,i}$ is equal to 0 or 1 respectively.

For later use, we define (for each $i \in V$) the set $C(i) := \{a \in C : i \in V(a)\}$, corresponding to those clauses that impose constraints on variable x_i . This set of clauses can be decomposed into two disjoint subsets

$$C^-(i) := \{a \in C(i) : J_{a,i} = 1\}, \quad C^+(i) := \{a \in C(i) : J_{a,i} = 0\}, \quad (2)$$

according to whether the clause is satisfied by $x_i = 0$ or $x_i = 1$, respectively. Moreover, for each pair $(a, i) \in E$, the set $C(i) \setminus \{a\}$ can be divided into two (disjoint) subsets, depending on whether their preferred assignment of x_i agrees (in which case $b \in C_a^+(i)$) or disagrees (in which case $b \in C_a^-(i)$) with the preferred assignment of x_i corresponding to clause a . More formally, we define

$$C_a^+(i) := \{b \in C(i) \setminus \{a\} : J_{a,i} = J_{b,i}\}, \quad C_a^-(i) := \{b \in C(i) \setminus \{a\} : J_{a,i} \neq J_{b,i}\}. \quad (3)$$

Our focus is on random ensembles of k -SAT instances: for a given clause density $\alpha > 0$, a random instance is obtained by sampling $m = \alpha n$ clauses uniformly and

with replacement from the set of all k -clauses on n variables. In terms of the factor graph representation, this procedure samples a random (n, m) -bipartite graph, in which each clause $a \in C$ has degree k .

2.1.3. Markov Random Fields and Marginalization. The k -SAT problem can also be associated with a particular distribution defined as a Markov random field. Recall that a given instance of k -SAT can be specified by the collection of clause functions $\{\psi_a : a \in C\}$, as defined in equation (1). Using these functions, let us define a probability distribution over binary sequences via

$$p(x) := \frac{1}{Z} \prod_{a \in C} \psi_a(x), \quad (4)$$

where $Z := \sum_{x \in \{0,1\}^n} \prod_{a \in C} \psi_a(x)$ is the normalization constant. Note that this definition makes sense if and only if the k -SAT instance is satisfiable, in which case the distribution (4) is simply the uniform distribution over satisfying assignments.

This Markov random field representation (4) of any satisfiable formula motivates an approach for finding a satisfying assignment based on marginalization. In particular, suppose that we had an oracle that could compute exactly the marginal probability

$$p(x_i) = \sum_{\{x_j, j \in V \setminus \{i\}\}} p(x_1, x_2, \dots, x_n), \quad (5)$$

for a particular variable x_i . Note that this marginal reveals the existence of SAT configurations with $x_i = 0$ (if $p(x_i = 0) > 0$) or $x_i = 1$ (if $p(x_i = 1) > 0$). Therefore, a SAT configuration could be obtained by a recursive marginalization-decimation procedure, consisting of computing the marginal $p(x_i)$, appropriately setting x_i (i.e., decimating), and then reiterating the procedure on the modified Markov random field.

Of course, exact marginalization is computationally intractable in general [Cooper 1990; Dagum and Luby 1993], which motivates the use of efficient algorithms for approximate marginalization. An example of such an algorithm is what we will refer to as the “naive belief propagation algorithm”. The belief propagation (BP) algorithm, described in detail in Appendix A, can be applied to a MRF of the form (4) to estimate the marginal probabilities. Even though the BP algorithm is not exact, an intuitively reasonable approach is to set the variable that has the largest bias towards a particular value, and repeat. In fact, this marginalization-decimation approach based on naive BP finds a satisfying assignment for α up to approximately 3.9 for $k = 3$; for higher α , however, the iterations for BP typically fail to converge [Mézard et al. 2002; Aurrell et al. 2004; Braunstein et al. 2005].

2.2. SURVEY PROPAGATION. In contrast to the naive BP approach, a marginalization-decimation approach based on *survey propagation* appears to be effective in solving random k -SAT problems even close to the satisfiability threshold [Mézard et al. 2002; Braunstein et al. 2005]. Here we provide an explicit description of what we refer to as the $\text{SP}(\rho)$ family of algorithms, where setting the parameter $\rho = 1$ yields the pure form of survey propagation. For any given $\rho \in [0, 1]$, the algorithm involves updating messages from clauses to variables, as well as from variables to clauses. Each clause $a \in C$ passes a real number $\eta_{a \rightarrow i} \in [0, 1]$ to each of its variable neighbors $i \in V(a)$. In the other direction,

Message from clause a to variable i :

$$\eta_{a \rightarrow i} = \prod_{j \in V(a) \setminus \{i\}} \left[\frac{\Pi_{j \rightarrow a}^u}{\Pi_{j \rightarrow a}^u + \Pi_{j \rightarrow a}^*} \right]. \quad (6)$$

Message from variable i to clause a :

$$\Pi_{i \rightarrow a}^u = \left[1 - \rho \prod_{b \in C_a^u(i)} (1 - \eta_{b \rightarrow i}) \right] \prod_{b \in C_a^s(i)} (1 - \eta_{b \rightarrow i}). \quad (7a)$$

$$\Pi_{i \rightarrow a}^s = \left[1 - \prod_{b \in C_a^s(i)} (1 - \eta_{b \rightarrow i}) \right] \prod_{b \in C_a^u(i)} (1 - \eta_{b \rightarrow i}). \quad (7b)$$

$$\Pi_{i \rightarrow a}^* = \prod_{b \in C_a^s(i)} (1 - \eta_{b \rightarrow i}) \prod_{b \in C_a^u(i)} (1 - \eta_{b \rightarrow i}). \quad (7c)$$

FIG. 4. SP(ρ) updates.

each variable $i \in V$ passes a triplet of real numbers $\Pi_{i \rightarrow a} = (\Pi_{i \rightarrow a}^u, \Pi_{i \rightarrow a}^s, \Pi_{i \rightarrow a}^*)$ to each of its clause neighbors $a \in C(i)$. The precise form of the updates are given in Figure 4.

We pause to make a few comments about these SP(ρ) updates:

- (1) Although we have omitted the time step index for simplicity, Eqs. (6) and (7) should be interpreted as defining a recursion on (η, Π) . The initial values for η are chosen randomly in the interval $(0, 1)$.
- (2) The idea of the ρ parameter is to provide a smooth transition from the original naive belief propagation algorithm to the survey propagation algorithm. As shown in Braunstein et al. [2005], setting $\rho = 0$ yields the belief propagation updates applied to the probability distribution (4), whereas setting $\rho = 1$ yields the pure version of survey propagation.

2.2.1. Intuitive “Warning” Interpretation. To gain intuition for these updates, it is helpful to consider the pure SP setting of $\rho = 1$. As described by Braunstein et al. [2005], the messages in this case have a natural interpretation in terms of probabilities of warnings. In particular, at time $t = 0$, suppose that the clause a sends a warning message to variable i with probability $\eta_{a \rightarrow i}^0$, and a message without a warning with probability $1 - \eta_{a \rightarrow i}^0$. After receiving all messages from clauses in $C(i) \setminus \{a\}$, variable i sends a particular symbol to clause a saying either that it can’t satisfy it (“u”), that it can satisfy it (“s”), or that it is indifferent (“*”), depending on what messages it got from its other clauses. There are four cases:

- (1) If variable i receives warnings from $C_a^u(i)$ and no warnings from $C_a^s(i)$, then it cannot satisfy a and sends “u”.
- (2) If variable i receives warnings from $C_a^s(i)$ but no warnings from $C_a^u(i)$, then it sends an “s” to indicate that it is inclined to satisfy the clause a .
- (3) If variable i receives no warnings from either $C_a^u(i)$ or $C_a^s(i)$, then it is indifferent and sends “*”.
- (4) If variable i receives warnings from both $C_a^u(i)$ and $C_a^s(i)$, a contradiction has occurred.

The updates from clauses to variables are especially simple: in particular, any given clause sends a warning if and only if it receives “u” symbols from all of its other variables.

In this context, the real-valued messages involved in the pure SP(1) all have natural probabilistic interpretations. In particular, the message $\eta_{a \rightarrow i}$ corresponds to the probability that clause a sends a warning to variable i . The quantity $\Pi_{j \rightarrow a}^u$ can be interpreted as the probability that variable j sends the “u” symbol to clause a , and similarly for $\Pi_{j \rightarrow a}^s$ and $\Pi_{j \rightarrow a}^*$. The normalization by the sum $\Pi_{j \rightarrow a}^u + \Pi_{j \rightarrow a}^s + \Pi_{j \rightarrow a}^*$ reflects the fact that the fourth case is a failure, and hence is excluded a priori from the probability distribution.

Suppose that all of the possible warning events were independent. In this case, the SP message update Eqs. (6) and (7) would be the correct estimates for the probabilities. This independence assumption is valid on a graph without cycles, and in that case the SP updates do have a rigorous probabilistic interpretation. It is not clear if the equations have a simple interpretation in the case $\rho \neq 1$.

2.2.2. Decimation Based on Survey Propagation. Supposing that these survey propagation updates are applied and converge, the overall conviction of a value at a given variable is computed from the incoming set of equilibrium messages as

$$\begin{aligned} \mu_i(1) &\propto \left[1 - \rho \prod_{b \in C^+(i)} (1 - \eta_{b \rightarrow i}) \right] \prod_{b \in C^-(i)} (1 - \eta_{b \rightarrow i}), \\ \mu_i(0) &\propto \left[1 - \rho \prod_{b \in C^-(i)} (1 - \eta_{b \rightarrow i}) \right] \prod_{b \in C^+(i)} (1 - \eta_{b \rightarrow i}), \\ \mu_i(*) &\propto \prod_{b \in C^+(i)} (1 - \eta_{b \rightarrow i}) \prod_{b \in C^-(i)} (1 - \eta_{b \rightarrow i}). \end{aligned}$$

To be consistent with their interpretation as (approximate) marginals, the triplet $\{\mu_i(0), \mu_i(*), \mu_i(1)\}$ at each node $i \in V$ is normalized to sum to one. We define the *bias* of a variable node as $B(i) := |\mu_i(0) - \mu_i(1)|$.

The marginalization-decimation algorithm based on survey propagation [Braunstein et al. 2005] consists of the following steps:

- (1) Run SP(1) on the SAT problem. Extract the fraction β of variables with the largest biases, and set them to their preferred values.
- (2) Simplify the SAT formula, and return to Step (1).

Once the maximum bias over all variables falls below a pre-specified tolerance, the Walk-SAT algorithm is applied to the formula to find the remainder of the assignment (if possible). Intuitively, the goal of the initial phases of decimation is to find a cluster; once inside the cluster, the induced problem is considered easy to solve, meaning that any “local” algorithm should perform well within a given cluster.

3. Markov Random Fields Over Partial Assignments

In this section, we show how a large class of message-passing algorithms—including the SP(ρ) family as a particular case—can be recovered by applying

the well-known belief propagation algorithm to a novel class of Markov random fields (MRFs) associated with any k -SAT problem. We begin by introducing the notion of a partial assignment, and then use it to define the family of MRFs over these assignments.

3.1. PARTIAL ASSIGNMENTS. Suppose that the variables $x = \{x_1, \dots, x_n\}$ are allowed to take values in $\{0, 1, *\}$, which we refer to as a *partial assignment*. It will be convenient, when discussing the assignment of a variable x_i with respect to a particular clause a , to use the notation $s_{a,i} := 1 - J_{a,i}$ and $u_{a,i} := J_{a,i}$ to indicate, respectively, the values that are *satisfying* and *unsatisfying* for the clause a . With this set-up, we have the following:

Definition 1. A partial assignment x is *invalid* for a clause a if either

- (a) all variables are unsatisfying (i.e., $x_i = u_{a,i}$ for all $i \in V(a)$), or
- (b) all variables are unsatisfying except for exactly one index $j \in V(a)$, for which $x_j = *$.

Otherwise, the partial assignment x is valid for clause a , and we denote this event by $\text{VAL}_a(x_{V(a)})$. We say that a partial assignment is *valid* for a formula if it is valid for all of its clauses.

The motivation for deeming case (a) invalid is clear, in that any partial assignment that does not satisfy the clause must be excluded. Note that case (b) is also invalid, since (with all other variables unsatisfying) the variable x_j is effectively forced to $s_{a,i}$, and so cannot be assigned the $*$ symbol.

For a valid partial assignment, the subset of variables that are assigned either 0 or 1 values can be divided into *constrained* and *unconstrained* variables in the following way:

Definition 2. We say that a variable x_i is the *unique satisfying variable* for a clause if it is assigned $s_{a,i}$ whereas all other variables in the clause (i.e., the variables $\{x_j : j \in V(a) \setminus \{i\}\}$) are assigned $u_{a,j}$. A variable x_i is *constrained* by clause a if it is the unique satisfying variable.

We let $\text{CON}_{i,a}(x_{V(a)})$ denote an indicator function for the event that x_i is the unique satisfying variable in the partial assignment $x_{V(a)}$ for clause a . A variable is *unconstrained* if it has 0 or 1 value, and is not constrained by any clause. Thus for any partial assignment the variables are divided into stars, constrained and unconstrained variables. We define the three sets

$$\begin{aligned} S_*(x) &:= \{i \in V : x_i = *\} \\ S_c(x) &:= \{i \in V : x_i \text{ constrained}\} \\ S_o(x) &:= \{i \in V : x_i \text{ unconstrained}\} \end{aligned} \quad (8)$$

of $*$, constrained and unconstrained variables, respectively. Finally, we use $n_*(x)$, $n_c(x)$ and $n_o(x)$ to denote the respective sizes of these three sets.

Various probability distributions can be defined on valid partial assignments by giving different weights to stars, constrained and unconstrained variables, which we denote by ω_c, ω_* and ω_o respectively. Since only the ratio of the weights matters, we set $\omega_c = 1$, and treat ω_o and ω_* as free non-negative parameters (we generally take them in the interval $[0, 1]$). We define the weights of partial assignments in the following way: invalid assignments x have weight $W(x) = 0$, and for any valid

assignment x , we set

$$W(x) := (\omega_o)^{n_o(x)} \times (\omega_*)^{n_*(x)}. \quad (9)$$

Our primary interest is the probability distribution given by $p_W(x) \propto W(x)$. In contrast to the earlier distribution p , it is important to observe that this definition is valid for any SAT problem, whether or not it is satisfiable, as long as $\omega_* \neq 0$, since the all- $*$ vector is always a valid partial assignment. Note that, if $\omega_o = 1$ and $\omega_* = 0$, then the distribution $p_W(x)$ is the uniform distribution on satisfying assignments. Another interesting case that we will discuss is that of $\omega_o = 0$ and $\omega_* = 1$, which corresponds to the uniform distribution over valid partial assignments without unconstrained variables.

3.2. ASSOCIATED MARKOV RANDOM FIELDS. Given our set-up thus far, it is not at all obvious whether or not the distribution p_W can be decomposed as a Markov random field based on the original factor graph. Interestingly, we find that p_W does indeed have such a Markov representation for any choices of $\omega_o, \omega_* \in [0, 1]$. Obtaining this representation requires the addition of another dimension to our representation, which allows us to assess whether a given variable is constrained or unconstrained. We define the *parent set* of a given variable x_i , denoted by P_i , to be the set of clauses for which x_i is the unique satisfying variable. Immediate consequences of this definition are the following:

- (a) If $x_i = 0$, then we must have $P_i \subseteq C^-(i)$.
- (b) If $x_i = 1$, then there must hold $P_i \subseteq C^+(i)$.
- (c) The setting $x_i = *$ implies that $P_i = \emptyset$.

We will say that a pair (x_i, P_i) is valid if it satisfies conditions (a), (b) and (c). Note also that $P_i = \emptyset$ means that x_i cannot be constrained. For each $i \in V$, let $\mathcal{P}(i)$ be the set of all possible parent sets of variable i . Due to the restrictions imposed by our definition, P_i must be contained in either $C^+(i)$ or $C^-(i)$ but not both. Therefore, the cardinality² of $\mathcal{P}(i)$ is $2^{|C^-(i)|} + 2^{|C^+(i)|} - 1$.

Our extended Markov random field is defined on the Cartesian product space $\mathcal{X}_1 \times \dots \times \mathcal{X}_n$, where $\mathcal{X}_i := \{0, 1, *\} \times \mathcal{P}(i)$. The distribution factorizes as a product of compatibility functions at the variable and clause nodes of the factor graph, which are defined as follows:

3.2.1. Variable Compatibilities. Each variable node $i \in V$ has an associated compatibility function of the form:

$$\Psi_i(x_i, P_i) := \begin{cases} \omega_o & : P_i = \emptyset, x_i \neq * \\ \omega_* & : P_i = \emptyset, x_i = * \\ 1 & : \text{for any other valid } (x_i, P_i). \end{cases} \quad (10)$$

The role of these functions is to assign weight to the partial assignments according to the number of unconstrained and star variables, as in the weighted distribution p_W .

3.2.2. Clause Compatibilities. The compatibility functions at the clause nodes serve to ensure that only valid assignments have non-zero probability, and that the parent sets $P_{V(a)} := \{P_i : i \in V(a)\}$ are consistent with the assignments

²Note that it is necessary to subtract one so as not to count the empty set twice.

$x_{V(a)} := \{x_i : i \in V(a)\}$ in the neighborhood of a . More precisely, we require that the partial assignment $x_{V(a)}$ is valid for a (denoted by $\text{VAL}_a(x_{V(a)}) = 1$) and that for each $i \in V(a)$, exactly one of the two following conditions holds:

- (a) $a \in P_i$ and x_i is constrained by a or
- (b) $a \notin P_i$ and x_i is not constrained by a .

The following compatibility function corresponds to an indicator function for the intersection of these events:

$$\Psi_a(x_{V(a)}, P_{V(a)}) := \text{VAL}_a(x_{V(a)}) \times \prod_{i \in V(a)} \delta(\text{Incl}[a \in P_i], \text{CON}_{a,i}(x_{V(a)})). \quad (11)$$

We now form a Markov random field over partial assignments and parent sets by taking the product of variable (10) and clause (11) compatibility functions

$$p_{\text{gen}}(x, P) \propto \prod_{i \in V} \Psi_i(x_i, P_i) \prod_{a \in C} \Psi_a(x_{V(a)}, P_{V(a)}). \quad (12)$$

With these definitions, some straightforward calculations show that $p_{\text{gen}} = p_W$.

3.3. SURVEY PROPAGATION AS AN INSTANCE OF BELIEF PROPAGATION. We now consider the form of the belief propagation (BP) updates as applied to the MRF p_{gen} defined by Eq. (12). We refer the reader to Appendix A for the definition of the BP algorithm on a general factor graph. The main result of this section is to establish that the $\text{SP}(\rho)$ family of algorithms are equivalent to belief propagation as applied to p_{gen} with suitable choices of the weights ω_o and ω_* . In the interests of readability, most of the technical lemmas will be presented in the appendix.

We begin by introducing some notation necessary to describe the BP updates on the extended MRF. The BP message from clause a to variable i , denoted by $M_{a \rightarrow i}(\cdot)$, is a vector of length $|\mathcal{X}_i| = 3 \times |\mathcal{P}(i)|$. Fortunately, due to symmetries in the variable and clause compatibilities defined in Eqs. (10) and (11), it turns out that the clause-to-variable message can be parameterized by only three numbers, $\{M_{a \rightarrow i}^u, M_{a \rightarrow i}^s, M_{a \rightarrow i}^*\}$, as follows:

$$M_{a \rightarrow i}(x_i, P_i) = \begin{cases} M_{a \rightarrow i}^s & \text{if } x_i = s_{a,i}, P_i = S \cup \{a\} \text{ for some } S \subseteq C_a^s(i), \\ M_{a \rightarrow i}^u & \text{if } x_i = u_{a,i}, P_i \subseteq C_a^u(i), \\ M_{a \rightarrow i}^* & \text{if } x_i = s_{a,i}, P_i \subseteq C_a^s(i) \text{ or } x_i = *, P_i = \emptyset, \\ 0 & \text{otherwise,} \end{cases} \quad (13)$$

where $M_{a \rightarrow i}^s, M_{a \rightarrow i}^u$ and $M_{a \rightarrow i}^*$ are elements of $[0, 1]$.

Now turning to messages from variables to clauses, it is convenient to introduce the notation $P_i = S \cup \{a\}$ as a shorthand for the event

$$a \in P_i \text{ and } S = P_i \setminus \{a\} \subseteq C_a^s(i),$$

where it is understood that S could be empty. In Appendix B, we show that the variable-to-clause message $M_{i \rightarrow a}$ is fully specified by values for pairs (x_i, P_i) of

six general types:

$$\{(s_{a,i}, S \cup \{a\}), (s_{a,i}, \emptyset \neq P_i \subseteq C_a^s(i)), \\ (u_{a,i}, \emptyset \neq P_i \subseteq C_a^u(i)), (s_{a,i}, \emptyset), (u_{a,i}, \emptyset), (*, \emptyset)\}.$$

The BP updates themselves are most compactly expressed in terms of particular linear combinations of such basic messages, defined in the following way:

$$R_{i \rightarrow a}^s := \sum_{S \subseteq C_a^s(i)} M_{i \rightarrow a}(s_{a,i}, S \cup \{a\}) \quad (14)$$

$$R_{i \rightarrow a}^u := \sum_{P_i \subseteq C_a^u(i)} M_{i \rightarrow a}(u_{a,i}, P_i) \quad (15)$$

$$R_{i \rightarrow a}^* := \sum_{P_i \subseteq C_a^s(i)} M_{i \rightarrow a}(s_{a,i}, P_i) + M_{i \rightarrow a}(*, \emptyset). \quad (16)$$

Note that $R_{i \rightarrow a}^s$ is associated with the event that x_i is the unique satisfying variable for clause a ; $R_{i \rightarrow a}^u$ with the event that x_i does not satisfy a ; and $R_{i \rightarrow a}^*$ with the event that x_i is neither unsatisfying nor uniquely satisfying (i.e., either $x_i = *$, or $x_i = s_{a,i}$ but is not the only variable that satisfies a).

With this terminology, the BP algorithm on the extended MRF can be expressed in terms of the following recursions on the triplets $(M_{a \rightarrow i}^s, M_{a \rightarrow i}^u, M_{a \rightarrow i}^*)$ and $(R_{i \rightarrow a}^s, R_{i \rightarrow a}^u, R_{i \rightarrow a}^*)$:

BP updates on extended MRF:

Messages from clause a to variable i

$$M_{a \rightarrow i}^s = \prod_{j \in V(a) \setminus \{i\}} R_{j \rightarrow a}^u \\ M_{a \rightarrow i}^u = \prod_{j \in V(a) \setminus \{i\}} (R_{j \rightarrow a}^u + R_{j \rightarrow a}^*) + \sum_{k \in V(a) \setminus \{i\}} (R_{k \rightarrow a}^s - R_{k \rightarrow a}^*) \prod_{j \in V(a) \setminus \{i,k\}} R_{j \rightarrow a}^u - \prod_{j \in V(a) \setminus \{i\}} R_{j \rightarrow a}^u \\ M_{a \rightarrow i}^* = \prod_{j \in V(a) \setminus \{i\}} (R_{j \rightarrow a}^u + R_{j \rightarrow a}^*) - \prod_{j \in V(a) \setminus \{i\}} R_{j \rightarrow a}^u.$$

Messages from variable i to clause a :

$$R_{i \rightarrow a}^s = \prod_{b \in C_a^s(i)} M_{b \rightarrow i}^u \left[\prod_{b \in C_a^s(i)} (M_{b \rightarrow i}^s + M_{b \rightarrow i}^*) \right] \\ R_{i \rightarrow a}^u = \prod_{b \in C_a^s(i)} M_{b \rightarrow i}^u \left[\prod_{b \in C_a^s(i)} (M_{b \rightarrow i}^s + M_{b \rightarrow i}^*) - (1 - \omega_o) \prod_{b \in C_a^s(i)} M_{b \rightarrow i}^* \right] \\ R_{i \rightarrow a}^* = \prod_{b \in C_a^s(i)} M_{b \rightarrow i}^u \left[\prod_{b \in C_a^s(i)} (M_{b \rightarrow i}^s + M_{b \rightarrow i}^*) - (1 - \omega_o) \prod_{b \in C_a^s(i)} M_{b \rightarrow i}^* \right] + \omega_* \prod_{b \in C_a^s(i) \cup C_a^u(i)} M_{b \rightarrow i}^*.$$

We provide a detailed derivation of these BP equations on the extended MRF in Appendix B. Since the messages are interpreted as probabilities, we only need their ratio, and we can normalize them to any constant. At any iteration, approximations to the local marginals at each variable node $i \in V$ are given by (up to a normalization

constant):

$$\begin{aligned} F_i(0) &\propto \prod_{b \in C^+(i)} M_{b \rightarrow i}^u \left[\prod_{b \in C^-(i)} (M_{b \rightarrow i}^s + M_{b \rightarrow i}^*) - (1 - \omega_o) \prod_{b \in C^-(i)} M_{b \rightarrow i}^* \right] \\ F_i(1) &\propto \prod_{b \in C^-(i)} M_{b \rightarrow i}^u \left[\prod_{b \in C^+(i)} (M_{b \rightarrow i}^s + M_{b \rightarrow i}^*) - (1 - \omega_o) \prod_{b \in C^+(i)} M_{b \rightarrow i}^* \right] \\ F_i(*) &\propto \omega_* \prod_{b \in C(i)} M_{b \rightarrow i}^* \end{aligned}$$

We demonstrate their derivation also in Appendix B.

The following theorem establishes that the $\text{SP}(\rho)$ family of algorithms is equivalent to belief propagation on the extended MRF:

THEOREM 2. *For all $\omega_* \in [0, 1]$, the BP updates on the extended (ω_*, ω_o) -MRF p_{gen} are equivalent to the $\text{SP}(\omega_*)$ family of algorithms under the following restrictions:*

- (a) *the constraint $\omega_o + \omega_* = 1$ is imposed, and*
- (b) *all messages are initialized such that $M_{a \rightarrow i}^u = M_{a \rightarrow i}^*$ for every edge (a, i) .*

PROOF. Under the constraint $\omega_o + \omega_* = 1$, if we initialize $M_{a \rightarrow i}^u = M_{a \rightarrow i}^*$ on every edge, then there holds $R_{i \rightarrow a}^s = R_{i \rightarrow a}^*$ and consequently $M_{a \rightarrow i}^u = M_{a \rightarrow i}^*$ remains true at the next iteration. Initializing the parameters in this way and imposing the normalization $M_{a \rightarrow i}^s + M_{a \rightarrow i}^* = 1$ leads to the following recurrence equations

$$M_{a \rightarrow i}^s = \frac{\prod_{j \in V(\omega) \setminus \{i\}} R_{j \rightarrow a}^u}{\prod_{j \in V(\omega) \setminus \{i\}} (R_{j \rightarrow a}^s + R_{j \rightarrow a}^u)},$$

where

$$\begin{aligned} R_{i \rightarrow a}^u &= \prod_{b \in C_a^+(i)} (1 - M_{b \rightarrow i}^s) \left[1 - \omega_* \prod_{b \in C_a^+(i)} (1 - M_{b \rightarrow i}^s) \right], \text{ and} \\ R_{i \rightarrow a}^* &= \prod_{b \in C_a^+(i)} (1 - M_{b \rightarrow i}^*). \end{aligned}$$

These updates are equivalent to $\text{SP}(\omega_*)$ by setting $\eta_{a \rightarrow i} = M_{a \rightarrow i}^s$, $\Pi_{i \rightarrow a}^u = R_{i \rightarrow a}^u$, and $\Pi_{i \rightarrow a}^s + \Pi_{i \rightarrow a}^* = R_{i \rightarrow a}^*$. \square

3.3.1. Remarks.

- (1) Theorem 2 is a generalization of the result of Braunstein and Zecchina [2004], who showed that $\text{SP}(1)$ is equivalent to belief propagation on a certain MRF.
- (2) The essence of Theorem 3 is that the pure survey propagation algorithm, as well as all the ρ -variants thereof, are all equivalent to belief propagation on our extended MRF with suitable parameter choices. This equivalence is important for a number of reasons:

- (a) Belief propagation is a widely-used algorithm for computing approximations to marginal distributions in general Markov random fields [Yedidia et al. 2003; Kschischang 2001]. It also has a variational interpretation as an

iterative method for attempting to solve a non-convex optimization problem based on the Bethe approximation [Yedidia et al. 2003]. Among other consequences, this variational interpretation leads to other algorithms that also solve the Bethe problem, but unlike belief propagation, are guaranteed to converge [Welling and Teh 2001; Yuille 2002; Wainwright and Jordan 2003].

- (b) Given the link between SP and extended MRFs, it is natural to study combinatorial and probabilistic properties of the latter objects. In Section 4, we show how so-called ‘‘cores’’ arise as fixed points of $\text{SP}(1)$, and we prove a weight-preserving identity that shows how the extended MRF for general ρ is a ‘‘smoothed’’ version of the naive MRF.
- (c) Finally, since BP (and hence SP) is computing approximate marginals for the MRF, it is natural to study other ways of computing marginals and examine if these lead to an effective way for solving random k -SAT problems. We begin this study in Section 4.5.

- (3) The initial messages have very small influence on the behavior of the algorithm, and they are typically chosen to be uniform random variables in $(0, 1)$. In practice, for $\omega_o + \omega_* = 1$ if we start with different values for $M_{a \rightarrow i}^u$ and $M_{a \rightarrow i}^*$ they soon converge to become equal.

- (4) If we restrict our attention to 3-SAT, the equations have simpler form. In particular for a clause a on x_i, x_j, x_k , the messages to variable node i are:

$$\begin{aligned} M_{a \rightarrow i}^s &= R_{j \rightarrow a}^u R_{k \rightarrow a}^u \\ M_{a \rightarrow i}^u &= R_{j \rightarrow a}^* R_{k \rightarrow a}^* + R_{j \rightarrow a}^s R_{k \rightarrow a}^u + R_{j \rightarrow a}^u R_{k \rightarrow a}^* \\ M_{a \rightarrow i}^* &= R_{j \rightarrow a}^* R_{k \rightarrow a}^* + R_{j \rightarrow a}^* R_{k \rightarrow a}^u + R_{j \rightarrow a}^u R_{k \rightarrow a}^* \end{aligned}$$

4. Combinatorial Properties

This section is devoted to an investigation of the combinatorial properties associated with the family of extended Markov random fields defined in the previous section. We begin by defining an acyclic directed graph on all valid partial assignments. Of particular interest are the minimal elements in the resulting partial ordering. We refer to these as *cores*.

4.1. DIRECTED GRAPH AND PARTIAL ORDERING. The vertex set of the directed graph G consists of all valid partial assignments. The edge set is defined in the following way: for a given pair of valid partial assignments x and y , the graph includes a directed edge from x to y if there exists an index $i \in V$ such that (i) $x_j = y_j$ for all $j \neq i$; and (ii) $y_i = *$ and $x_i \neq y_i$. We label the edge between x and y with the index i , corresponding to the fact that y is obtained from x by adding one extra $*$ in position i .

This directed graph G has a number of properties:

- (a) Valid partial assignments can be separated into different levels based on their number of star variables. In particular, assignment x is in level $n_*(x)$. Thus, every edge is from an assignment in level $l - 1$ to one in level l , where l is at most n .
- (b) The out-degree of any valid partial assignment x is exactly equal to its number of unconstrained variables $n_o(x)$.

(c) It is an acyclic graph so that its structure defines a partial ordering; in particular, we write $y < x$ if there is a directed path in G from x to y . Notice that all directed paths from x to y are labeled by indices in the set $T = \{i \in V : x_i \neq y_i = *\}$, and only the order in which they appear is different.

Given the partial ordering defined by G , it is natural to consider elements that are minimal in this partial ordering. For any valid partial assignment x and a subset $S \subseteq V$, let $\gamma_S(x)$ be the minimal $y < x$, such that the path from x to y is labeled only by indices in S . In particular, $\gamma_V(x)$ is a minimal assignment in the order. It is easy to show that there always exists a unique $\gamma_S(x)$.

PROPOSITION 4. *For any valid assignment x and $S \subseteq V$, there is a unique minimal $y < x$ such that the path from x to y is labeled only by indices in S . Furthermore, $S_o(y) \cap S = \emptyset$ and $S_*(y) = S_*(x) \cup T$, where $T \subseteq S$ is the set of labels on any path from x to y .*

PROOF. To prove the second assertion in the proposition statement for a minimal y , suppose there exists $i \in S \cap S_o(y)$. Then there must be an outgoing edge from y labeled by an element in S , which contradicts the assumed minimality of y . The equivalence $S_*(y) = S_*(x) \cup T$ follows directly from the definition of G and its edge labels.

To establish the uniqueness statement, suppose that there are two minimal such assignments y_1 and y_2 , and the paths from x to y_1 and y_2 are labeled by sets of indices $T_1, T_2 \subseteq S$ respectively. If $T_1 = T_2$, then $y_1 = y_2$, so let us assume that T_1 and T_2 are distinct. Without loss of generality, we may take $T_1 \setminus T_2 \neq \emptyset$. Consider a particular path from x to y_1 , with labels t_1, t_2, \dots, t_r , where $r = |T_1|$. Let t_i be the first label such that $t_i \notin T_2$. Then its corresponding variable is unconstrained when the variables indexed by $\{t_1, \dots, t_{i-1}\} \cup S_*(x) \subseteq T_2 \cup S_*(x)$ are assigned $*$, therefore it is unconstrained in y_2 . This implies that there exists an edge out of y_2 that is labeled by $t_i \in S$, which contradicts the assumption that y_2 is minimal. \square

We define a *core assignment* to be a valid partial assignment $y \in \{0, 1, *\}^n$ that contains no unconstrained variables. We say that a core assignment y is *non-trivial* if $n_*(y) < n$, so that it has at least one constrained $\{0, 1\}$ variable. Under this definition, it follows that for any partial assignment x , the associated minimal element $\gamma_V(x)$ is a core assignment.

It is easy to see that assignments that belong to the same connected component in the graph of solutions (in which two solutions are defined to be connected if they differ in exactly one variable) have the same core. This follows from the fact that neighboring solutions that differ in variable i are both above the partial assignment that differs from them in x_i where it is $*$. The converse statement is not true, that is, solutions from different connected components of the solution graph may have the same core.

Given a valid ordinary assignment $z \in \{0, 1\}^n$, an interesting object is the subgraph of partial assignments that lie below it in the partial ordering. It can be seen that any pair of elements in this subgraph have both a unique maximal element and a unique minimal element, so that any such subgraph is a lattice [Stanley 1997].

In examples shown in Figure 5, only a subset of the partial assignments is shown, since even for small formulas the space of partial assignments is quite large. For the first formula all satisfying assignments have a trivial core. For the second one, on the other hand, there are assignments with nontrivial cores.




FIG. 5. Portion of the directed graph on partial assignments for two different formulas: (a) $(\bar{x}_1 \vee \bar{x}_2 \vee x_3) \wedge (x_2 \vee x_3 \vee \bar{x}_4)$. Highlighted in the lattice below the satisfying assignment $z = (1, 1, 1, 1)$, whose core is trivial (i.e., $\gamma_V(z) = (*, *, *, *)$). (b) $(\bar{x}_1 \vee x_2 \vee x_3) \wedge (x_1 \vee \bar{x}_2 \vee x_3) \wedge (x_2 \vee \bar{x}_3 \vee x_4) \wedge (x_2 \vee x_3 \vee x_5) \wedge (x_1 \vee x_5 \vee \bar{x}_4)$. The satisfying assignment $z = (0, 0, 0, 0, 1)$ has the nontrivial core $\gamma_V(z) = (0, 0, 0, *, *)$. For the same formula, there are other satisfying assignments, for example, $(1, 0, 1, 0, 1)$ which have a trivial core.

4.2. PURE SURVEY PROPAGATION AS A PEELING PROCEDURE. As a particular case of Theorem 2, setting $\omega_* = 1$ and $\omega_o = 0$ yields the extended MRF that underlies the SP(1) algorithm. In this case, the only valid assignments with positive weight are those without any unconstrained variables—namely, core assignments. Thus, the distribution p_W for $(\omega_o, \omega_*) = (0, 1)$ is simply uniform over the core assignments. The following result connects fixed points of SP(1) to core assignments:

PROPOSITION 4. *For a valid assignment x , let $SP(1)$ be initialized by*

$$\Pi_{i \rightarrow a}^\mu = \delta(x_i, u_{a,i}), \quad \Pi_{i \rightarrow a}^\nu = \delta(x_i, s_{a,i}), \quad \Pi_{i \rightarrow a}^* = 0.$$

Then, within a finite number of steps, the algorithm converges and the output fields are

$$\mu_i(b) = \delta(y_i, b),$$

where $y = \gamma_V(x)$ and $b \in \{0, 1, *\}$.

PROOF. We say that a variable i belongs to the core if $y_i \neq *$. We say that a clause a belongs to the core if all the variables in the clause belong to the core. We first show by induction that

- I. If a and i belong to the core and y_i is not the unique satisfying variable for a , then $\Pi_{i \rightarrow a}^u = \delta(x_i, u_{a,i})$ and $\Pi_{i \rightarrow a}^s = \delta(x_i, s_{a,i})$, and
- II. If a and i belong to the core and y_i is the unique satisfying variable for a , then $\eta_{a \rightarrow i} = 1$.

Clearly, property I holds at time 0. Therefore, it suffices to prove that if property I holds at time t , then so does II and that if property II holds at time t , then property I holds at time $t + 1$.

Suppose that property I holds at time t . Let a and i belong to the core such that y_i is the unique satisfying variable of the clause a . By the induction hypothesis for all $j \in V(a) \setminus \{i\}$ it holds that $\Pi_{j \rightarrow a}^u = \delta(x_j, u_{a,j}) = 1$. This implies that $\eta_{a \rightarrow i} = 1$ as needed.

Suppose that property II holds at time t . Let a and i belong to the core such that y_i is not unique satisfying for a . By the assumption, it follows that there exists b which belongs to the core such that y_i is the unique satisfying variable for b . This implies by the induction hypothesis that $\eta_{b \rightarrow i} = 1$. It is now easy to see that at update $t + 1$: $\Pi_{i \rightarrow a}^u = \delta(x_i, u_{a,i})$ and $\Pi_{i \rightarrow a}^s = \delta(x_i, s_{a,i})$. Note that the claim above implies that for all times t and all i such that $y_i \neq *$, it holds that $\mu_i(b) = \delta(y_i, b)$. Let i_1, i_2, \dots, i_s be a “peeling-path” from x to y . In other words, the variable i_1 is not uniquely satisfying any clause. Once, this variable is set to $*$, the variable i_2 is not uniquely satisfying any clause etc. We claim that for all $1 \leq t \leq s$, for all updates after time t and for all clauses a such that $i_t \in V(a)$ it holds that $\eta_{a \rightarrow i_t} = 0$. The proof follows easily by induction on t . This in turn implies that if for all updates after time t $\mu_{i_t}(b) = \delta(y_i, *)$, from which the result follows. \square

Thus, $\text{SP}(1)$, when suitably initialized, simply strips the valid assignment x down to its core $\gamma_V(x)$. Moreover, Proposition 5, in conjunction with Theorem 3, leads to viewing the pure form of survey propagation $\text{SP}(1)$ as performing an approximate marginalization over cores. Therefore, our results raise the crucial question: do cores exist for random formulas? Motivated by this perspective, Achlioptas and Ricci-Tersenghi [2006] have answered this question affirmatively for k -SAT with $k \geq 9$. In Section 5, we show that cores, if they exist, must be “large” in a suitable sense (see Proposition 8). In the following section, we explore the case $k = 3$ via experiments on large instances.

4.3. PEELING EXPERIMENTS. We have performed a large number of the following experiments:

- (1) starting with a satisfying assignment x , change a random one of its unconstrained variables to $*$,
- (2) repeat until there are no unconstrained variables.




FIG. 6. Evolution of the number of unconstrained variables in the peeling process: start with a satisfying assignment, change a random unconstrained variable to $*$ and repeat. Plotted is the result of an experiment for $n = 100,000$, for random formulas with $k = 3$ and $\alpha = \{2, 2.5, 3, 3.5, 4, 4.1, 4.2\}$. In particular, core assignments are on the x -axis, and satisfying assignments are on the y -axis.

This procedure, which we refer to as “peeling”, is equivalent to taking a random path from x in G , by choosing at each step a random outgoing edge. Any such path terminates at the core $\gamma_V(x)$. It is interesting to examine at each step of this process the number of unconstrained variables (equivalently, the number of outgoing edges in the graph G). For $k = 3$ SAT problems, Figure 6 shows the results of such experiments for $n = 100,000$, and using different values of α . The plotted curves are the evolution of the number of unconstrained variables as the number of $*$'s increases. We note that for $n = 100$ and α close to the threshold, satisfying assignments often correspond to core assignments; a similar observation was also made by Braunstein and Zecchina [2004]. In contrast, for larger n , this correspondence is rarely the case. Rather, the generated curves suggest that $\gamma_V(x)$ is almost always the all- $*$ assignment, and moreover that for high density α , there is a critical level in G where the out-degrees are very low. Increasing α results in failure of the algorithm itself, rather than in the formation of real core assignments.

For $k = 2$, the event that there is a path in G from a satisfying assignment to the all- $*$ assignment has a very natural interpretation. In particular, it is equivalent to the event that the *pure-literal rule* succeeds in finding an assignment. The pure-literal rule [Rosenthal et al. 1999] is an algorithm consisting of the following steps: assign 1 to a variable if it only appears positively in a clause, and 0 if it only appears negatively in a clause, reduce the formula, and repeat the procedure. It is straightforward to check that the sequence of variables given by the labels on any path from the all- $*$ assignment to a satisfying assignment can be identified with a sequence of steps of the pure-literal type. Furthermore, it is known [Rosenthal et al. 1999] that there is a phase transition for the event that the pure-literal rule succeeds at $\alpha = 1$.

Interestingly, as mentioned earlier, for $k \geq 9$, it has been shown [Achlioptas and Ricci-Tersenghi 2006] that there exist values $\alpha < \alpha_c$ such that this peeling procedure provably results in a nontrivial core assignment with high probability. The fact that we do not observe core assignments for $k = 3$, and yet the algorithm




FIG. 7. Performance of BP for different choices of (ω_0, ω_*) as applied to a particular randomly chosen formula with $n = 10000$, $k = 3$, $\alpha = 4.2$. Four distinct cases can be distinguished: (i) BP converges and the decimation steps yields a complete solution, (ii) BP converges and the decimation steps yield a partial solution, completed by using Walk-SAT, (iii) BP converges, but the decimation steps don't lead to a solution, and (iv) BP does not converge.

is successful, means that an alternative explanation is required. Accordingly, we propose studying the behavior of $SP(\rho)$ for $\rho \in (0, 1)$. Our experimental results, consistent with similar reports from S. Kirkpatrick (personal communication, 2004), show that $SP(\rho)$ tends to be most effective in solving k -SAT for values of $\rho < 1$. If so, the good behavior of $SP(1)$ may well follow from the similarity of $SP(1)$ updates to $SP(\rho)$ updates for $\rho \approx 1$. To further explore this issue, the effects of varying the weight distribution (ω_0, ω_*) , and consequently the parameter ρ , are discussed in the following section.

4.4. WEIGHT DISTRIBUTION AND SMOOTHING. One of the benefits of our analysis is that it suggests a large pool of algorithms to be investigated. One option is to vary the values of ω_0 and ω_* . A “good” setting of these parameters should place significant weight on precisely those valid assignments that can be extended to satisfying assignments. At the same time, the parameter setting clearly affects the level of connectivity in the space of valid assignments. Connectivity most likely affects the performance of belief propagation, as well as any other algorithm that we may apply to compute marginals or sample from the distribution.

Figure 7 shows the performance of belief propagation on the extended MRF for different values of (ω_0, ω_*) , and applied to particular random formula with $n = 10,000$, $k = 3$ and $\alpha = 4.2$. The most successful pairs in this case were $(0.05, 0.95)$, $(0.05, 0.9)$, $(0.05, 0.85)$, and $(0.05, 0.8)$. For these settings of the parameters, the decimation steps reached a solution, so a call to WalkSAT was not needed. For weights satisfying $\omega_0 + \omega_* > 1$, the behavior is very predictable: even when the algorithm converges, the choices that it makes in the decimation steps lead to a contradiction. Note that there is a sharp transition in algorithmic behavior as the weights cross the line $\omega_0 + \omega_* = 1$, which is representative of the more general behavior.

The following result provides some justification for the excellent performance in the regime $\omega_0 + \omega_* \leq 1$.

THEOREM 6. *If $\omega_0 + \omega_* = 1$, then $\sum_{y \leq x} W(y) = \omega_*^{n_A(x)}$ for any valid assignment x . If $\omega_0 + \omega_* < 1$, then $\sum_{y \leq x} W(y) \geq (\omega_*)^{n_A(x)}$ for any valid assignment x .*

It should be noted that Theorem 6 has a very natural interpretation in terms of a “smoothing” operation. In particular, the (ω_0, ω_*) -MRF may be regarded as a smoothed version of the uniform distribution over satisfying assignments, in which the uniform weight assigned to each satisfying assignment is spread over the lattice associated with it.³

The remainder of this section is devoted to the proof of Theorem 6.

PROOF. We start with the case $\omega_0 + \omega_* = 1$. Let A denote the set of partial assignments z such that $z_j \in \{x_j, *\}$ for all $j \in V$. We refer to these as the set of assignments consistent with x . Let $B = \{y : y \leq x\}$ be the set of valid assignments that are reachable from x . Notice that all $y \in B$ are valid and consistent with x , but not every valid assignment in A is reachable from x . We will let $S_*(z)$ denote the set of variables assigned $*$ both for valid and invalid assignments z .

We define a map between all assignments consistent with x and the set of reachable ones. Let $\sigma : A \rightarrow B$ be defined as

$$\sigma(z) := \gamma_{S_*(z)}(x).$$

(See Figure 8 for an illustration.) Notice that if $y \in B$ then $\sigma(y) = y$. The map is, of course, many-to-one. We define what we'll show is the reverse map. For $y \in B$, let

$$\tau(y) := \{z \in A : S_*(z) = S_*(y) \cup T, T \subseteq S_c(y)\}.$$

LEMMA 7. *For any $y \in B$ and $z \in A$, $z \in \tau(y)$ if and only if $\sigma(z) = y$.*

PROOF. Let $z \in \tau(y)$ so that $S_*(z) = S_*(y) \cup T$ for some $T \subseteq S_c(y)$. $\sigma(z) = \gamma_{S_*(z)}(x)$ is the minimal valid assignment such that the path from x to it is labeled only by elements in $S_*(z)$. We'll show that y satisfies these properties, and therefore by Proposition 4, $y = \sigma(z)$. Any path from x to y (which exists since $y \in B$) is labeled by $S_*(y) \setminus S_*(x) \subseteq S_*(z)$. Furthermore, for every $i \in S_*(z)$, $i \notin S_0(y)$ so there is no outgoing edge from y labeled by an element in $S_*(z)$. Therefore, y is minimal.

Let $y = \sigma(z) = \gamma_{S_*(z)}(x)$. By Proposition 4, there is no $i \in S_*(z)$ such that $i \in S_0(y)$. Therefore, $S_*(z) \subseteq S_*(y) \cup S_c(y)$. Further, we have that $S_*(y) \subseteq S_*(z) \cup S_*(x) = S_*(z)$; therefore $S_*(z) = S_*(y) \cup T$ for some $T \subseteq S_c(y)$. Hence, $z \in \tau(y)$. \square

For a set of partial assignments X , let $W(X) = \sum_{x \in X} W(x)$. Let $W^{(0)}(z) = (\omega_*)^{n_A(z)} \times (\omega_0)^{n_A(z)}$, denote the weight of any partial assignment, if the formula had no clauses. For such a formula, all partial assignments are valid. Observe that

³ Note, however, that any partial assignment that belongs to two or more lattices is assigned a weight only once. Otherwise, the transformation would be a convolution operation in a strict sense.




FIG. 8. The directed graph G and the map σ for the formula $(x_1 \vee x_2 \vee x_3) \wedge (\bar{x}_2 \vee \bar{x}_3 \vee x_4)$ and the satisfying assignment $(0, 0, 1, 0)$. The solid arrows denote edges in G and the dashed arrows denote σ .

if we restrict our attention to the assignments that are consistent with x ,

$$\begin{aligned} W^\theta(A) &= \sum_{z \in A} W^\theta(z) \\ &= \sum_{S \subseteq V \setminus S_*(x)} (\omega_*)^{|S_*(x)|+|S|} \times (\omega_o)^{|S_*(x)|-|S|} \\ &= (\omega_*)^{|S_*(x)|} \times (\omega_* + \omega_o)^{|S_*(x)|} \\ &= (\omega_*)^{n_*(x)}. \end{aligned}$$

We show that when clauses are added to the formula, the total weight under x is preserved as long as x is still valid. In particular, when an assignment z that is consistent with x becomes invalid, it passes its weight to an assignment that is still valid, namely $\sigma(z)$, which has fewer $*$ variables than z .

$$\begin{aligned} W(y) &= (\omega_*)^{n_*(y)} \times (\omega_o)^{n_o(y)} \times 1^{n_c(y)} \\ &= (\omega_*)^{n_*(y)} \times (\omega_o)^{n_o(y)} \times (\omega_* + \omega_o)^{n_c(y)} \\ &= \sum_{T \subseteq S_c(y)} (\omega_*)^{n_*(y)+|T|} \times (\omega_o)^{n_o(y)+n_c(y)-|T|} \\ &= \sum_{T \subseteq S_c(y)} W^\theta(z : S_*(z) = S_*(y) \cup T) \\ &= W^\theta(\{z : S_*(z) = S_*(y) \cup T, T \subseteq S_c(y)\}) \\ &= W^\theta(\tau(y)). \end{aligned} \tag{17}$$

Finally, we have:

$$\sum_{y \subseteq x} W(y) = \sum_{y \subseteq x} W^\theta(\tau(y)) = W^\theta(A) = (\omega_*)^{n_*(x)}$$

where we used the fact that the sets $\tau(y)$ for $y \in B$ partition A by Lemma 7.

The proof of the case $\omega_o + \omega_* < 1$ is similar except that Eq. (17) becomes an inequality:

$$W(y) = (\omega_o)^{n_o(y)} \times (\omega_*)^{n_*(y)} \times 1^{n_c(y)} \geq \sum_{T \subseteq S_c(y)} W^\theta(\tau(y)).$$

When an assignment z that is consistent with x becomes invalid, it passes more than its own weight to $\sigma(z)$.

4.5. GIBBS SAMPLING. Based on our experiments, the algorithm $\text{SP}(\rho)$ is very effective for appropriate choices of the parameter ρ . The link provided by Theorem 6 suggests that the distribution p_W , for which $\text{SP}(\rho)$ —as an instantiation of belief propagation on the extended MRF—is computing approximate marginals, must possess good “smoothness” properties. One expected consequence of such “smoothness” is that algorithms other than BP should also be effective in computing approximate marginals. Interestingly, rigorous conditions that imply (rapid) convergence of BP [Tatikonda and Jordan 2002]—namely, uniqueness of Gibbs measures on the computation tree—are quite similar to conditions implying rapid convergence of Gibbs samplers, which are often expressed in terms of “uniqueness”, “strong spatial mixing”, and “extremality” (see, e.g., Martinelli [1999] and Berger et al. [2005]).

In this section, we explore the application of sampling methods to the extended MRF as a means of computing unbiased stochastic approximations to the marginal distributions, and hence biases at each variable. More specifically, we implemented a Gibbs sampler for the family of extended MRFs developed in Section 3. The Gibbs sampler performs a random walk over the configuration space of the extended MRF—that is, on the space of partial valid assignments. Each step of the random walk entails picking a variable x_i uniformly at random, and updating it randomly to a new value $b \in \{0, 1, *\}$ according to the conditional probability $p_W(x_i = b | (x_j : j \neq i))$. By the construction of our extended MRF (see Eq. (12)), this conditional probability is an (explicit) function of the variables x_j and x_i appear together in a clause, and of the variables x_k such that x_k and x_j appear together in a clause, where x_j and x_i appear together in a clause.

It is of interest to compare the approximate marginals computed by the $\text{SP}(\beta)$ family of algorithms (to which we refer as *pseudomarginals*) to the (stochastic) estimates computed by Gibbs sampler. Given the manner in which the SP pseudomarginals are used in the decimation procedure, the most natural comparison is between the biases $\mu_i(0) - \mu_i(1)$ provided by the $\text{SP}(\beta)$ algorithm, and the biases $\tau_i(0) - \tau_i(1)$ associated with the Gibbs sampler (where τ_i are the approximate marginals obtained from Gibbs sampling on the extended MRF p_{gen} with parameter ρ , which we denote by $\text{MRF}(\rho)$). The results of such comparisons for the SP parameter $\beta \in \{0.95, 0.9, 0.7, 0.5\}$ and the Gibbs sampling parameter $\rho \in \{0.4, 0.5, 0.7, 0.9\}$ are shown in Figure 9. Comparisons are made for each pair (β, ρ) in these sets, and over a range of clause densities $\alpha \in \{4.2, 4.1, 4.0, 3.8, 3.6, 3.4\}$. For fairly dense formulas (e.g., $\alpha \geq 4.0$), the general trend is that the $\text{SP}(\beta)$ biases with larger β agree most closely with the Gibbs biases with ρ relatively smaller (i.e., $\rho < \beta$). For lower clause densities (e.g., $\alpha = 3.4$), the agreement between the $\text{SP}(\beta)$ and Gibbs biases on $\text{MRF}(\rho)$ when $\beta = \rho$ is substantially closer.

SAT α	Gibbs ρ			
	0.4	0.5	0.7	0.9
4.2	0.0493	0.1401	0.3143	0.4255
4.1	0.0297	0.1142	0.3015	0.4046
4.0	0.0874	0.0416	0.2765	0.3873
3.8	0.4230	0.4554	0.1767	0.0737
3.6	0.4032	0.4149	0.1993	0.0582
3.4	0.4090	0.4010	0.2234	0.0821

(a) Comparison to SP(0.95)

SAT α	Gibbs ρ			
	0.4	0.5	0.7	0.9
4.2	SP fails	SP fails	SP fails	SP fails
4.1	0.0230	0.0985	0.3236	0.4341
4.0	0.0493	0.0079	0.3273	0.4309
3.8	0.0531	0.0194	0.2860	0.4104
3.6	0.0980	0.0445	0.2412	0.3887
3.4	0.0365	0.0356	0.1301	0.3869

(c) Comparison to SP(0.7)

SAT α	Gibbs ρ			
	0.4	0.5	0.7	0.9
4.2	SP fails	SP fails	SP fails	SP fails
4.1	0.1925	0.2873	0.3989	0.4665
4.0	0.0483	0.1092	0.2986	0.4179
3.8	0.0924	0.0372	0.3235	0.4323
3.6	0.0184	0.0304	0.2192	0.4009
3.4	0.0323	0.0255	0.0718	0.3613

(d) Comparison to SP(0.5)

FIG. 9. Comparison of $SP(\beta)$ pseudomarginals for $\beta \in \{0.95, 0.9, 0.7, 0.5\}$ to marginals estimated by Gibbs sampling on weighted MRFs with $\rho \in \{0.4, 0.5, 0.7, 0.9\}$ for the range of SAT problems $\alpha \in \{4.2, 4.1, 4.0, 3.8, 3.6, 3.4\}$. Each entry in each table shows the average ℓ_1 error between the biases computed from the $SP(\beta)$ pseudomarginals compared to the biases computed from Gibbs sampling applied to $MRF(\rho)$. Calculations were based on top 50 most biased nodes on a problem of size $n = 1000$. The bold entry within each row (corresponding to a fixed α) indicates the $MRF(\rho)$ that yields the smallest ℓ_1 error in comparison to the SP biases.

5. Expansion Arguments for Random Formulas

This section is devoted to the study of properties of the MRF on random formulas. We will use simple random graph arguments in order to obtain typical properties of cores, as well as the behavior of Gibbs sampling or message-passing algorithms applied to the MRF associated with a randomly chosen formula. Throughout this section, we denote p_W^ϕ to denote the MRF distribution for a fixed formula ϕ . Otherwise, we write $\mathbb{P}^{n,m}$ for the uniform measure on k -sat formulas with n variables and m clauses, and $\mathbb{P}^{n,\alpha}$ for the uniform measure on k -sat formulas with n variables and $m = \alpha n$ clauses. We often drop n, m , and/or α when they are clear from the context. Finally, we use $\mathbb{E}_W^\phi, \mathbb{E}^{n,m}$ and $\mathbb{E}^{n,\alpha}$ to denote expectations with respect to the distributions $p_W^\phi, \mathbb{P}^{n,m}$ and $\mathbb{P}^{n,\alpha}$ respectively.

5.1. SIZE OF CORES. We first prove a result that establishes that cores, if they exist, are typically at least a certain linear fraction $c(\alpha, k)$ of the total number n of variables.

PROPOSITION 8. Let ϕ be a random k -sat formula with $m = \alpha n$ clauses where $k \geq 3$. Then, for all positive integers C , it holds that

$$\mathbb{P}^{n,\alpha}[\phi \text{ has a core with } C \text{ clauses}] \leq \left(\frac{e^2 \alpha C^{k-2}}{n^{k-2}} \right)^C, \quad (18)$$

Consequently, if we define $c(\alpha, k) := (\alpha e^2)^{-1/(k-2)}$, then with $\mathbb{P}^{n,\alpha}$ -probability tending to one as $n \rightarrow +\infty$, there are no cores of size strictly less than $c(\alpha, k)n$.

PROOF. Suppose that the formula ϕ has a core with C clauses. Note that the variables in these clauses all lie in some set of at most C variables. Thus, the probability that a core with C clauses exist is bounded by the probability that there is a set of C clauses all whose variables lie in some set of size $\leq C$. This probability is bounded by

$$\binom{m}{C} \binom{n}{C} \left(\frac{C}{n} \right)^C,$$

which can be upper bounded by

$$\left(\frac{em}{C} \right)^C \left(\frac{en}{C} \right)^C \binom{C}{n}^{Ck} = \left(\frac{e^{-2} \alpha C^{k-2}}{n^{k-2}} \right)^C,$$

as needed. \square

5.2. (META)-STABILITY OF THE ALL * ASSIGNMENT FOR SMALL ρ . By definition, the extended MRF for $\rho = 1$ assigns positive mass to the all-* vector. Moreover, Proposition 8 implies that the size of cores (when they exist) is typically linear in n . It follows that the state space of the MRF for $\rho = 1$ typically satisfies one of the following properties:

- Either the state space is trivial, meaning that it contains only the all * state, or
- The state space is disconnected with respect to all random walks based on updating a small linear fraction of the coordinates in each step.

The goal of this section is to establish that a similar phenomenon persists when ρ is close to 1 (i.e., when $1 - \rho$ is small).

We begin by introducing some notions from the analysis of the mixing properties of Markov chains. Let T be a reversible chain with respect to a measure p on a state space Ω . For sets $A, B \subset \Omega$, write

$$q_T(A, B) = \sum_{x \in A, y \in B} p(x) T_{x \rightarrow y} = \sum_{x \in A, y \in B} p(y) T_{y \rightarrow x}.$$

The *conductance* of the chain T is defined as

$$c(T) = \inf_{S \subset \Omega} \left\{ \frac{q_T(S, S^c)}{p(S)(1 - p(S))} \right\}.$$

It is well known that $c(T)/2$ is an upper bound on the spectral gap of the chain T and that $2/c(T)$ is a lower bound on the mixing time of the chain. We note moreover that the definition of T implies that for every two sets A, B it holds that $q_T(A, B) \leq \min\{p(A), p(B)\}$.

Definition 9. Consider a probability measure p on a space Ω of strings of length n . Let T be a Markov chain on Ω . The radius of T denoted by $r(T)$ is defined by

$$r(T) := \sup\{d_H(x, y) : T_{x,y} > 0\}, \quad (19)$$

where d_H is the Hamming distance. We let the radius r -conductance of p denote by $c(r, p)$ be

$$c(r, p) := \sup\{c(T) : T \text{ is reversible with respect to } p \text{ and } r(T) \leq r\}. \quad (20)$$

Now returning to the random k -SAT problem, we write p_ρ for the measure $p_W = p_W^\phi$ with $\omega_* = \rho$ and $\omega_o = 1 - \rho$.

PROPOSITION 10. *Consider a randomly chosen k -SAT formula with density α . Then there exists a $\rho_0 \in (0, 1)$ such that if $\rho > \rho_0$ then $\mathbb{P}^n[\phi \in A_n \cup B_n] \rightarrow 1$ as $n \rightarrow +\infty$ where A_n and B_n are the following events:*

(I) A_n consists of all the formulas ϕ satisfying $p_\rho^\phi[n - n_*(x)] \leq 2\sqrt{(1 - \rho)n} \geq 1 - \exp(-\Omega(n))$.

(II) B_n consists of all the formulas ϕ for which the measure p_ρ^ϕ satisfies $c\sqrt{(1 - \rho)n, p_\rho} \leq \exp(-\Omega(n))$.

PROOF. We let δ be a small positive number to be determined, and set $1 - \rho = \delta^2$. As it suffices to work with ratios of probabilities, we use the unnormalized weight $W^\phi(x)$ instead of $p_W^\phi(x)$. The proof requires the following:

LEMMA 11. *Let d be an integer satisfying $\delta n \leq d \leq 2\delta n$. For δ sufficiently small, it holds that with \mathbb{P}^n probability going to 1 as $n \rightarrow \infty$*

$$\sum_{d=\delta n}^{2\delta n} \frac{W^\phi[n - n_* = d]}{\rho^{3n}} = \exp(-\Omega(n)). \quad (21)$$

PROOF. See Appendix C.1. \square

To establish the proposition, it suffices to show that for any formula ϕ for which Eq. (21) of Lemma 11 is valid, then one of either condition (I) or condition (II) must hold.

- (i) First suppose that $W^\phi[n - n_*(x) > 2\delta n] \leq \rho^{3n/2}$. In this case, condition (I) in the statement of the proposition follows immediately.
- (ii) Otherwise, we may take $W^\phi[n - n_*(x) > 2\delta n] \geq \rho^{3n/2}$. In this case, we can apply the conductance bound in order to bound the gap of any operator with radius $\leq \delta n$. Take the set A to be all x with $n - n_*(x) < \delta n$ and B be the set of all x with $\delta n \leq n - n_*(x) \leq 2\delta n$. Let T be any Markov chain with radius δn that is reversible with respect to p_W . Then, we have $q_T(A, A^c) = q_T(A, B) \leq p(B)$. In addition, it holds that $W^\phi[n - n_*(x) < \delta n] \geq \rho^n$ (since if x is the all- $*$ assignment, we have $W^\phi(x) = \rho^n$); moreover, if we take n sufficiently large, then we have $W^\phi[\delta n \leq n - n_*(x) \leq 2\delta n] \leq \rho^{3n}$ by Lemma 11. Combining these inequalities, we obtain that the conductance of T is bounded above by

$$\begin{aligned} \frac{q(A, A^c)}{p(A)p(A^c)} &\leq \frac{p(B)}{p(A)p(A^c)} \\ &\leq \frac{W^\phi[\delta n \leq n - n_*(x) \leq 2\delta n]}{W^\phi[n - n_*(x) < \delta n]W^\phi[n - n_*(x) > 2\delta n]} \\ &\leq \frac{\rho^{3n}}{\rho^n \rho^{\frac{3n}{2}}} = \rho^{n/2}, \end{aligned}$$

which implies condition (II).

5.3. MESSAGE-PASSING ALGORITHMS ON RANDOM ENSEMBLES. The analysis of the preceding section demonstrated that for values of ρ close to 1, any random sampling technique based on local moves (e.g., Gibbs sampling), if started at the all $*$ assignment, will take exponentially long to get to an assignment with more than a negligible fraction of non- $*$. This section is devoted to establishing an analogous claim for the belief propagation updates on the extended Markov random fields. More precisely, we prove that if ρ is sufficiently close to 1, then running belief propagation with initial messages that place most of their mass on $*$ will result in assignments that also place most of the mass on $*$.

This result is proved in the ‘‘density-evolution’’ setting e.g., [Richardson and Urbanke 2001] (i.e., the number of iterations is taken to be less than the girth of the graph, so that cycles have no effect). More formally, we establish the following:

THEOREM 12. *For every formula density $\alpha > 0$, arbitrary scalars $\epsilon'' > 0$ and $\delta > 0$, there exists $\rho' < 1, \epsilon' \in (0, \epsilon'')$ and $\gamma > 0$ such that for all $\rho \in (\rho', 1]$ and $\epsilon \in (0, \epsilon')$, the algorithm $SP(\rho)$ satisfies the following condition.*

Consider a random formula ϕ , a random clause b and a random variable i that belongs to the clause b . Then with probability at least $1 - \delta$, if $SP(\rho)$ is initialized with all messages $\eta_{a \rightarrow j}^0 < \epsilon$, then the inequality $\eta_{b \rightarrow i}^t < \epsilon'$ holds for all iterations $t = 0, 1, \dots, \gamma \log n$.

The first step of the proof is to compare the SP iterations to simpler ‘‘sum-product’’ iterations.

LEMMA 13. *For any $\rho \in [0, 1]$, the $SP(\rho)$ iterations satisfy the inequality:*

$$\eta_{a \rightarrow i}^{t+1} \leq \prod_{j \in V(a) \setminus \{i\}} \min \left(1, (1 - \rho) + \rho \sum_{b \in C(j) \setminus \{a\}} \eta_{b \rightarrow j}^t \right)$$

PROOF. See Appendix C.2. \square

Since our goal is to bound the messages $\eta_{a \rightarrow i}^{t+1}$, Lemma 13 allows us to analyze the simpler message-passing algorithm with updates specified by:

$$\eta_{a \rightarrow i}^{t+1} = \prod_{j \in V(a) \setminus \{i\}} \min \left(1, (1 - \rho) + \rho \sum_{b \in C(j) \setminus \{a\}} \eta_{b \rightarrow j}^t \right). \quad (22)$$

The next step is to bound the probability of ‘‘short-cycles’’ in the computation tree corresponding to the message-passing updates specified in Eq. (22). More formally, given a formula ϕ , we define a directed graph $G(\phi) = (V, E)$, in which the vertex set V consists of messages $\eta_{a \rightarrow i}$. The edge set E includes the edge $\eta_{a \rightarrow i} \rightarrow \eta_{b \rightarrow j}$ belongs to E if and only if $j \in V(a) \setminus \{i\}$ and $b \in C_a^u(i)$. In words, the graph $G(\phi)$ includes an edge between the $\eta_{a \rightarrow i}$ and $\eta_{b \rightarrow j}$ if the latter is involved in the update of $\eta_{a \rightarrow i}$ specified in Eq. (22).

LEMMA 14. *Let $G(\phi)$ be the random graph generated by choosing a formula ϕ uniformly at random with αn clauses and n variables. Let v be a vertex of $G(\phi)$ chosen uniformly at random. For all clause densities $\alpha > 0$, there exists $\gamma > 0$ such that with probability $1 - o(1)$, the vertex v does not belong to any directed cycle of length smaller than $\gamma \log n$ in $G(\phi)$.*

PROOF. The proof is based on standard arguments from random graph theory for example, Janson et al. [2000]. \square

Our analysis of the recursion (22) on the computation tree in based on an edge exposure technique that generates a neighborhood of a vertex v in the graph $G(\phi)$ for a random ϕ . More specifically, pick a clause a and a variable i in a at random. Now for each variable $j \in V(a) \setminus \{i\}$, expose all clauses b containing j (but not any other of the variables appearing so far). Then for each such b , we look at all variables $k \in V(b) \setminus \{j\}$, and so on. We consider the effect of repeating this exposure procedure over $t = \gamma \log n$ steps. When the vertex $\eta_{a \rightarrow i}$ does not belong to cycles shorter than t in $G(\phi)$, such an analysis yields a bound on $\eta_{a \rightarrow i}^t$.

Note that each clause can expose at most $k - 1$ variables. Recall that we generate the formula ϕ by choosing each of the $N_c = 2^k \binom{n}{k}$ clauses with probability $\alpha n / N_c$. The distribution of the number of clauses exposed for each variable is thus dominated by $\text{Bin}(M_c, \alpha n / N_c)$ where $M_c = 2^k \binom{n}{k-1}$. An equivalent description of this process is the following: each vertex $v = \eta_{a \rightarrow i}$ exposes X_v neighbors $\eta_{b \rightarrow j}$, where the distribution of the collection $\{X_v\}$ is dominated by a collection $\{Y_v\}$ of i.i.d. random variables. Moreover, the Y_v 's are jointly distributed as the sum of $k - 1$ i.i.d. $\text{Bin}(M_c, \alpha n / N_c)$ variables.

The proof requires the following lemma on branching processes.

LEMMA 15. *Consider a branching process where each vertex gives birth to Y children. Assume further that the branching process is stopped after m levels and let $K > 0$ be given.*

The notion of a good vertex is defined inductively as follows. All vertices at level m are good. A vertex at level $m - 1$ is good if it has ℓ children and $\ell \leq K$. By induction for $s \geq 2$, we call a vertex at level $m - s$ good if v has ℓ children v_1, \dots, v_ℓ with $\ell \leq K$ and

- (a) *Either all of v_1, \dots, v_ℓ have at most K children, of which all are good; or*
- (b) *all of v_1, \dots, v_ℓ have at most K children, of which all but one are good.*

Denote by $p(m, K)$ the probability that the root of the branching process is good. Then

$$\inf_{0 \leq m < \infty} p(m, K) = 1 - \exp(-\Omega(K)).$$

PROOF. See Appendix C.3. \square

We are now equipped to complete the proof of Theorem 12. Using Lemma 14, first choose $\gamma = \gamma(\alpha)$ such that a random vertex in $G(\phi)$ does not belong to cycles shorter than $\gamma \log n$ with probability $1 - o(1)$. Next, use Lemma 15 to choose K such that the probability $\inf_{0 \leq m < \infty} p(m, K)$ that the root of the branching process is good is at least $1 - \delta/2$.

Next, we define a pair of functions θ and ζ (each mapping $R \times R$ to the real line) in the following way:

$$\theta(\epsilon, \rho) := ((1 - \rho) + K\rho\epsilon), \quad \zeta(\epsilon, \rho) := \theta(\theta(\epsilon, \rho), \rho) \times \theta(\theta(\epsilon, \rho)^2, \rho).$$

Setting $\epsilon' := \min(\epsilon'', \frac{1}{2K^3})$, observe that $\theta(\epsilon', 1) = K\epsilon'$ and therefore $\theta^2(\epsilon', 1) \leq \frac{\epsilon'}{4}$ and

$$\zeta(\epsilon', 1) = \theta(K\epsilon', 1)\theta((K\epsilon')^2, 1) = (K^2\epsilon')(K^4\epsilon'^2) = K^6\epsilon'^3 \leq \frac{\epsilon'}{4}.$$

It now follows by continuity that there exists $\rho' < 1$ such that for all $1 \geq \rho \geq \rho'$ it holds that

$$\theta^2(\epsilon', \rho) \leq \frac{\epsilon'}{2}, \quad \zeta(\epsilon', \rho) \leq \frac{\epsilon'}{2}. \quad (23)$$

We claim that the statement of the theorem holds with the choices of γ, ϵ', ρ' above. Indeed, choose a formula ϕ with density α at random and let $v = \eta_{a \rightarrow i}$ be a random vertex of $G(\phi)$. With probability at least $1 - \delta/2$, the vertex v does not belong to any cycle shorter than $t = \gamma \log n$.

Since v does not belong to any such cycle, the first t levels of the computation tree of v may be obtained by the exposure process defined above. We will then compare the computation tree to an exposure process where each variable gives birth to exactly $\text{Bin}(M_c, \alpha n / N_c)$ clauses. Since the messages are generated according to (22), any bound derived on the values of non-* messages for the larger tree implies the same bound for the real computation tree.

We now claim that if v is a good vertex on that tree, then the message at v after t iterations—namely, $\eta_{a \rightarrow i}^t$ —is at most ϵ' . Since a vertex of the tree is good with probability $1 - \delta/2$, proving this claim will establish the theorem.

We prove this claim by induction on s , where $m - s$ is the level of w . For $s = 0$, the claim follows immediately from the initialization of the messages. For $s = 1$, observe that Eq. (22) implies that if $w = \eta_{b \rightarrow j}$ is good at level $m - 1$, then

$$\eta_{b \rightarrow j} \leq \theta^{k-1}(\rho, \epsilon) \leq \theta^2(\rho, \epsilon') \leq \frac{\epsilon'}{2}.$$

For the general induction step, assume that $w = \eta_{b \rightarrow j}$ at level $m - s$ is good and $s \geq 2$. There are two cases to consider:

- (i) w has all its grandchildren good. In this case, we repeat the argument above twice to obtain $\eta_{b \rightarrow j} \leq \epsilon'$.
- (ii) Exactly one of $w = \eta_{b \rightarrow j}$ grandchildren is not good. Let $y' = \eta_{d' \rightarrow \ell'}$ denote the grandchild and $y = \eta_{d \rightarrow \ell}$ denote y parent. Then, by Eq. (22):

$$\eta_{d \rightarrow \ell} \leq (1 - \rho) + K\rho\epsilon' = \theta(\epsilon', \rho).$$

Using (13) again yields

$$\begin{aligned} \eta_{d \rightarrow \ell} &\leq ((1 - \rho) + K\rho\theta(\epsilon', \rho))(1 - \rho) + K\rho\theta^2(\epsilon', \rho))^{\ell-2} \\ &\leq ((1 - \rho) + K\rho\theta(\epsilon', \rho))(1 - \rho) + K\rho\theta^2(\epsilon', \rho)) = \zeta(\epsilon', \rho) \leq \epsilon'/2, \end{aligned}$$

which completes the proof.

6. Conclusion

The survey propagation algorithm, recently introduced by Mézard et al. [2002] for solving random instances of k -SAT problems, has sparked a great deal of excitement and research in both the statistical physics and computer science communities e.g., Braunstein et al. [2005, 2003], Braunstein and Zecchina [2004], Aurell et al. [2004], Achlioptas and Ricci-Tersenghi [2006], Mora et al. [2005], Parisi [2002], and Wainwright and Maneva [2005]. This article provides a new interpretation of

the survey propagation algorithm—namely, as an instance of the well-known belief propagation algorithm but as applied to a novel probability distribution over the partial satisfiability assignments associated with a k -SAT formula. The perspective of this article reveals the combinatorial structure that underlies the survey propagation algorithm, and we established various results on the form of these structures and the behavior of message-passing algorithms, both for fixed instances and over random ensembles.

The current work suggests various questions and open issues for further research. As we described, a large family of Markov random fields over partial assignments, is associated with any k -SAT problem as specified by the parameter ρ (or more generally, the parameters ω_o and ω_*). Further analysis of survey propagation and its generalizations requires a deeper understanding of the following two questions. First, for what parameter choices do the marginals of the associated Markov random field *yield useful information* about the structure of satisfiability assignments? Second, for what parameter choices do efficient message-passing algorithms like belief propagation *yield accurate approximations* to these marginals? Our results show that the success of SP-like algorithms depends on a delicate balance between these two factors. (For instance, the marginals of the uniform distribution over SAT assignments clearly contain useful information, but belief propagation fails to yield good approximations for sufficiently large clause densities.)

Although this current article has focused on a particular kind of constraint satisfaction problem (namely, k -SAT), it is likely that message-passing algorithms, in the flavor of the survey propagation method, can also be useful for solving other classes of constraint satisfaction problems (CSPs). For instance, some recent work [Wainwright and Maneva 2005; Ciliberti et al. 2005] has shown how forms of message-passing related to survey propagation can be used to obtain excellent approximate solutions for the MAX-XORSAT problem; practical methods for solving this problem have important applications in lossy data compression and digital watermarking [Martinian and Wainwright 2006]. Thus, it is natural to speculate whether message-passing approaches might also be brought to bear fruitfully upon other classes of CSPs. Another interesting link is that between decimation based on survey propagation, and other techniques for solving CSPs. Indeed, as pointed out by one of the referees, decimation combined with message-passing can be viewed as type of depth-first search combined with a “largest bias” heuristic [Russell and Norvig 2003].

Finally, these questions touch upon a broader collection of issues, all related to a deeper understanding of constraint satisfaction problems and especially the relationship between finite instances and their asymptotic analysis. Given the fundamental role that constraint satisfaction problems play in diverse branches of computer science, further progress on these issues is of broad interest.

Appendixes

A. Belief Propagation on a Generic Factor Graph

Given a subset $S \subseteq \{1, 2, \dots, n\}$, we define $x_S := \{x_i \mid i \in S\}$. Consider a probability distribution on n variables x_1, x_2, \dots, x_n , that can be factorized as

$$p(x_1, x_2, \dots, x_n) = \frac{1}{Z} \prod_{i=1}^n \psi_i(x_i) \prod_{a \in C} \psi_a(x_{V(a)}), \quad (24)$$

where for each $a \in C$ the set $V(a)$ is a subset of $\{1, 2, \dots, n\}$; and $\psi_i(x_i)$ and $\psi_a(x_{V(a)})$ are non-negative real functions, referred to as compatibility functions, and

$$Z := \sum_x \left[\prod_{i=1}^n \psi_i(x_i) \prod_{a \in C} \psi_a(x_{V(a)}) \right] \quad (25)$$

is the normalization constant or partition function. A factor graph representation of this probability distribution is a bipartite graph with vertices V corresponding to the variables, called *variable nodes*, and vertices C corresponding to the sets $V(a)$ and called *function nodes*. There is an edge between a variable node i and function node a if and only if $i \in V(a)$. We write also $a \in C(i)$ if $i \in V(a)$.

Suppose that we wish to compute the marginal probability of a single variable i for such a distribution, as defined in Eq. (5). The belief propagation or sum-product algorithm [Kschischang 2001] is an efficient algorithm for computing the marginal probability distribution of each variable, assuming that the factor graph is acyclic. The essential idea is to use the distributive property of the sum and product operations to compute independent terms for each subtree recursively. These recursions can be cast as a message-passing algorithm, in which adjacent nodes on the factor graph exchange intermediate values. Let each node only have access to its corresponding compatibility function. As soon as a node has received messages from all neighbors below it, it can send a message up the tree containing the term in the computation corresponding to it. In particular, let the vectors $M_{i \rightarrow a}$ denote the message passed by variable node i to function node a ; similarly, the quantity $M_{a \rightarrow i}$ denotes the message that function node a passes to variable node i . The messages from function to variables are updated in the following way:

$$M_{a \rightarrow i}(x_i) \propto \sum_{x_{V(a) \setminus \{i\}}} \left[\psi_a(x_{V(a)}) \prod_{j \in V(a) \setminus \{i\}} M_{j \rightarrow a}(x_j) \right]. \quad (26)$$

In the other direction, the messages from variable nodes to function nodes are updated as follows:

$$M_{i \rightarrow a}(x_i) \propto \psi_i(x_i) \prod_{b \in C(i) \setminus \{a\}} M_{b \rightarrow i}(x_i). \quad (27)$$

It is straightforward to show that for a factor graph without cycles, these updates will converge after a finite number of iterations. Upon convergence, the local marginal distributions at variable nodes and function nodes can be computed, using the message fixed point \hat{M} , as follows:

$$F_i(x_i) \propto \psi_i(x_i) \prod_{b \in C(i)} \hat{M}_{b \rightarrow i}(x_i) \quad (28a)$$

$$F_a(x_{V(a)}) \propto \psi_a(x_{V(a)}) \prod_{j \in V(a)} \hat{M}_{j \rightarrow a}(x_j). \quad (28b)$$

The same updates, when applied to a graph with cycles, are no longer exact due to presence of cycles. An exact algorithm will generally require exponential time. For certain problems, including error-control coding, applying belief propagation to a graph with cycles gives excellent results. Since there are no leaves on graphs with cycles, usually the algorithm is initialized by sending random messages on all

edges, and is run until the messages converge to some fixed value [Kschischang 2001], or if the messages do not converge, until some fixed number of iterations.

B. Derivation of BP Updates on the Extended MRF

B.1. MESSAGES FROM VARIABLES TO CLAUSES. We first focus on the update of messages from variables to clauses. Recall that we use the notation $P_i = S \cup \{a\}$ as a shorthand for the event

$$a \in P_i \text{ and } S = P_i \setminus \{a\} \subseteq C_d^s(i),$$

where it is understood that S could be empty.

LEMMA 16 (VARIABLE TO CLAUSE MESSAGES). *The variable-to-clause message vector $M_{i \rightarrow a}$ is fully specified by values for pairs (x_i, P_i) of the form:*

$$\{(s_{a,i}, S \cup \{a\}), (s_{a,i}, \emptyset \neq P_i \subseteq C_d^s(i)), \\ (u_{a,i}, \emptyset \neq P_i \subseteq C_d^u(i)), (s_{a,i}, \emptyset), (u_{a,i}, \emptyset), (*, \emptyset)\}.$$

Specifically, the updates for these five pairs take the following form:

$$M_{i \rightarrow a}(s_{a,i}, P_i = S \cup \{a\}) = \prod_{b \in S} M_{b \rightarrow i}^s \prod_{b \in C_d^s(i) \setminus S} M_{b \rightarrow i}^* \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^u \quad (29a)$$

$$M_{i \rightarrow a}(s_{a,i}, \emptyset \neq P_i \subseteq C_d^s(i)) = \prod_{b \in P_i} M_{b \rightarrow i}^s \prod_{b \in C_d^s(i) \setminus P_i} M_{b \rightarrow i}^* \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^u \quad (29b)$$

$$M_{i \rightarrow a}(u_{a,i}, \emptyset \neq P_i \subseteq C_d^u(i)) = \prod_{b \in P_i} M_{b \rightarrow i}^s \prod_{b \in C_d^s(i) \setminus P_i} M_{b \rightarrow i}^* \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^u \quad (29c)$$

$$M_{i \rightarrow a}(s_{a,i}, P_i = \emptyset) = \omega_0 \prod_{b \in C_d^s(i)} M_{b \rightarrow i}^* \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^u \quad (29d)$$

$$M_{i \rightarrow a}(u_{a,i}, P_i = \emptyset) = \omega_0 \prod_{b \in C_d^s(i)} M_{b \rightarrow i}^* \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^u \quad (29e)$$

$$M_{i \rightarrow a}(*, P_i = \emptyset) = \omega_* \prod_{b \in C(i) \setminus \{a\}} M_{b \rightarrow i}^* \quad (29f)$$

PROOF. The form of these updates follows immediately from the definition (10) of the variable compatibilities in the extended MRF, and the BP message update (27). \square

B.2. FORMS OF R QUANTITIES AND FINAL APPROXIMATIONS TO LOCAL MARGINALS. In this section, we compute the specific forms of the linear sums of messages defined in Eqs. (14)–(16). First, we use definition (14) and Lemma (16) to compute the form of $R_{i \rightarrow a}^s$:

$$\begin{aligned} R_{i \rightarrow a}^s &:= \sum_{S \subseteq C_d^s(i)} M_{i \rightarrow a}(s_{a,i}, P_i = S \cup \{a\}) \\ &= \sum_{S \subseteq C_d^s(i)} \prod_{b \in S} M_{b \rightarrow i}^s \prod_{b \in C_d^s(i) \setminus S} M_{b \rightarrow i}^* \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^u \\ &= \prod_{b \in C_d^s(i)} M_{b \rightarrow i}^u \left[\prod_{b \in C_d^s(i)} (M_{b \rightarrow i}^s + M_{b \rightarrow i}^*) \right]. \end{aligned}$$

Similarly, Definition (15) and Lemma (16) allow us to compute the following form of $R_{i \rightarrow a}^u$:

$$\begin{aligned} R_{i \rightarrow a}^u &:= \sum_{S \subseteq C_d^u(i)} M_{i \rightarrow a}(u_{a,i}, P_i = S) \\ &= \sum_{S \subseteq C_d^u(i), S \neq \emptyset} \prod_{b \in S} M_{b \rightarrow i}^s \prod_{b \in C_d^s(i) \setminus S} M_{b \rightarrow i}^* \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^u \\ &\quad + \omega_0 \prod_{b \in C_d^s(i)} M_{b \rightarrow i}^* \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^u \\ &= \prod_{b \in C_d^s(i)} M_{b \rightarrow i}^u \left[\prod_{b \in C_d^u(i)} (M_{b \rightarrow i}^s + M_{b \rightarrow i}^*) - (1 - \omega_0) \prod_{b \in C_d^s(i)} M_{b \rightarrow i}^* \right]. \end{aligned}$$

Finally, we compute $R_{i \rightarrow a}^*$ using Definition (16) and Lemma (16):

$$\begin{aligned} R_{i \rightarrow a}^* &:= \left[\sum_{S \subseteq C_d^s(i)} M_{i \rightarrow a}(s_{a,i}, P_i = S) \right] + M_{i \rightarrow a}(*, P_i = \emptyset) \\ &= \left[\sum_{S \subseteq C_d^s(i), S \neq \emptyset} \prod_{b \in S} M_{b \rightarrow i}^s \prod_{b \in C_d^s(i) \setminus S} M_{b \rightarrow i}^* \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^u \right] \\ &\quad + \omega_0 \prod_{b \in C_d^s(i)} M_{b \rightarrow i}^* \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^u + \omega_* \prod_{b \in C_d^s(i)} M_{b \rightarrow i}^* \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^* \\ &= \prod_{b \in C_d^u(i)} M_{b \rightarrow i}^u \left[\prod_{b \in C_d^s(i)} (M_{b \rightarrow i}^s + M_{b \rightarrow i}^*) - (1 - \omega_0) \prod_{b \in C_d^s(i)} M_{b \rightarrow i}^* \right] \\ &\quad + \omega_* \prod_{b \in C_d^s(i) \cup C_d^u(i)} M_{b \rightarrow i}^*. \end{aligned}$$

The estimates of the local marginals are derived similarly. Here we just demonstrate the derivation of $F_i(1)$:

$$\begin{aligned} F_i(1) &= \sum_{P_i \subseteq C(i)} F_i(1, P_i) \\ &\propto \sum_{P_i \subseteq C(i)} \prod_{a \in C^+(i)} M_{a \rightarrow i}(s_{a,i}, P_i) \prod_{a \in C^-(i)} M_{a \rightarrow i}(u_{a,i}, P_i) \\ &= \prod_{a \in C^-(i)} M_{a \rightarrow i}^u \left[\sum_{P_i \subseteq C(i), P_i \neq \emptyset} \prod_{a \in P_i} M_{a \rightarrow i}^s \prod_{a \in C^+(i) \setminus P_i} M_{a \rightarrow i}^* + \omega_0 \prod_{a \in C^+(i)} M_{a \rightarrow i}^* \right] \\ &= \prod_{a \in C^-(i)} M_{a \rightarrow i}^u \left[\prod_{a \in C^+(i)} (M_{a \rightarrow i}^s + M_{a \rightarrow i}^*) - (1 - \omega_0) \prod_{a \in C^+(i)} M_{a \rightarrow i}^* \right] \end{aligned}$$

B.3. CLAUSE TO VARIABLE UPDATES. In this section, we derive the form of the clause to variable updates.

LEMMA 17 (CLAUSE TO VARIABLE MESSAGES). *The updates of messages from clauses to variables in the extended MRF take the following form:*

$$M_{a \rightarrow i}^s = \prod_{j \in V(a) \setminus \{i\}} R_{j \rightarrow a}^u \quad (30a)$$

$$M_{a \rightarrow i}^u = \prod_{j \in V(a) \setminus \{i\}} (R_{j \rightarrow a}^u + R_{j \rightarrow a}^*) + \sum_{k \in V(a) \setminus \{i\}} (R_{k \rightarrow a}^s - R_{k \rightarrow a}^*) \prod_{j \in V(a) \setminus \{i, k\}} R_{j \rightarrow a}^u - \prod_{j \in V(a) \setminus \{i\}} R_{j \rightarrow a}^u \quad (30b)$$

$$M_{a \rightarrow i}^* = \prod_{j \in V(a) \setminus \{i\}} (R_{j \rightarrow a}^u + R_{j \rightarrow a}^*) - \prod_{j \in V(a) \setminus \{i\}} R_{j \rightarrow a}^u. \quad (30c)$$

PROOF.

(i) We begin by proving Eq. (30a). When $x_i = s_{a,i}$ and $P_i = S \cup \{a\}$ for some $S \subseteq C_a^s(i)$, then the only possible assignment for the other variables at nodes in $V(a) \setminus \{i\}$ is $x_j = u_{a,j}$ and $P_j \subseteq C_a^u(j)$. Accordingly, using the BP update equation (26), we obtain the following update for $M_{a \rightarrow i}^s = M_{a \rightarrow i}(s_{a,i}, P_i = S \cup \{a\})$:

$$\begin{aligned} M_{a \rightarrow i}^s &= \prod_{j \in V(a) \setminus \{i\}} \sum_{P_j \subseteq C_a^u(j)} M_{j \rightarrow a}(u_{a,j}, P_j) \\ &= \prod_{j \in V(a) \setminus \{i\}} R_{j \rightarrow a}^u. \end{aligned}$$

(ii) Next we prove Eq. (30c). In the case $x_i = *$ and $P_i = \emptyset$, the only restriction on the other variables $\{x_j : j \in V(a) \setminus \{i\}\}$ is that they are not all unsatisfying. The weight assigned to the event that they are all unsatisfying is

$$\begin{aligned} &\sum_{\{S_j \subseteq C_a^u(j) : j \in V(a) \setminus \{i\}\}} \prod_{j \in V(a) \setminus \{i\}} M_{j \rightarrow a}(u_{a,j}, S_j) \\ &= \prod_{j \in V(a) \setminus \{i\}} \left[\sum_{S_j \subseteq C_a^u(j)} M_{j \rightarrow a}(u_{a,j}, S_j) \right] \\ &= \prod_{j \in V(a) \setminus \{i\}} R_{j \rightarrow a}^u. \end{aligned} \quad (31)$$

On the other hand, the weight assigned to the event that each is either unsatisfying, satisfying or $*$ can be calculated as follows. Consider a partition $J^u \cup J^s \cup J^*$ of the set $V(a) \setminus \{i\}$, where J^u, J^s and J^* corresponds to the subsets of unsatisfying, satisfying and $*$ assignments respectively. The weight $W(J^u, J^s, J^*)$ associated with this partition takes the form

$$\sum_{\{S_j \subseteq C_a^u(j) : j \in J^u\}} \sum_{\{S_j \subseteq C_a^s(j) : j \in J^s\}} \prod_{j \in J^u} M_{j \rightarrow a}(u_{a,j}, S_j) \prod_{j \in J^s} M_{j \rightarrow a}(s_{a,j}, S_j) \prod_{j \in J^*} M_{j \rightarrow a}(*, \emptyset).$$

Simplifying by distributing the sum and product leads to

$$\begin{aligned} W(J^u, J^s, J^*) &= \prod_{j \in J^u} \left[\sum_{S_j \subseteq C_a^u(j)} M_{j \rightarrow a}(u_{a,j}, S_j) \right] \\ &\quad \times \prod_{j \in J^s} \left[\sum_{S_j \subseteq C_a^s(j)} M_{j \rightarrow a}(s_{a,j}, S_j) \right] \prod_{j \in J^*} M_{j \rightarrow a}(*, \emptyset) \\ &= \prod_{j \in J^u} R_{j \rightarrow a}^u \prod_{j \in J^s} [R_{j \rightarrow a}^s - M_{j \rightarrow a}(*, \emptyset)] \prod_{j \in J^*} M_{j \rightarrow a}(*, \emptyset), \end{aligned}$$

where we have used the definitions of $R_{j \rightarrow a}^u$ and $R_{j \rightarrow a}^*$ from Section B.2. Now summing $W(J^u, J^s, J^*)$ over all partitions $J^u \cup J^s \cup J^* = V(a) \setminus \{i\}$ yields

$$\begin{aligned} &\sum_{J^u \cup J^s \cup J^*} W(J^u, J^s, J^*) \\ &= \sum_{J^u \subseteq V(a) \setminus \{i\}} \prod_{j \in J^u} R_{j \rightarrow a}^u \\ &\quad \times \sum_{J^s \cup J^* = V(a) \setminus \{J^u \cup i\}} \left\{ \prod_{j \in J^s} [R_{j \rightarrow a}^s - M_{j \rightarrow a}(*, \emptyset)] \prod_{j \in J^*} M_{j \rightarrow a}(*, \emptyset) \right\} \\ &= \sum_{J^u \subseteq V(a) \setminus \{i\}} \prod_{j \in J^u} R_{j \rightarrow a}^u \prod_{j \in V(a) \setminus \{J^u \cup i\}} R_{j \rightarrow a}^* \\ &= \prod_{j \in V(a) \setminus \{i\}} [R_{j \rightarrow a}^u + R_{j \rightarrow a}^*], \end{aligned} \quad (32)$$

where we have used the binomial identity twice. Overall, equations (31) and (32) together yield that

$$M_{a \rightarrow i}^* = \prod_{j \in V(a) \setminus \{i\}} [R_{j \rightarrow a}^u + R_{j \rightarrow a}^*] - \prod_{j \in V(a) \setminus \{i\}} R_{j \rightarrow a}^u,$$

which establishes Eq. (30c).

(iii) Finally, turning to Eq. (30b), for $x_i = u_{a,i}$ and $P_i \subseteq C_a^u(i)$, there are only two possibilities for the values of $x_{V(a) \setminus \{i\}}$:

- (a) either there is one satisfying variable and everything else is unsatisfying, or
- (b) there are at least two variables that are satisfying or $*$.

We first calculate the weight $W(A)$ assigned to possibility (a), again using the BP update equation (26):

$$\begin{aligned} W(A) &= \sum_{k \in V(a) \setminus \{i\}} \sum_{S^k \subseteq C_a^s(k)} M_{k \rightarrow a}(s_{a,k}, S^k \cup \{a\}) \prod_{j \in V(a) \setminus \{i, k\}} \sum_{S_j \subseteq C_a^u(j)} M_{j \rightarrow a}(u_{a,j}, S_j) \\ &= \sum_{k \in V(a) \setminus \{i\}} R_{k \rightarrow a}^s \prod_{j \in V(a) \setminus \{i, k\}} R_{j \rightarrow a}^u, \end{aligned} \quad (33)$$

where we have used the definitions of $R_{k \rightarrow a}^s$ and $R_{k \rightarrow a}^u$ from Section B.2.

We now calculate the weight $W(B)$ assigned to possibility (b) in the following way. From our calculations in part (ii), we found that the weight assigned to the event that each variable is either unsatisfying, satisfying or $*$ is $\prod_{j \in V(\alpha) \setminus \{i\}} [R_{j \rightarrow a}^u + R_{j \rightarrow a}^*]$. The weight $W(B)$ is given by subtracting from this quantity the weight assigned to the event that there are *not* at least two $*$ or satisfying assignments. This event can be decomposed into the disjoint events that either all assignments are unsatisfying (with weight $\prod_{j \in V(\alpha) \setminus \{i\}} R_{j \rightarrow a}^u$ from part (ii)); or that exactly one variable is $*$ or satisfying. The weight corresponding to this second possibility is

$$\begin{aligned} & \sum_{k \in V(\alpha) \setminus \{i\}} \left[M_{k \rightarrow a}(*, \emptyset) + \sum_{S^k \subseteq C_a^*(k)} M_{k \rightarrow a}(S_{k,a}, S^k) \right] \\ & \times \prod_{j \in V(\alpha) \setminus \{i,k\}} \sum_{S^j \subseteq C_j^*(\alpha)} M_{j \rightarrow a}(u_{j,a}, S^j) \\ & = \sum_{k \in V(\alpha) \setminus \{i\}} R_{k \rightarrow a}^* \prod_{j \in V(\alpha) \setminus \{i,k\}} R_{j \rightarrow a}^u \\ & W(B) = \prod_{j \in V(\alpha) \setminus \{i\}} [R_{j \rightarrow a}^u + R_{j \rightarrow a}^*] - \sum_{k \in V(\alpha) \setminus \{i\}} R_{k \rightarrow a}^* \prod_{j \in V(\alpha) \setminus \{i,k\}} R_{j \rightarrow a}^u \\ & \quad - \prod_{j \in V(\alpha) \setminus \{i\}} R_{j \rightarrow a}^u. \end{aligned} \quad (34)$$

Combining our calculations so far, we have

$$\begin{aligned} W(B) &= \prod_{j \in V(\alpha) \setminus \{i\}} [R_{j \rightarrow a}^u + R_{j \rightarrow a}^*] - \sum_{k \in V(\alpha) \setminus \{i\}} R_{k \rightarrow a}^* \prod_{j \in V(\alpha) \setminus \{i,k\}} R_{j \rightarrow a}^u \\ &\quad - \prod_{j \in V(\alpha) \setminus \{i\}} R_{j \rightarrow a}^u. \end{aligned} \quad (34)$$

Finally, summing together the forms of $W(A)$ and $W(B)$ from Eqs. (33) and (34), respectively, and then factoring yields the desired Eq. (30b). \square

C. Proofs for Random Formulae

C.1. PROOF OF LEMMA 11. In order to prove (21), it suffices by the Markov inequality to show that for every integer d in the interval $[\delta n, 2\delta n]$, it holds that

$$\frac{\mathbb{E}^n[W^\alpha[n - n_* = d]]}{\rho^{3n}} = \exp(-\Omega(n)). \quad (35)$$

To establish (35), consider a fixed set of d variables. The average W -weight assigned to the event that this set of size d constitutes all the nonstar variables is bounded by

$$\rho^{n-d} \sum_{r=0}^d (1-\rho)^{d-r} \binom{d}{r} \left(\frac{\alpha n}{r}\right) (d/n)^{kr},$$

where r represents the number of constrained variables. We obtain this bound by the following reasoning. First, the $n-d$ variables assigned $*$ all receive weight ρ . Otherwise, if r out of the remaining d variables are constrained, there must be r clauses chosen from a total of αn , and each such clause must have all of its k variables chosen from within the set of d non-star variables.

Consequently, the total probability of having d non-star variables is bounded by

$$\rho^{n-d} \binom{n}{d} \sum_{r=0}^d (1-\rho)^{d-r} \binom{d}{r} \left(\frac{\alpha n}{r}\right) \left(\frac{d}{n}\right)^{kr}$$

$$\begin{aligned} & \leq \rho^{n-d} \left(\frac{en}{d}\right)^d \sum_{r=0}^d (1-\rho)^{d-r} \left(\frac{ed}{r}\right)^r \left(\frac{\alpha en}{r}\right)^r \left(\frac{d}{n}\right)^{kr} \\ & = \rho^{n-d} \left(\frac{(1-\rho)en}{d}\right)^d \sum_{r=0}^d \left(\frac{e^2 d^{k+1} \alpha}{r^2 (1-\rho)n^{k-1}}\right)^r. \end{aligned}$$

Recalling that $1-\rho = \delta^2$ and $d \in [\delta n, 2\delta n]$, we obtain that the last expression is at most

$$\begin{aligned} \rho^{n-2\delta n} \left(\frac{\delta^2 en}{\delta n}\right)^d \sum_{r=0}^{2\delta n} \left(\frac{e^2 (2\delta n)^{k+1} \alpha}{r^2 \delta^2 n^{k-1}}\right)^r &= \rho^{n-2\delta n} (\delta e)^d \sum_{r=0}^{2\delta n} \left(\frac{e^2 2^{k+1} \delta^{k-1} n^2 \alpha}{r^2}\right)^r \\ &\leq \rho^{n-2} (\delta e)^{\delta n} \sum_{r=0}^{2\delta n} \left(\frac{2^{k+1} \alpha \delta^{k-1} n^2 e^2}{r^2}\right)^r, \end{aligned}$$

where the final inequality is valid when $\delta e < 1$. A straightforward calculation yields that the function $g(r) := \left(\frac{2^{k+1} \alpha \delta^{k-1} e^2 n^2}{r^2}\right)^r$ is maximized at $r^* = \sqrt{2^{k+1} \alpha \delta^{k-1} n}$ and the associated value is $g(r^*) = e^{2r^*}$. Consequently, the sum above is bounded by

$$\begin{aligned} 2\delta n \rho^{n-2\delta n} (\delta e)^{\delta n} \exp(2r^*) &= 2\delta n \rho^{n-2\delta n} \left[\delta \exp\left(1 + \frac{2r^*}{\delta n}\right) \right]^{\delta n} \\ &= 2\delta n \rho^{n-2\delta n} \left[\delta \exp\left(1 + \sqrt{2^{k+3} \alpha \delta^{k-3}}\right) \right]^{\delta n} \\ &\leq 2\delta n \rho^{n-2\delta n} \left[\delta \exp\left(1 + \sqrt{2^{k+3} \alpha}\right) \right]^{\delta n}. \end{aligned} \quad (36)$$

This expression is exponentially smaller than ρ^{3n} for large n if

$$\left[\delta \exp\left(1 + \sqrt{2^{k+3} \alpha}\right) \right]^\delta < \rho^3 = (1-\delta^2)^3.$$

Inequality (36) holds for sufficiently small $\delta > 0$, which establishes the lemma.

C.2. PROOF OF LEMMA 13. It will be useful to denote $\prod_{b \in C_j^*(i)} (1 - \eta_{b \rightarrow i})$ by $P_s(i)$ and $\prod_{b \in C_a^*(i)} (1 - \eta_{b \rightarrow i})$ by $P_u(j)$. With this notation, the j th term in (6) is given by

$$\begin{aligned} \frac{\prod_{j \rightarrow a}^u}{\prod_{j \rightarrow a}^s + \prod_{j \rightarrow a}^* + \prod_{j \rightarrow a}^s} &= \frac{(1-\rho P_u(j)) P_s(j)}{(1-\rho P_u(j)) P_s(j) + (1-P_s(j)) P_u(j) + P_s(j) P_u(j)} \\ &= \frac{(1-\rho P_u(j)) P_s(j)}{P_s(j) + P_u(j) - \rho P_s(j) P_u(j)} \leq 1 - \rho P_u(j). \end{aligned}$$

We therefore conclude that

$$\eta_{a \rightarrow i} \leq \prod_{j \in V(\alpha) \setminus \{i\}} (1 - \rho P_u(j)).$$

On the other hand, we have $P_u(j) = \prod_{b \in C_u^c(i)} (1 - \eta_{b \rightarrow i}) \geq \max\left(0, 1 - \sum_{b \in C_u^c(i)} \eta_{b \rightarrow i}\right)$, so that

$$1 - \rho P_u(j) \leq \min\left(1, (1 - \rho) + \rho \sum_{b \in C_u^c(i)} \eta_{b \rightarrow i}\right).$$

This yields the bound $\eta_{a \rightarrow i}^{t+1} \leq \prod_{j \in V(a) \setminus \{i\}} \min\left(1, (1 - \rho) + \rho \sum_{b \in C_u^c(i)} \eta_{b \rightarrow j}\right)$, from which Eq. (22) follows.

C.3. PROOF OF LEMMA 15. We start by estimating the probability that a vertex is bad by induction. Let g_K denote the probability that v has more than K children, or that one of v 's children has more than K children. Clearly,

$$g_K \leq (K+1)\mathbb{P}[Y \geq K] \leq (K+1)(k-1)\mathbb{P}\left[\text{Bin}\left(M_c, \frac{\alpha n}{N_c}\right) \geq \frac{K}{k-1}\right] \leq \exp(-\Omega(K)). \quad (37)$$

Write $q(m, K) = 1 - p(m, K)$ and note that $q(0, K) = 0$ and $q(1, K) \leq g_K$. By induction, a vertex can be bad for two reasons: it has too many descendants in the two levels below it, or it has two bad descendants in the two levels below it. We may thus bound the probability of a vertex being bad as

$$q(s, K) \leq g_K + \mathbb{P}[\text{Bin}(K^2, q(s-2, K)) \geq 2]. \quad (38)$$

Note also that

$$\mathbb{P}[\text{Bin}(K^2, q(s-2, K)) \geq 2] \leq K^4 q(s-2, K)^2. \quad (39)$$

Combining (38) and (39) yields

$$q(s, K) \leq g_K + K^4 q(s-2, K)^2. \quad (40)$$

By (37), when K is sufficiently large $K^4(2g_K)^2 < g_K$. Thus, when K is sufficiently large, it follows from Eq. (40) that

$$q(s, K) \leq 2g_K$$

for all s . Finally when K is sufficiently large, $p(s, K) \geq 1 - 2g_K$ for all s and $1 - 2g_K \geq 1 - \exp(-\Omega(K))$ as needed.

ACKNOWLEDGMENT. We would like to thank Dimitris Achlioptas, Federico Ardila, Scott Kirkpatrick, Andrea Montanari, Mark Mézard, Giorgio Parisi and Alistair Sinclair for helpful discussions.

REFERENCES

- ACHLIOPTAS, D., AND PERES, Y. 2003. The threshold for random k -SAT is $2^{1/2} \log 2 - o(k)$. In *Proceedings of the Symposium on the Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, CA, 223–231.
- ACHLIOPTAS, D., AND RICCI-TERSENGHI, F. 2006. On the solution-space geometry of random constraint satisfaction problems. In *Proceedings of the Symposium on Theory of Computing*. ACM, New York, 130–139.
- AURELL, E., GORDON, U., AND KIRKPATRICK, S. 2004. Comparing beliefs, surveys, and random walks. In *Proceedings of NIPS*. MIT Press, Cambridge, MA.

- BERGER, N., KENYON, C., MOSSEL, E., AND PERES, Y. 2005. Glauber dynamics on trees and hyperbolic graphs. *Probab. Theory Relat. Fields*, 131(3):311–340. (Extended abstract by Kenyon, Mossel and Peres appeared in *Proceedings of FOCS*.)
- BRAUNSTEIN, A., MÉZARD, M., WEIGT, M., AND ZECCHINA, R. 2003. Constraint satisfaction by survey propagation. Technical report. (Preprint at URL: <http://lanl.arXiv.org/cond-mat/0212451>.)
- BRAUNSTEIN, A., MÉZARD, M., AND ZECCHINA, R. 2005. Survey propagation: An algorithm for satisfiability. *Rand. Struct. Algor.*, 27:201–226.
- BRAUNSTEIN, A., AND ZECCHINA, R. 2004. Survey propagation as local equilibrium equations. *J. Stat. Mech. Theory and Experiments (JSTAT)* 7, 06007.
- CHVATAL, V., AND REED, B. 1992. Mick gets some (the odds are on his side). In *Proceedings of the Symposium on the Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, CA, 620–627.
- CILIBERTI, S., MÉZARD, M., AND ZECCHINA, R. 2005. Message-passing algorithms for non-linear nodes and data compression. Technical report, November 2005. arXiv:cond-mat/0508723.
- COOK, S. 1971. The complexity of theorem-proving procedures. In *Proceedings of the Symposium on Theory of Computing*. ACM, New York, 151–158.
- COOPER, G. 1990. The computational complexity of probabilistic inference using Bayesian belief networks. *Artif. Intell.* 42, 393–405.
- COUGHLAN, J., AND FERREIRA, S. 2002. Finding deformable shapes using loopy belief propagation. In *Proceedings of European Conference on Computer Vision*. Springer-Verlag, New York, 453–468.
- DAGUM, P., AND LUBY, M. 1993. Approximate probabilistic reasoning in Bayesian belief networks is NP-hard. *Artif. Intell.* 60, 141–153.
- DE LA VEGA, W. F. 1992. On random 2-sat. Unpublished manuscript.
- DECHTER, R. 2003. *Constraint processing*. Morgan-Kaufmann, Palo Alto, CA.
- DUBOIS, O., BOUFKHAJ, Y., AND MANDLER, J. 2000. Typical random 3-sat formulae and the satisfiability threshold. In *Proceedings of SODA*. ACM, New York, 126–127.
- FEIGE, U. 2002. Relations between average case complexity and approximation complexity. In *Proceedings of Symposium on Theory of Computing*. ACM, New York, 534–543.
- FREEMAN, W. T., PASZTOR, E. C., AND CARMICHAEL, O. T. 2000. Learning low-level vision. *Int. J. Comput. Vision* 40, 1, 25–47.
- FRIEDGUT, E. 1999. Necessary and sufficient conditions for sharp thresholds of graph properties and the k -problem. *J. Amer. Math. Soc.* 12, 1017–1054.
- GALLAGER, R. G. 1963. *Low-density parity check codes*. MIT Press, Cambridge, MA.
- GOERTZ, A. 1996. A remark on random 2-sat. *J. Comput. Syst. Sci.* 53, 469–486.
- JANSON, S., LUCZAK, T., AND RUCINSKI, A. 2000. *Random graphs*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York.
- KAPORIS, A., KIROUSIS, L. M., AND LALAS, E. G. 2000. The probabilistic analysis of a greedy satisfiability algorithm. In *Proceedings of 10th Annual European Symposium on Algorithms*. Springer-Verlag, Berlin, Germany, 574–585.
- KSCHISCHANG, F., FREY, B., AND LOELIGER, H.-A. 2001. Factor graphs and the sum-product algorithm. *IEEE Trans. Inf. Theory* 47, 498–519.
- KUMAR, V. 1992. Algorithms for constraint-satisfaction problems: A survey. *AI Mag.* 13, 1, 32–44.
- LEVIN, L. 1986. Average case for complete problems. *SIAM J. Comput.* 15, 1, 285–286.
- MANEVA, E., MOSSEL, E., AND WAINWRIGHT, M. J. 2005. A new look at survey propagation and its generalizations. In *Proceedings of SODA*. 1089–1098.
- MARTINELLI, F. 1999. Lectures on Glauber dynamics for discrete spin models. In *Lectures on probability theory and statistics (Saint-Flour, 1997)*. Lecture Notes in Math., vol. 1717. Springer-Verlag, Berlin, Germany, 93–191.
- MARTINIAN, E., AND WAINWRIGHT, M. J. 2006. Low density codes can achieve the Wyner-Ziv and Gelfand-Pinsker bounds. In *International Symposium on Information Theory*. IEEE Computer Society Press, Los Alamitos, CA.
- MÉZARD, M., PARISI, G., AND ZECCHINA, R. 2002. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297, 812. (Scienceexpress published on-line 27-June-2002; 10.1126/science.1073287)
- MÉZARD, M., RICCI-TERSENGHI, F., AND ZECCHINA, R. 2003. Two solutions to diluted p -spin models and XORSAT problems. *J. Stat. Phys.* 111, 505.
- MÉZARD, M., AND ZECCHINA, R. 2002. Random k -satisfiability: From an analytic solution to an efficient algorithm. *Phys. Rev. E*, 66.

- MONASSON, R., AND ZECCHINA, R. 1997. Statistical mechanics of the random k -satisfiability model. *Phys. Rev. E* 3, 1357–1370.
- MORA, T., MÉZARD, M., AND ZECCHINA, R. 2005. Clustering of solutions in the random satisfiability problem. *Phys. Rev. Lett.* 94.
- PARISI, G. 2002. On local equilibrium equations for clustering states. Technical report. Preprint at URL: <http://lanl.arXiv.org/cs.CC/0212047>.
- PEARL, J. 1988. *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Morgan-Kaufmann, Palo Alto, CA.
- RICHARDSON, T., AND URBANKE, R. 2001. The capacity of low-density parity check codes under message-passing decoding. *IEEE Trans. Info. Theory* 47, 599–618.
- ROSENTHAL, J., PLOTKIN, J., AND FRANCO, J. 1999. The probability of pure literals. *J. Computat. Logic* 9, 501–513.
- RUSSELL, S., AND NORVIG, P. 2003. *Artificial Intelligence: A Modern Approach*. Prentice-Hall, Englewood Cliffs, NJ, 2nd edition.
- SELMAN, B., KAUTZ, H., AND COHEN, B. 1996. Local search strategies for satisfiability testing. In D. S. Johnson and M. A. Trick, editors, *Cliques, coloring, and satisfiability: Second DIMACS implementation challenge*. (October 11–13, 1993), Providence, RI. American Mathematical Society.
- STANLEY, R. P. 1997. *Enumerative Combinatorics*, Vol. 1. Cambridge University Press, Cambridge, UK.
- TATIKONDA, S., AND JORDAN, M. I. 2002. Loopy belief propagation and Gibbs measures. In *Proceedings of UAI*. 493–500.
- WAINWRIGHT, M. J., AND JORDAN, M. I. 2003. Graphical models, exponential families, and variational methods. Technical report, UC Berkeley, Department of Statistics, No. 649. Preprint at URL: <http://www.eecs.berkeley.edu/~wainwrig/Papers/WaiJorVariational03.ps>.
- WAINWRIGHT, M. J., AND MANEVA, E. 2005. Lossy source encoding via message-passing and decimation over generalized codewords of LDGM codes. In *Proceedings of IEEE International Symposium on Information Theory*. IEEE Computer Society Press, Los Alamitos, CA, 1493–1497.
- WANG, J. 1997. Average case computational complexity theory. In *Complexity Theory Retrospective*, L. Hemaspaandra and A. Selman, Eds., Vol. II, Springer-Verlag. 295–328.
- WELLING, M., AND TEH, Y. 2001. Belief optimization: A stable alternative to loopy belief propagation. In *Proceedings of UAI*. 554–561.
- YEDIDIA, J., FREEMAN, W. T., AND WEISS, Y. 2005. Constructing free energy approximations and generalized belief propagation algorithms. *IEEE Trans. Inf. Theory* 51, 7, 2282–2312.
- YEDIDIA, J. S., FREEMAN, W. T., AND WEISS, Y. 2003. Understanding belief propagation and its generalizations. In *Exploring Artificial Intelligence in the New Millennium*, Science and Technology books, New York, Chapt. 8, 236–239.
- YUILLE, A. 2002. CCCP algorithms to minimize the Bethe and Kikuchi free energies: Convergent alternatives to belief propagation. *Neural Computat.* 14, 1691–1722.