

Multi-Objective 3-SAT with Survey-Propagation

Cyril Furtlehner, Marc Schoenauer

► **To cite this version:**

Cyril Furtlehner, Marc Schoenauer. Multi-Objective 3-SAT with Survey-Propagation. NIPS 2010 Workshop on Discrete Optimization in Machine Learning: Structures, Algorithms and Applications (DISCML), Dec 2010, Whisler, Canada. 2010. <inria-00533149>

HAL Id: inria-00533149

<https://hal.inria.fr/inria-00533149>

Submitted on 5 Nov 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Multi-Objective 3-SAT with Survey-Propagation

Cyril Furtlehner*
cyril.furtlehner@inria.fr

Marc Schoenauer†
marc.schoenauer@inria.fr

Abstract

An original approach to multi-objective optimization is introduced, using a message-passing algorithm to sample the Pareto set, i.e. the set of Pareto-non-dominated solutions. Several heuristics are proposed and tested on a simple bi-objective 3-SAT problem. The first one is based on a straightforward deformation of the Survey-Propagation (SP) equation to locally encode a Pareto trade-off. A simple heuristic is then tested, which combines an elimination procedure of clauses with the usual decimation of variables used in the SP algorithm, and is able to sample different regions of the Pareto-front. We study in more details the compliance of these deformed equations with basic Belief-Propagation (BP) properties. This first leads to an explicit Markov Random Field (MRF) of valid warning configuration, for which the SP equations are basic BP equations. This observation is then generalized to the multi-objective context. Numerical experiments on artificial problems up to 10^5 variables are presented and discussed.

1 Introduction

Message passing algorithms based on the Belief-Propagation (BP) have proved very successful in particular in combinatorial optimization, to solve for example the random K-SAT problem with SP [1]. However, it is well-known that most real-world problems are in fact multi-objective and we are not aware of any work addressing multi-objective problems with message-passing algorithms. The goal of this paper is to extend the message-passing algorithms strategy from single to multi-objective context in the constraint satisfaction domain. The aim of multi-objective optimization is to sample the Pareto set, i.e. the set of solutions that are not dominated in the Pareto sense in the decision space, and the Pareto front, i.e. the corresponding points in the objective space (each objective being a coordinate). The Pareto dominance relation defines a partial order on the decision space: a solution a dominates a solution b if a is better than b on at least one criterion, without being worse on any other. The knowledge of (a good approximation of) the Pareto set allows the user to make an informed decision, knowing exactly what an increase on a given objective will cost in terms of the other objectives. For combinatorial optimization problems, message passing heuristics can be set up in principle, once a uniform measure is defined on the set of solutions, typically in the form of an MRF. Our guiding principle then for addressing the multi-objective context is to search for a MRF, accurately approximating the Pareto set and at the same time suitable to run message passing algorithms. The paper is organized as follows: in section 2 we give a brief introduction to the SP algorithm and underlying assumptions. In section 3 we define the benchmark problem and discuss how the Pareto dominance can be inserted locally into the SP equations. In section 4 we discuss the compliance of these equation with BP and how the Pareto front can be estimated on single problem instances. In section 5, a simple heuristic based on the modified equations is presented along with numerical results.

*INRIA-Saclay, LRI Bât. 490, F-91405 Orsay(France)

†INRIA-Saclay, LRI Bât. 490, F-91405 Orsay(France)

2 Random 3-SAT Problems and SP

The 3-SAT problem is a decision problem involving a set \mathcal{V} of N binary decision variables $(x_i)_{i \in [1, N]}$ (*FALSE* or *TRUE*), subject to a conjunction of a set \mathcal{F} of M constraints or clauses. Defined in conjunctive normal form the problem reads, $C_{\mathcal{V}, \mathcal{F}} = \bigwedge_{a=1}^M C_a(x_a)$ where $x_a = \{x_i, x_j, x_k\}$ is a subset of \mathcal{V} , and clause C_a appears as the disjunction of three literals where each literal corresponds to a negated or non-negated variable. The clause is *SAT* if at least one of its literal is *TRUE*. The clause density $\alpha \stackrel{\text{def}}{=} M/N$ measures the difficulty of the problem. The random *SAT* is a family of problems indexed by this control parameter, a given instance being obtained by taking at random the subset x_a of variables attached to any given clause a and the sign of each literal is also taken at random. The phase diagram of random K-SAT has been determined and refined over the years mainly with help of mean-field considerations [1, 2, 3, 4]. Various clustering phenomena taking place in the solutions space give its structure to the phase diagram. Schematically for 3-SAT, in the thermodynamic limit, a sharp *SAT-UNSAT* transition is occurring at $\alpha = \alpha_c \simeq 4.267$ (the probability for the problem to be *SAT* drops discontinuously from 1 to 0); for $\alpha \leq 3.86$ there is a *SAT* phase, corresponding to a giant cluster of nearby solutions, while the domain $\alpha \in [3.86, 4.267]$ is referred to as the hard SAT phase, the space of solutions being fragmented into distant clusters.

The SP equations [5] assume a 1-RSB phase in which solutions are grouped into well-separated clusters, these clusters being parametrized (presumably in a non-unique way) by a set of binary variables $w_{a \rightarrow i} \in \{0, 1\}$ called warnings, attached to each link relating a clause a to a variable i on the factor graph [6]. When a variable receives such a message it should adopt the value requested by the clause sending this message. A given configuration of warnings is valid iff no variable receives contradictory warnings; a clause sends a warning to one of its neighbours if its other neighbors received incompatible warnings with the requirement of that clause. Fixing in a self-consistent way the values of these warnings is actually equivalent to run BP algorithm on a MRF associated to SAT assignments [2]. Let $J_{ai} \in \{-1, 1\}$ say whether a variable x_i is negated (-1) or not ($+1$) in clause a and let $\tau_{aib} \stackrel{\text{def}}{=} \frac{1+J_{ai}J_{bi}}{2} \in \{0, 1\}$ indicate if clause a and b have compatible requirements ($\tau_{aib} = 1$) or not ($\tau_{aib} = 0$) w.r.t. variable i . The self-consistency rule for the warnings reads:

$$w_{a \rightarrow i} = \prod_{j \in a \setminus i} \Pi_{j \rightarrow a}^-(w_j) \quad (1)$$

with $w_j = \{w_{b \rightarrow j}, b \ni j\}^1$, and

$$\Pi_{i \rightarrow a}^+ \stackrel{\text{def}}{=} \prod_{b \ni i \setminus a} (\bar{w}_{b \rightarrow i} + \tau_{aib} w_{b \rightarrow i}) - \Pi_{i \rightarrow a}^0, \quad \Pi_{i \rightarrow a}^- \stackrel{\text{def}}{=} \prod_{b \ni i \setminus a} (\bar{w}_{b \rightarrow i} + \bar{\tau}_{aib} w_{b \rightarrow i}) - \Pi_{i \rightarrow a}^0, \quad (2)$$

with $\Pi_{i \rightarrow a}^0(w_i) \stackrel{\text{def}}{=} \prod_{b \ni i \setminus a} \bar{w}_{b \rightarrow i}$. In the hard SAT phase, this schema is actually not working because of the clustering of the solutions. The SP algorithm finds a uniform measure on the valid warning assignments by propagating instead the probability $\eta_{a \rightarrow i} \stackrel{\text{def}}{=} P(w_{a \rightarrow i} = 1)$, called the survey. Assuming probabilistic independence of warnings sent to a given variable, the SP equation then reads,

$$\eta_{a \rightarrow i} = \prod_{j \in a \setminus i} \frac{\Pi_{j \rightarrow a}^-(\eta_j)}{\Pi_{j \rightarrow a}^0(\eta_j) + \Pi_{j \rightarrow a}^+(\eta_j) + \Pi_{j \rightarrow a}^-(\eta_j)}. \quad (3)$$

where again η_j is the set of surveys received by j . The denominator here corresponds to a conditioning on non-contradictory warnings under the independent law defined by the set of surveys.

The fixed-point solution can then be used to simplify SAT formulas by fixing the most polarized variable. Iterating this procedure constitutes the SP-decimation algorithm, which ends when the fixed point degenerates to all surveys being identically zero. At this point the reduced problem is expected to be very easy to solve with a local search algorithm. In the *UNSAT* phase, the problem (aka *MAXSAT*) is instead to find configurations with lowest possible number of violated clauses. The SP equations are structurally different in that case, they involve a real parameter y , but still lead to an efficient distributed algorithm called SP-Y [7].

¹ $j \in b$ is a shorthand notation expressing that j is neighbour to b

3 SP Deformed Equations with Local Pareto Constraints

The bi-objective benchmark problem that we consider in this paper consists simply in having two sets of clauses \mathcal{F}_0 and \mathcal{F}_1 instead of a single one, while keeping a single set \mathcal{V} of variables. For simplicity \mathcal{F}_0 and \mathcal{F}_1 are taken to be of equal size $M/2$ with $M/N < 2\alpha_c$, which means that each sub-problem $(\mathcal{V}, \mathcal{F}_\mu)$, $\mu \in \{1, 2\}$ taken independently is in the *SAT* phase while the junction of the two $(\mathcal{V}, \mathcal{F} = \mathcal{F}_0 + \mathcal{F}_1)$ is in the *UNSAT* one.

To adapt the SP equations to this multi-objective context we consider the Pareto dominance relation between solutions at the local level, by comparing two solutions separated by a single variable flip: we can say that a variable is Pareto optimal if under a flip it cannot increase the number of *SAT* clauses of one objective without strictly increasing the number of *UNSAT* clauses for the other one. With the chosen value of the clause density, each sub-problem taken alone can be made *SAT*, henceforth the Pareto set contains solutions for which one of the 2 sub-problem is *SAT*. This leads us to consider the ensemble of valid warning configuration in which a variable cannot receive contradictory warning emitted from the same sub-problem. We are looking for warning configurations which may have mutual conflicts between sub-problems, but for which internal conflicts, i.e., contradictory warnings send by clauses pertaining to the same sub-problem, are excluded. Then a variable may be in three different situations which all imply a local Pareto equilibrium:

- the variable is unconstrained, it does not receive any warning and can take either *TRUE* or *FALSE* value without modifying any of the objective.
- the variable receives at least one warning but without any contradiction, so that it takes the value obeying to the warnings.
- the variable receives at least one warning from \mathcal{F}_0 and \mathcal{F}_1 and these are contradictory. In that case the variable can chose to conform to either \mathcal{F}_0 or \mathcal{F}_1 . Under a flip, one of the sub problem will lose at least one *SAT* clause while the other will gain at least one.

To cope with this new specification, considering first the quantity $\Pi_{i \rightarrow a}^c$ associated to these type of contradictions:

$$\Pi_{i \rightarrow a}^c(w_i) = \Pi_{i \rightarrow a}^{+-}(w_i) + \Pi_{i \rightarrow a}^{-+}(w_i),$$

where $\Pi_{i \rightarrow a}^{+-}$ and $\Pi_{i \rightarrow a}^{-+}$ represent warning configuration where i receive at least one a -compatible warning from sub-problem \mathcal{F}_a containing a and one a -incompatible warning from the other sub-problem $\bar{\mathcal{F}}_a$, or vice versa, and are expressed as products of $\Pi_{i \rightarrow a}^\pm$ (2) restricted to these clause subsets. A variable, submitted to two incompatible requests from the two sub-problems has now the freedom to choose to which one it obeys. Averaging over this choice induces anyway some correlations between warnings which are difficult to handle, so we fix from the beginning the choice that will take each variable in case of a contradiction. Let $\theta_i \in \{0, 1\}$ represent this binary choice. The rule to send a warning is then determined by the relation

$$w_{a \rightarrow i} = \prod_{j \in a \setminus a} (\Pi_{j \rightarrow a}^-(w_j) + \bar{\theta}_{ai} \Pi_{j \rightarrow a}^c(w_j)), \quad (4)$$

where $\theta_{ai} = \theta_a \theta_i + \bar{\theta}_a \bar{\theta}_i \in \{0, 1\}$ if $\theta_a \in \{0, 1\}$ gives the appartenance set \mathcal{F}_a of a .

The SP equations are then adapted as follows, by taking into account conflicting sets of warning but with some penalization factor $q \in [0, 1]$:

$$\eta_{a \rightarrow i} = \prod_{j \in a \setminus i} \frac{\Pi_{j \rightarrow a}^-(\eta_j) + q \bar{\theta}_{ai} \Pi_{j \rightarrow a}^c(\eta_j)}{\Pi_{j \rightarrow a}^0(\eta_j) + \Pi_{j \rightarrow a}^+(\eta_j) + \Pi_{j \rightarrow a}^-(\eta_j) + q \Pi_{j \rightarrow a}^c(\eta_j)}. \quad (5)$$

Given the set of surveys one has in principle access to a certain number of quantities. In particular the complexity Σ , i.e. the log number of clusters can serve to tune q such that Σ remains slightly positive. Also the probability P_a^v of each clause to be violated, which decompose into one contribution coming from the environment of the clause, and the second term represents the direct impact of the clause, causing some new variables to be under contradiction (see [8] for details). This quantity, ΔP_a^v will be useful when trying to identify which clauses are the most difficult to satisfy.

4 BP compliance

The equations presented so far, although having simple rules suffers from an important drawback which we describe now. Some compatibilities between surveys, at the basis of the BP schema are not satisfied, this preventing us from an exact evaluation of P_a^v as well as Σ , and henceforth a reliable estimation of the Pareto front. This motivates a closer investigation of the compliance of these equations with the basic BP equations. This question has been addressed in various ways for SP, first in [9], using a dual formulation on an extended factor graph and in [10] by introducing the notion of cover. We describe here another connection holding directly at the level of warnings [8].

4.1 The Case of SP

Consider that the attribute of a variable node i involved in the factor graph representation of this problem is the set of messages $w_i \stackrel{\text{def}}{=} \{w_{a \rightarrow i}, a \ni i\}$, while those of the factor nodes are the set of incoming warning on variables attached to a , namely $w_a \stackrel{\text{def}}{=} \{w_{b \rightarrow i}, i \in a, b \ni i\}$. The MRF associated to the uniform measure of valid warning configurations is then given by:

$$P(w) = \frac{1}{Z} \prod_{a \in \mathcal{F}} C_a(w_a) \prod_{i \in \mathcal{V}} C_i(w_i) \quad (6)$$

with

$$\begin{aligned} C_i(w_i) &\stackrel{\text{def}}{=} \Pi_i^0(w_i) + \Pi_i^+(w_i) + \Pi_i^-(w_i), \\ C_a(w_a) &\stackrel{\text{def}}{=} \prod_{i \in a} \bar{w}_{a \rightarrow i} (\Pi_{i \rightarrow a}^0 + \Pi_{i \rightarrow a}^+ + \Pi_{i \rightarrow a}^-) - \prod_{i \in a} \bar{w}_{a \rightarrow i} \Pi_{i \rightarrow a}^- \\ &\quad + \sum_{i \in a} w_{a \rightarrow i} (\Pi_{i \rightarrow a}^0 + \Pi_{i \rightarrow a}^+) \prod_{j \in a \setminus i} \bar{w}_{a \rightarrow j} \Pi_{j \rightarrow a}^-. \end{aligned}$$

$C_a(w_a)$ is defined in such a way to encode the rule (1) for emitting or not a message. Note that configurations in which a clause emits more that one message are excluded. Running the BP on this MRF results in the following message update rule:

$$m_{a \rightarrow i}(w_i) \propto \sum_{w_a/w_i} C_a(w_a) \prod_{j \in a/i} n_{j \rightarrow a}(w_j) \quad (7)$$

$$n_{i \rightarrow a}(w_i) = C_i(w_i) \prod_{b \ni i/a} m_{b \rightarrow i}(w_i) \quad (8)$$

This BP schema is well defined but potentially heavy because variables w_i are d_i -dimensional Boolean vectors, if d_i is the connectivity of i . A direct relation to SP is obtain from the following lemma [8].

Lemma 4.1. *Let $\eta_{a \rightarrow i} \in [0, 1]$, the update rule (7) is stable with respect to the following parametrization of the message,*

$$\begin{aligned} m_{a \rightarrow i}(w_i) &= w_{a \rightarrow i} (\Pi_{i \rightarrow a}^0(w_i) + \Pi_{i \rightarrow a}^+(w_i)) \eta_{a \rightarrow i} \\ &\quad + \bar{w}_{a \rightarrow i} (\Pi_{i \rightarrow a}^0(w_i) + \Pi_{i \rightarrow a}^+(w_i) + \Pi_{i \rightarrow a}^-(w_i)) (1 - \eta_{a \rightarrow i}). \end{aligned} \quad (9)$$

with $\eta_{a \rightarrow i}$ satisfying the SP update rules.

This insure in particular that the BP-based entropy formula is correct as well as the various probabilities associated to variables and factor nodes obtained from the surveys.

4.2 Generalization

In our bi-objective context, the MRF associated to the uniform measure of valid warning configurations is given by (6) with

$$C_i(w_i) \stackrel{\text{def}}{=} \Pi_i^0(w_i) + \Pi_i^+(w_i) + \Pi_i^-(w_i) + q \Pi_i^c(w_i)$$

and $C_a(w_a)$ enforcing the self-consistent rules (4) used to send or not a warning. The minimal parametrization of the messages to cope with this MRF involves now 3 real independent messages, instead of a single one for SP. These are the probability coefficients required to account for the 4 relevant states of a variable in this case, whether it receives a warning or not from a ($w_{a \rightarrow i} = 1$ or $w_{a \rightarrow i} = 0$) and whether it is forced to contradict a or not. (see table 1). We denote by $x_{a \rightarrow i}$,

	$w_{a \rightarrow i} = 0$	$w_{a \rightarrow i} = 1$
i SAT a	$x_{a \rightarrow i}$	$z_{a \rightarrow i}$
i UNSAT a	$y_{a \rightarrow i}$	$t_{a \rightarrow i}$

Table 1: Different states of variable i w.r.t clause a and associated surveys

$y_{a \rightarrow i}$, $z_{a \rightarrow i}$ and $t_{a \rightarrow i}$ the associated probabilistic surveys with $x_{a \rightarrow i} + y_{a \rightarrow i} + z_{a \rightarrow i} + t_{a \rightarrow i} = 1$ and let $\Pi_{i \rightarrow a}^x(w_i)$, $\Pi_{i \rightarrow a}^y(w_i)$, $\Pi_{i \rightarrow a}^z(w_i)$ and $\Pi_{i \rightarrow a}^t(w_i)$ the corresponding indicator functions on these states (see [8] for details). The clause constraints with help of these operators take the form:

$$C_a(w_a) = \prod_{i \in a} \Pi_{i \rightarrow a}^x + \sum_{i \in a} \Pi_{i \rightarrow a}^y \prod_{j \in a/i} \Pi_{j \rightarrow a}^x + \prod_{i \in a} \Pi_{i \rightarrow a}^t + \sum_{i \in a} \Pi_{i \rightarrow a}^z \prod_{j \in a/i} \Pi_{j \rightarrow a}^y \quad (10)$$

and the message is parametrized in the following way:

$$m_{a \rightarrow i}(w_i) = x_{a \rightarrow i} \Pi_{i \rightarrow a}^x(w_i) + y_{a \rightarrow i} \Pi_{i \rightarrow a}^y(w_i) + z_{a \rightarrow i} \Pi_{i \rightarrow a}^z(w_i) + t_{a \rightarrow i} \Pi_{i \rightarrow a}^t(w_i).$$

We do not explicit here the update rule of the surveys, which are a bit more involved than in the preceding deformed SP equations, but again, once a set of surveys satisfying these equations is found, quantities like the probability P_a^v for a clause to be violated can be computed exactly. These expression can be used to estimate the Pareto front of a given problem instance by computing the expected number

$$\mathcal{E}_\mu(q, \theta) \stackrel{\text{def}}{=} \mathbb{E}[\#\text{UNSAT}_\mu] = \sum_{a \in \mathcal{F}_\mu} P_a^v, \quad \mu \in \{0, 1\}$$

of *UNSAT* clause for each sub-problem, given the penalty q and a set $\theta = \{\theta_i, i \in \mathcal{V}\}$ of binary choices. Therefore, for a given choice of (q, θ) we can compute its corresponding estimate $(\mathcal{E}_0, \mathcal{E}_1, \Sigma)$. The set of non-dominated parameters choice regarding $(\mathcal{E}_0, \mathcal{E}_1)$ and for which $\Sigma \geq 0$ constitutes the estimation of the Pareto front corresponding to the 4-surveys equations.

5 Numerical experiments

Numerical experiments have been run using the deformed SP equation (5) to find Pareto solutions and compared them with *MAXSAT* solution obtained with SP-Y. The procedure is as follows:

- (i) **clause elimination**: based on ΔP_a with highest value, a small set of clauses are successively selected to be taken aside from the problem. *Nelim*, the total number of eliminated clauses is fixed a priori, and in practice the best results are obtained with a lower value than the one required to make the problem *SAT*.
- (ii) **variable decimation**: as in the original SP algorithm, the variables with the highest polarization are fixed sequentially, until the problem becomes paramagnetic or until convergence is lost.
- (iii) **resolution of the reduced problem** *WALKSAT* is run on the reduced problem a given number of times in order to generate a cloud of solutions.

During both the elimination and the decimation stages, the penalty q is maintained at convergence threshold. The position of the solution found on the Pareto front depends on how the clauses are selected in the elimination procedure. Here, it is implicitly determined by the choice of θ_i for each variable i in Eq. (5) before letting SP converge. Among many possible heuristics, the one giving the best results so far consists in eliminating n_0 clauses from problem \mathcal{F}_0 by letting $\theta_i = 1$ uniformly, and then to flip to $\theta_i = 0$ uniformly to eliminate $n_1 = \text{Nelim} - n_0$ clauses from objective \mathcal{F}_1 . The comparison with SP-Y is made by running SP-Y with backtracking and with different values of the pseudo inverse temperature y around the optimal y^* for which the complexity vanishes. For $\alpha < 4.4$, the Pareto front which is obtained with the best value for *Nelim* is not far from being optimal in the *MAXSAT* region, and scales up gracefully with the problem size N (see Figure. 1). However, the performance degrades when α increases [8]. Ideally, on Figure. 1, the Pareto front should enter the region below the Gardner energy [7], which is not the case yet. The 4-surveys equations of section 4.2 yield the Pareto-front estimate shown on Figure.1-left.

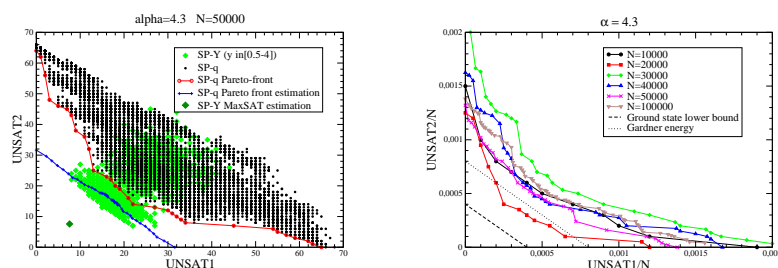


Figure 1: The Pareto front obtained with deformed SP equations compared with solutions obtained from SP-Y, along with the Pareto front estimation and the SP-Y maxsat estimate (left). Rescaled Pareto-front obtained when N is varied at $\alpha = 4.3$ (right).

6 Conclusion and perspectives

This paper has investigated the possibility to associate a MRF to the Pareto set of a specific multi-objective problem (3 -SAT). An original method to estimate this set, based on message passing algorithm, has been proposed, as well as a simple heuristic that is able to sample the Pareto set with reasonably good performance. Still, a gap with an optimal performance remains which may have several possible origins: a simplifying assumption in the MRF definition; an additional simplification in the SP equation; no backtracking techniques during the elimination and decimation stage of the procedure. These issues require additional work to be tackled, but more generally from this study we can foresee how our approach could be generalized to more than two objectives and other type of multi-objective combinatorial problems. Any consistent message passing equations (in the sense discussed in section 4.2), like the one used in SP-Y or the one proposed in section 4.2, contains useful information on the clauses, which could be used in principle during the elimination procedure. This procedure seems quite efficient in our case, but could probably be improved with backtracking techniques. Additionally, the proposed approach is completely generic for constraint satisfaction problems at large, so the idea would be to use it in combination with single objective message passing optimizers, as a basic tool to sample Pareto front of multi-objective constraint satisfactions problems with many different objectives.

References

- [1] M. Mézard and R. Zecchina. The random K -satisfiability problem: from an analytic solution to an efficient algorithm. *Phys.Rev.E* 66, page 56126, 2002.
- [2] S. Mertens, M. Mézard, and R. Zecchina. Threshold values of random k -sat from the cavity method. *Random Structures and Algorithms*, 28:340–373, 2006.
- [3] A. Montanari, G. Parisi, and F. Ricci-Tersenghi. Instability of one-step replica-symmetry-broken phase in satisfiability problems. *J. Phys. A*, 37:2073, 2004.
- [4] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, and L. Zdeborová. Gibbs states and the set of solutions of random constraint satisfaction problems. *PNAS*, 104:10318–10323, 2007.
- [5] Alfredo Braunstein, Marc Mézard, and Riccardo Zecchina. Survey propagation: an algorithm for satisfiability. *CoRR*, cs.CC/0212002, 2002.
- [6] F. R. Kschischang, B. J. Frey, and H. A. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Trans. on Inf. Th.*, 47(2):498–519, 2001.
- [7] D. Battaglia, M. Kolár, and R. Zecchina. Minimizing energy below the glass thresholds. *Phys. Rev. E*, 70:036107, 2004.
- [8] C. Furtlehner and M. Schoenauer. Multi-objective 3-sat problems with message passing techniques. Technical Report 7424, Inria, 2010.
- [9] A. Braunstein and R. Zecchina. Survey propagation as local equilibrium equations. *J. Stat. Mech.*, page P06007, 2004.
- [10] E. Maneva, E. Mossel, and M. J. Wainwright. A new look at survey propagation and its generalizations. *Journal of the ACM*, 54(4):2–41, 2007.