Trust Estimation in autonomic networks: a message passing approach

Stefano Ermon *

* *Department of Computer Science, Cornell University, Ithaca NY 14850 USA (e-mail: ermonste@cs.cornell.edu).*

Abstract: The ability to maintain belief relationship among entities in autonomic networks is considered a major challenge. In this work we tackle the problem by casting it into the framework of Estimation Theory as an inference problem on a Markov Random Field. A fully distributed algorithm based on message passing techniques is then proposed, where messages are not considered as abstract intermediate results of a computation, but as real messages exchanged by the nodes in the network.
With this case study we therefore demonstrate that Markov Random Field theory used in combination with Message Passing algorithms constitutes a powerful theoretical framework for the development of algorithms for information distribution and fusion.

Keywords: Autonomic Networks, Trust management, Estimation, Belief Propagation

1. INTRODUCTION

Interest in the development of decentralized self-managing and self-configuring networks, broadly referred to as autonomic networks, is surging both from a research and an industrial perspective. This new research field, aimed both at increasing robustness and at reducing the need of human intervention in the management and the deployment effort, is drawing ideas from several existing disciplines, including network management, artificial intelligence, control and game theory.

The lack of a pre-defined and fixed infrastructure with a centralized control and the highly dynamic nature of autonomic systems pose a number of new challenges ahead. Security in particular represents a critical issue, thus it is not surprising that a major effort of the networking community is currently devoted at defining and introducing security services into these new communication architectures. As pointed out in Selcuk et al. (2004); Blaze et al. (1996); Abdul-Rahman and Hailes (1998), the most important security challenge within the autonomic network paradigm is that of establishing the trustworthiness status of the nodes in the network. In fact since the overall performance of an autonomic network depends on the collaborations that take place among self-managing entities, it is fundamental to develop suitable models for establishing and maintaining trust relationships between them. As proposed in Sun et al. (2006), in this context we will broadly interpret trust as a belief relationship, where an entity is confident that another one will operate fairly, or as it is designed.

Since trust relationships are essential to predict the future behavior of other entities, a trust management system greatly influences both the specification of security policies and the effectiveness of the decision-making of most other protocols. In fact with a trust management system potential damage caused by malfunctioning or even malicious entities can be greatly reduced, mainly because most entities will avoid interacting with them, or at least they will do it in a cautious way.

The lack of infrastructure and of any authoritative entity in the network enforces the use of reputation-based systems, where trust is established by protocols that try to evaluate the previous behavior of the entities. In a setting where nodes have no prior knowledge of each other, trust information is obtained solely through continual self-monitoring and evaluation of past interactions. The individual knowledge gathered in this way is then distributed throughout the network and processed in a distributed manner.

Despite the growing importance of this problem, most state of the art trust management systems, such as those proposed in Selcuk et al. (2004); Theodorakopoulos and Baras (2004); Sun and Yang (2007); Abdul-Rahman and Hailes (1998); Michiardi and Molva (2002); Venkatraman et al. (2000); Buchegger (2002), are still mostly at an empirical level. As it is pointed out in Sun et al. (2006) and in Langheinrich (2003), most of the work on trust management in the literature is essentially based on heuristics and on simulation as evaluation method. The validation of the proposed systems is often an overlooked aspect, where not all solutions are actually verified and almost none are implemented and tested in a real environment. In this context, theoretical analysis is extremely rare and the comparison between different methods is therefore difficult to accomplish, mainly because of the great simulative effort that would be required. Solutions are often hard to compare even on a simulative basis, since they often rely on different hypothesis and are aimed at different application scenarios.

⋆ This research was in part supported by an NSF Expeditions in Computing grant 0832782.

Given the importance of a theoretical framework, the aim of this work is to provide a deeper understanding of the problem through a more mathematically sound approach. In particular for the sake of tractability we will focus on a non adversarial setting where all entities collaborate in the identification of malfunctioning nodes, with no malicious entities acting to disrupt the process. A practical example of such a setting is a Wireless Sensor Network where faulty sensors need to be recognized as unreliable by the entities they are interacting with and a malfunctioning node that is providing inaccurate measurements can perform a self-diagnosis only by querying its neighbors about the quality of its own measurements.

In the following sections we tackle the problem by casting it into the framework of Estimation Theory, as an inference problem on a Markov Random Field. A fully distributed algorithm based on message passing techniques is then proposed, where messages are not considered as abstract intermediate results of a computation, but as real messages exchanged by the nodes in the network.

2. THE TRUST ESTIMATION PROBLEM FORMULATION

The model we use to describe how trust information is obtained through the evaluation of previous interactions between entities, first proposed in Ermon et al. (2009), represents the starting point for the extensions presented in the following sections.

In our model, we consider a network consisting of N nodes, represented by a directed graph $G = (V, E)$, with $|V| = N$ and where edges connect entities that can communicate and therefore are assumed to be interacting.

A bit variable $T_i \in \{-1, 1\}$, representing a real trustworthiness status, is associated to each node i. We can therefore describe the trust status of the entire network with a *real trust vector* $T \in \{-1, 1\}^N$, adopting the convention

$$T_i = \begin{cases} 1 & \text{if node } i \text{ is trustworthy} \\ -1 & \text{otherwise} \end{cases}$$

The complete *real trust vector* T is unknown to the entities in the network, nonetheless it is in many cases useful to estimate it to improve the reliability and the performance of the network, as it is outlined in the introduction.

Even if T is unknown, nodes in the network can gather some evidence about it on the basis of the history of their previous interactions with their neighbors. In fact in this model we assume that the future behavior is statistically correlated with the past one, as well as with the real trustworthiness status of the nodes. In particular we assume that T is time invariant and it is related to an *opinion matrix* $C \in \mathbb{R}^{N \times N}$, a collection of random variables c_{ij}, representing respectively the opinion that node i has on node j based on their previous interactions. We therefore assume that the following equation holds

$$C = f(T, \omega), \qquad \omega \in \Omega \tag{1}$$

where Ω is a sample space and $f(\cdot)$ represents the way in which opinions among nodes are formed. Given the constraint that opinions are formed on the basis of the previous interactions with other nodes, we assume that c_{ij} is significant only if i and j are neighbors, and we set it to 0 otherwise. The outlined framework is very general and, with slight modifications, can be easily applied to many cases of practical interest, such as Wireless Sensor Networks and MANETs, by choosing a suitable function f.

In the following sections we will show how to design an algorithm that estimates T given the evidence of the opinions c_{ij}, with the remarkable property that it works in a distributed way, so that in each iteration only local information is used, without the need of any central coordination. The optimality of the algorithm is also guaranteed at least for certain network topologies.

2.1 The Gaussian case

For reasons of mathematical tractability, in the rest of the paper we will mainly consider the following particular case of equation (1):

$$c_{ij} = \begin{cases} T_i T_j + w_{ij} & \text{if } (i,j) \in E \\ 0 & \text{if } (i,j) \notin E \end{cases} \tag{2}$$

where $w_{ij} \sim \mathcal{N}(0, \sigma^2)$ is a Gaussian random variable that captures the uncertainty that affects the way in which opinions are formed.

It is easy to see that there is a symmetry in the behavior of the nodes, because $T_i T_j = 1$ both for a pair i, j of trustworthy nodes and for a pair of untrustworthy nodes. One possible interpretation is that untrustworthy nodes act in a similar way, since they tend to have good opinions of other nodes that should not be trusted.

The trust estimation problem consists in finding a trust vector $\hat{T} \in \{-1, 1\}^N$ that is a good estimate of the unknown *real trust vector* T, given the evidence of an *opinion matrix*. The most natural approach is to search for the configuration that is more likely to have generated a certain observed *opinion matrix* $\overline{C}$, or in other words the trust configuration with the highest a posteriori probability, given $\overline{C}$.

The likelihood $LH(S; \overline{C})$ of any configuration S given an *opinion matrix* $\overline{C}$ is by definition:

$$LH(T; \overline{C}) := p(T|\overline{C})$$

where $p(T|\overline{C})$ is the probability of T conditioned that $C = \overline{C}$. We can easily define the maximum likelihood estimator $g : \mathbb{R}^{N \times N} \to \{-1, 1\}^N$ as

$$g(\overline{C}) := \arg\max_{\hat{T}} LH(\hat{T}; \overline{C})$$

Observe that the Bayes rule yields

$$p(T|C) = \frac{p(C|T)p(T)}{p(C)}$$

where $p(T)$ is the a priori probability of the discrete random variable $T \in \{-1, 1\}^N$ while $p(C)$ and $p(C|T)$ are the density and conditional density of the continuous random variable $C \in \mathbb{R}^{N \times N}$. This shows that

$$g(\overline{C}) = \arg\max_{\hat{T}} p(\overline{C}|\hat{T})p(\hat{T})$$

because $p(C)$ does not depend on $\hat{T}$.

For the Gaussian model described in (2), assuming independence, we have that $p(C|T) = 0$ if C has a nonzero

entry in position $(i,j) \notin E$. If instead C has nonzero entries only in $(i,j) \in E$, then

$$p(C|T) = \prod_{(i,j)\in E} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(c_{ij}-T_iT_j)^2}{2\sigma^2}}$$
$$= q(C) e^{\frac{1}{\sigma^2}\sum_{(i,j)\in E} T_iT_jc_{ij}}$$

where $q(C)$ is a normalization constant independent of T.

Therefore maximizing $p(C|T)$ is equivalent to maximize $U(T;C) := \sum_{(i,j)\in E} T_iT_jc_{ij}$, so that

$$\arg\max_{\hat{T}} p(\overline{C}|\hat{T}) = \arg\max_{\hat{T}} U(\hat{T};\overline{C})$$

Suppose that the a priori probability distribution is Bernoulli-distributed with parameter p, namely

$$p := P(T_i = 1)$$

Then, assuming independence, we obtain that

$$p(T) = \gamma e^{-\lambda \sum_i T_i} \tag{3}$$

where γ normalizes to a probability distribution and

$$\lambda = -\frac{1}{2}\log\left(\frac{p}{1-p}\right)$$

Putting all together we obtain

$$LH(T;C) = q(C) e^{\frac{1}{\sigma^2}\sum_{(i,j)\in E} T_iT_jc_{ij}} \gamma e^{-\lambda\sum_i T_i}$$

$$= \gamma q(C) e^{\frac{1}{\sigma^2}(\sum_{(i,j)\in E} T_iT_jc_{ij} - \lambda\sigma^2\sum_i T_i)}$$

We conclude that the following proposition holds:

Proposition 1. The likelihood $LH(T;C)$ of a configuration T is proportional to a monotonic increasing function of

$$H(T) := \sum_{(i,j)\in E} T_iT_jc_{ij} - \eta\sum_i T_i \tag{4}$$

where $\eta = \lambda\sigma^2$.

We can therefore compute a maximum likelihood estimate of the *real trust vector* T by setting

$$\eta = \lambda\sigma^2 = -\frac{\sigma^2}{2}\log\left(\frac{p}{1-p}\right)$$

and by maximizing (4) over all possible configurations T.

3. STATISTICAL PHYSICS INTERPRETATION

3.1 The Ising Model

Equation (4) is very important because it represents the energy or Hamiltonian of a configuration S in an Ising Model (see Sherrington and Kirkpatrick (1975) for a detailed analysis of the model) in the presence of an external magnetic field of strength η that breaks the symmetry of the system. As the physical interpretation of the system confirms, when the a priori probability distribution of T is symmetrical, that is $p = 0.5$, the magnetic field disappears and the system becomes completely symmetrical.

Systems described by an Hamiltonian such as 4, where each variable T_j represents the direction of a spin variable, are more generally known in the statistical physics literature as Spin Glasses. Since the c_{ij} are modeled as random variables, systems in this class exhibit randomly distributed ferromagnetic and anti ferromagnetic interactions between spins, depending on the sign of the coupling coefficient c_{ij}. They represent the first studied class of systems with frustrated behavior, where the presence of conflicting interactions forbids simultaneous minimization of the interaction energies and hence the existence of a trivial global ground state.

The original problem of maximum likelihood estimation is easily reduced through proposition 1 to the problem of finding the maxima of (4). This optimization problem has been proved to be NP-Complete for generic graphs in Cipra (2000), and hence an exhaustive search for global maxima is widely believed to be computationally intractable.

An interesting approach to tackle the problem that we propose here is to embed the topology of the communication network directly into a graphical model so that we can make use of the classic message passing inference algorithms. The striking feature is that these techniques involve a set of messages that travel on the graphical model, that thanks to our symmetry with the communication network we can interpret as real messages exchanged by the entities in the network, thus naturally defining a protocol in which entities operate in a decentralized and fully distributed way.

3.2 Pairwise Random Markov Field

In addition to the Spin Glass interpretation, proposition 1 enables us to interpret the model in a more general setting as a Pairwise Random Markov Field.

A Markov Random Field (MRF) is a graphical model that captures the statistical dependence of several random variables $x_1, \ldots, x_N$ by the means of a graph, where each node is associated to a random variable and edges represent probabilistic dependency relationships. In particular in a Pairwise Markov Random Field the following factorization property is assumed to hold:

$$p\{x\} = p\{x_1, \ldots, x_N\} = \frac{1}{Z}\prod_{(i,j)} \psi_{ij}(x_i, x_j) \prod_i \phi_i(x_i) \tag{5}$$

and it is said to be pairwise because the overall probability distribution $p\{x\}$ is factored into two-variable dependencies $\psi_{ij}(x_i, x_j)$ relative to edges (i,j) in the graphical model.

The model described in section 2, as well as the general Ising model, is indeed a particular case of a Pairwise Markov Random Field. In fact we can define a graphical model by taking the undirected version of the communication graph $G = (V, E)$ and associating the random variable $x_i = T_i|C$ with alphabet $\{1, -1\}$ to each node. The factorization property defined by equation 5 is satisfied by choosing

$$\ln\psi_{ij}(x_i, x_j) = \frac{c_{ij} + c_{ji}}{\sigma^2} x_ix_j = J_{ij}x_ix_j \tag{6}$$

(notice that the graph of the graphical model is assumed to be undirected) and

$$\ln\phi_i(x_i) = -\lambda x_i \tag{7}$$

so that when $\{x\} = \{T|C\}$ we have that

$$p\{x\} = \frac{1}{Z} e^{\frac{H(T)}{\sigma^2}},$$

where the Hamiltonian $H(\cdot)$ is the same of equation 4.

4. MESSAGE PASSING ALGORITHMS

4.1 Estimating Marginals with Belief Propagation

After casting the model into the MRF framework, the first problem we address is that of computing marginals, or beliefs, defined as the a posteriori probabilities $p\{T_i|C\}$ of a single random variable $T_i|C$. In physical terms, the problem is analogous to computing *local magnetization* vectors

$$M_i = p\{T_i = 1|C\} - p\{T_i = -1|C\} = 1 - 2p\{T_i = -1|C\}$$

in the corresponding Ising model.

This task can be efficiently accomplished using Belief Propagation(BP), one of the best known algorithms for performing inference on graphical models. This algorithm has recently had a lot of attention in fields such as AI, computer vision, control and even coding theory, where it has been successfully used to decode turbocodes under the name of the sum product algorithm.

A key aspect of the algorithm for the application scenario proposed here is that it manages to distribute the global computation of the marginals into smaller local computations, whose results travel on the graphical model in the form of messages that are combined appropriately by the recipients. Remarkably, instead of considering these messages as abstract partial results of an algorithm, we can think of them as real messages exchanged by the nodes in the network, obtaining a powerful theoretical framework for the distribution and fusion of information among the nodes of a network.

According to BP, the nodes in the network try to reach an agreement on what their probability distributions should be by exchanging messages. Variables $m_{ij}(x_j)$ are introduced, where $m_{ij}(x_j)$ represents a message sent from node i to node j about what state node j should be into. The message is a vector with as many components as the alphabet of x_j, and each component intuitively represents how likely it is according to i that node j is in the corresponding state.

The belief $b_i(x_i) = p\{T_i = x_i|C\}$ at a node i is obtained as

$$b_i(x_i) = k\phi_i(x_i) \prod_{j \in N(i)} m_{ji}(x_i)$$

where k normalizes to a probability distribution. The messages are updated according to the following rule

$$m_{ij}(x_j) \leftarrow \sum_{x_i} \left(\phi_i(x_i)\psi_{ij}(x_i, x_j) \prod_{k \in N(i)\setminus\{j\}} m_{ki}(x_i) \right)$$

By using equations 7 and 6 the BP updating rules for our model become

$$b_i(x_i) = ke^{-\lambda x_i} \prod_{j \in N(i)} m_{ji}(x_i)$$

and

$$m_{ij}(x_j) \leftarrow \sum_{x_i} \left(e^{-\lambda x_i} e^{J_{ij} x_i x_j} \prod_{k \in N(i)\setminus\{j\}} m_{ki}(x_i) \right)$$

It is well known (for a proof see Pearl (1988)) that these rules give exact beliefs if the pairwise Markov Random Field is singly connected, that is if the underlying graph is a tree so that there are no loops. However the algorithm is well defined on any graph topology, even if there is no guarantee on the convergence and on the quality of the solution found.

4.2 Maximum a posteriori likelihood estimation with the Max-product algorithm

The problem of computing a maximum a posteriori likelihood estimate of the *real trust vector* T outlined in section 2 can be addressed with an algorithm very similar to Belief Propagation. This algorithm, known as the max-product algorithm, is once again a local message passing algorithm that works by exchanging messages among nodes, so that at every iteration, each node sends a message to each of its neighbors and receives one from each of them. The messages are defined as follows:

$$m_{ij}(x_j) \leftarrow \max_{x_i} \left(\psi_{ij}(x_i, x_j)\phi_i(x_i) \prod_{k \in N(i)\setminus\{j\}} m_{ki}(x_i) \right)$$

while the max marginals p_i are computed according to the following equation

$$p_i(x_i) = \phi_i(x_i) \prod_{j \in N(i)} m_{ji}(x_i)$$

By using equations 7 and 6 the updating rules for the max product algorithm become

$$m_{ij}(x_j) \leftarrow \max_{x_i} \left(e^{J_{ij} x_i x_j} e^{-\lambda x_i} \prod_{k \in N(i)\setminus\{j\}} m_{ki}(x_i) \right) \quad (8)$$

$$p_i(x_i) = e^{-\lambda x_i} \prod_{j \in N(i)} m_{ji}(x_i) \quad (9)$$

Each component of the vector $p_i(x_i)$ represents the probability of the most likely trust assignment to the entire network, when the trust status of node i is forced to be x_i. When the graph G is a tree, it is known that the algorithm is guaranteed to converge to a unique fixed point, such that it is possible to obtain the global most likely (a posteriori) configuration $\hat{T}$, that is defined component-wise by

$$\hat{T}_i = \arg\max_{x_i} p_i(x_i).$$

Unfortunately a tree shaped communication network is not a common scenario, but exactly as Belief Propagation the max product algorithm is well defined on any graph topology. However, when applied to a network with loops, it might not converge, and even when it does, there is no guarantee on the quality of the results obtained. Even if one can indeed find pathological examples of graphical models in which message passing algorithms fail, they have been successfully used on loopy graphical models in many applications arisen in coding theory, computer vision and medical diagnosis. In particular in Murphy et al. (1999) it is shown empirically that loopy Belief Propagation does converge to good approximations of the marginals in a

wide range of real world inference problems. Moreover in the simulative analysis presented in the next section we never experienced convergence problems, even in the case of a complete graph that intuitively represents a worst case scenario from the point of view of the presence of loops.

5. ANALYSIS

From a qualitative point of view, we can start the analysis by noticing that we cannot expect any topology-independent result. For example, in a network made by isolated vertices, we cannot do any better than just using the a priori knowledge. We will therefore need to fix a topology to be able to show some meaningful results.

Since we know that loops are a cause of trouble for message passing algorithms, we will focus our attention on a complete communication graph. Even if it is not representative of the topology of any interesting real world network, it should intuitively be close to a worst case scenario. Moreover most analytical results from Spin glasses theory are derived for this topology, and making use of them it is possible to show the following result:

Proposition 2. If $\eta \neq 0$

$$\lim_{N\to\infty} \mathbb{E}\left[\frac{h(\hat{T})}{N}\right] = 0$$

where

$$\hat{T} := \text{argmax}_{S\in\{1,-1\}^N} H(S)$$

and

$$h(\hat{T}) := |\{i : \hat{T}_i \neq T_i\}|$$

is the number of incorrect estimates given by $\hat{T}$.

A proof of proposition 2 can be found in Ermon et al. (2009), and it essentially descends from a non-trivial spin glass theory result that guarantees the convergence of the sequence

$$N^{-\frac{3}{2}}\mathbb{E}\left[\max_{S\in\{1,-1\}^N} \sum w_{ij}s_i s_j\right]$$

as N tends to infinity when the couplings w_{ij} are Gaussian distributed. In the context of estimation theory, proposition 2 essentially states that the maximum likelihood estimator is non-biased in the limit of large N, regardless of which method is used to actually compute it.

5.1 Simulative Analysis

From a simulative point of view, we are interested in measuring what is the fraction of nodes that the max-product algorithm is not able to correctly identify, in expectation. If S^* is the configuration returned by the max product algorithm, we are interested in the average error rate

$$\mathbb{E}\left[\frac{||S^* - T||_1}{2N}\right]$$

where the expectation is taken over all levels of randomness. The first experiment is performed by simulating the environment described by the Gaussian model presented in section 2.1, for various values of N and σ^2. The estimation algorithm used is the max product algorithm, where messages are initialized as constant functions and $O(N^2)$ messages are exchanged in total.

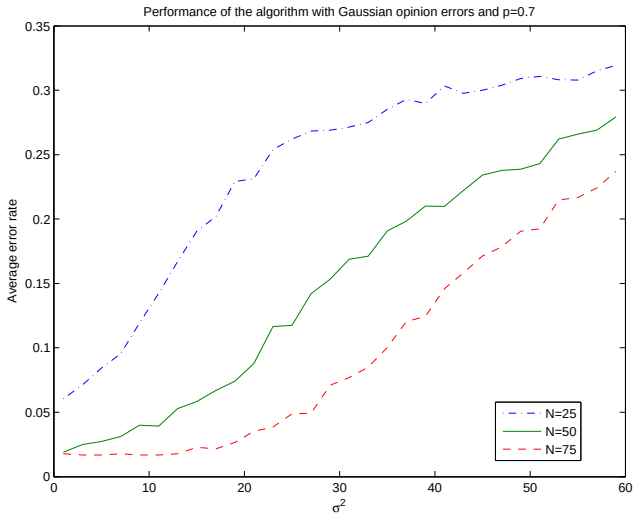

Fig. 1. Performance of the algorithm with a complete communication graph of N nodes for several values of N. The a priori probability p that a node is trustworthy is 0.7.

As we can note in figure 1, the algorithm does actually converge even in presence of loops, and moreover a number of exchanged messages polynomial in the size of the network suffices to reach a fixed point.

As one might expect, the performance of the algorithm decreases as does the quality of the a posteriori information (measured by a larger variance σ^2 on the opinions). As we can see from the case $N = 25$, when the a posteriori information is too noisy, the error rate becomes higher than the one obtained by the optimal estimator that is based solely on the a priori information S^*_{ap}:

$$S^*_{ap} = \begin{cases} [1,\dots,1] & \text{if } p > 0.5 \\ -[1,\dots,1] & \text{otherwise} \end{cases},$$

that clearly shows an average error rate of $(1-p)$.

Moreover the simulation data show that the error rate decreases as N grows, thus confirming proposition 2.

To test the robustness of the algorithm we consider another reasonable model for (1), where the errors are Bernoulli distributed. In particular we assume that if $(i,j) \in E$ then

$$c_{ij} = \begin{cases} T_iT_j & \text{with probability } 1-p_e \\ -T_iT_j & \text{with probability } p_e \end{cases} \quad (10)$$

In this model when a node is trustworthy ($T_i = 1$), $c_{ij} = T_j$ with probability $1 - p_e$, while the contrary holds when $T_i = -1$. Thus the parameter p_e is an error probability, representing how likely it is for a trustworthy node to misjudge a neighbor.

The results obtained with various error probabilities p_e and various number of nodes N are shown in figure 2. The trust estimation algorithm uses a value of

$$\sigma^2 = \mathbb{E}[(c_{ij} - T_iT_j)^2] = 4p_e \quad (11)$$

and it shows a low error rate at least until p_e approaches 0.4. The results are comparable with those obtained with model (2), when the variance of the error on the opinions is the same according to equation (11). However when $p_e > 0.5$, on average there are more wrong opinions than correct ones, and the algorithm is outperformed by the one

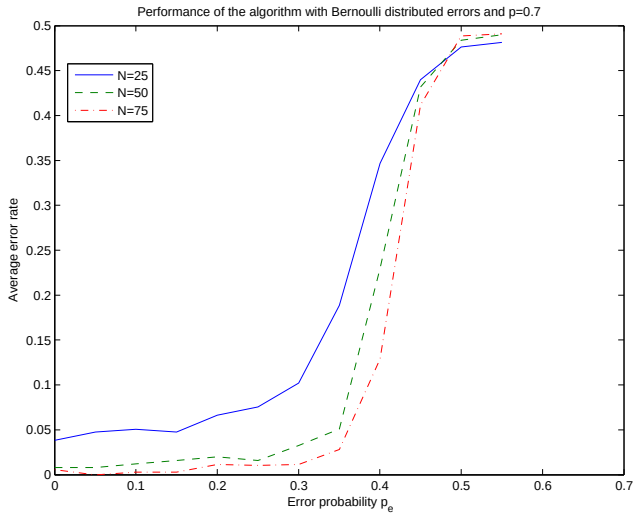

Fig. 2. Performance of the algorithm with a complete communication graph of N nodes for several values of N and opinions generated according to model (10). The a priori probability p that a node is trustworthy is 0.7.

based solely on the a priori information. The average error rate shows a sharp phase transition phenomenon around $p_e = 0.4$, that is typical of spin glasses systems. In the extreme and somewhat unrealistic situation where $p_e > 0.5$, it is possible to take advantage of the symmetry of the problem by considering $-C$ as the *opinion matrix*, thus obtaining the same accuracy as with an error probability of $1 - p_e$ and a curve symmetric with respect to $p_e = 0.5$.

6. CONCLUSIONS

The new paradigm of autonomic networks poses new challenges ahead, due to its self-managing, self-configuring and highly dynamic nature. Among all of them, in this paper we focused on the trust management system, that is arguably one of the most interesting ones.

In this work we presented a mathematically sound framework for trust evaluation based on Markov Random Field theory, and we proposed the use of a fully distributed algorithm based on message passing techniques. This algorithm is completely based on local interactions between nodes and can be implemented without any need for central coordination and demonstrates that Markov Random Field theory used in combination with Message Passing algorithms constitutes a powerful theoretical framework for the development of algorithms for information distribution and fusion.

REFERENCES

Abdul-Rahman, A. and Hailes, S. (1998). A distributed trust model. In *Proceedings of the 1997 workshop on New security paradigms*, 48–60. ACM New York, NY, USA.

Blaze, M., Feigenbaum, J., and Lacy, J. (1996). Decentralized trust management. In *1996 IEEE Symposium on Security and Privacy, 1996. Proceedings.*, 164–173.

Buchegger, S. (2002). Performance analysis of the CONFIDANT protocol. In *Proceedings of the 3rd ACM Int. Symp. on Mobile ad-hoc networking & computing*, 226–236. ACM New York, NY, USA.

Cipra, B. (2000). The Ising model is NP-complete. *SIAM News*, 33(6).

Ermon, S., Schenato, L., and Zampieri, S. (2009). Trust estimation in autonomic networks: a statistical mechanics approach. *48th IEEE Conference on Decision and Control, 2009*.

Langheinrich, M. (2003). When trust does not compute-the role of trust in ubiquitous computing. In *Workshop on Privacy at UBICOMP*.

Michiardi, P. and Molva, R. (2002). Core: a collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks. In *Advanced Communications and Multimedia Security: Sixth Joint Working Conference on Communications and Multimedia Security, 2002, Slovenia*, 107. Kluwer Academic Publishers.

Murphy, K., Weiss, Y., and Jordan, M. (1999). Loopy belief propagation for approximate inference: An empirical study. In *Proceedings of Uncertainty in AI*, 467–475.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann.

Selcuk, A., Uzun, E., and Pariente, M. (2004). A reputation-based trust management system for P2P networks. In *IEEE Int. Symp. on Cluster Computing and the Grid, CCGrid'04*, 251–258.

Sherrington, D. and Kirkpatrick, S. (1975). Solvable model of a spin-glass. *Physical review letters*, 35(26), 1792–1796.

Sun, Y., Han, Z., Yu, W., and Liu, K. (2006). A trust evaluation framework in distributed networks: Vulnerability analysis and defense against attacks. In *Proc. of IEEE Infocom*.

Sun, Y. and Yang, Y. (2007). Trust establishment in distributed networks: Analysis and modeling. In *IEEE International Conference on Communications, 2007. ICC'07*, 1266–1273.

Theodorakopoulos, G. and Baras, J. (2004). Trust evaluation in ad-hoc networks. In *Proceedings of the 3rd ACM workshop on Wireless security*, 1–10. ACM New York, NY, USA.

Venkatraman, M., Yu, B., and Singh, M. (2000). Trust and reputation management in a small-world network. In *Proceedings of Fourth Int. Conf. on MultiAgent Systems*, 449–450.