RC 22229 (W0111-013) November 2, 2001
Computer Science

IBM Research Report

Efficient fault diagnosis using local inference

Irina Rish, Mark A. Brodie, Haiqin Wang, Sheng Ma
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Efficient fault diagnosis using local inference

Irina Rish, Mark **Brodie, Haiqin** Wang, **Sheng** Ma
IBM T.J. Watson Research Center, 30 Saw Mill River Road, Hawthorne, NY 10532
rish,brodie,hwang,shengma@us.ibm.com

Abstract

This paper proposes a Bayesian network approach to fault diagnosis in distributed computer systems which use test transactions (*probes*). A lower bound on the MPE diagnosis accuracy is derived and analyzed with respect to the problem parameters. An empirical study of a local-inference approximation scheme (algorithm approx-mpe(1)) yields useful insights on the algorithm's behavior and its applicability to the considered class of diagnosis problems: the approximation quality is higher for higher-MPE ("higher-quality") networks and "degrades gracefully" with noise.

1 Introduction

Tbe increasing complexity of distributed computer systems makes their performance management more and more challenging, and requires efficient inference techniques for various diagnosis and prediction tasks. In this paper, we consider fault diagnosis in distributed systems using test transactions, or *probes*. A distributed system can be represented by a "dependency graph", where nodes can be either hardware elements (e.g., workstations, servers, routers) or software components/services, and links can represent both physical and logical connections between the elements (see Figure la). A probe is a specific transaction (e.g., *ping* or *traceroute* command, test mail message, or web-page access request), sent from a workstation ("probing station") to a server (or router) which can test a particular service (i.e, IP-connectivity, database- or web-access). A set of probe outcomes (e.g., response times or return codes) can be used to diagnose the states of system components (e.g., fault/no fault). Given the probe outcomes, the diagnosis task is to find the most-likely set of states of the elements.

We use a graphical probabilistic framework of Bayesian networks [9] that provides a compact factorized representation for multivariate probabilistic distributions, and a convenient tool for probabilistic inference. Since the exact inference in Bayesian networks is generally hard (NP-hard) [1], we focus on approximation techniques. Tbe complexity of inference is usually associated with a large size of probabilistic dependencies recorded during inference (clique size, or induced width). Thus, a popular approximation approach is to restrict the complexity by focusing only on local interactions. We use a local *inference* approach known as *mini-bucket* approximation for finding MPE [3, 8][1], which controls the

'A closely related example is local belief propagation [9], a linear-time approximation which became a surprisingly effective state-of-the art technique in error-correcting coding [4].

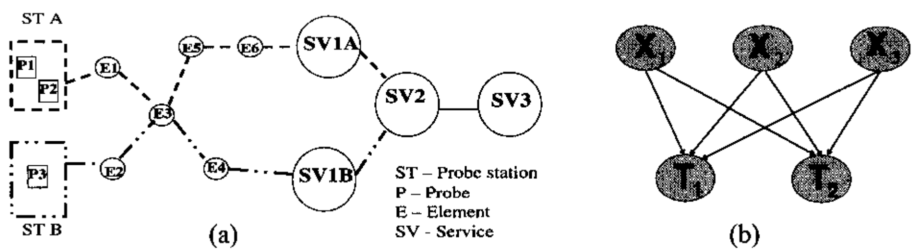

Figure 1: (a) An example of probing environment; (b) a two-layer Bayesian network structure for a set $\mathbf{X} = (X_1, X_2, X_3)$ of network elements and a set of probes $\mathbf{T} = (T_1, T_2)$.

accuracy/complexity trade-off by using a bound on the size of dependencies it creates. Particularly, its computational complexity is $O(n \cdot exp(i))$ where n is the number of nodes and i is the input parameter bounding the number of variables in a dependence [3].

In this paper, we (1) propose a noisy-AND Bayesian network formulation for fault diagnosis in distributed systems, (2) derive a lower bound on diagnosis accuracy for such networks, and (3) evaluate the performance of the mini-bucket approximation scheme with respect to network parameters, such as the level of noise. The accuracy of diagnosis is affected both by the quality of a model (diagnosis error based on the exact MPE solution) and by the accuracy of approximation to MPE. Our results suggest that the quality of approximation (1) is higher for higher-quality (i.e., higher-MPE) models, and (2) "degrades gracefully" with increasing noise. On the other hand, the computational complexity of the approximation used herein is linear in the number of nodes (mini-bucket scheme with $i = 1$ is used) versus exponential for exact inference. Thus, local inference appears to be a quite useful approach to handling diagnosis in large networks.

2 Fault diagnosis using probes: a noisy-AND Bayesian network

We now consider a simplified model of a computer network where each node (router, server, or workstation) can be in one of two states, 0 (fault) or 1 (no fault). The states of n network elements are denoted by a vector $\mathbf{X} = (X_1, ..., X_n)$ of *unobserved* boolean variables. Each probe, or test, T_j, originates at a particular node (probing workstation) and goes to some destination node (server or router). We also make an assumption that source routing is supported, i.e. we can specify the probe path in advance. A vector $\mathbf{T} = (T_1, ..., T_m)$ of *observed* boolean variables denoting the outcomes (0 - failure, 1 - OK) of m probes. Lower-case letters, such as x_i and t_j, denote the values of the corresponding variables, i.e. $\mathbf{x} = (x_1, ..., x_n)$ denotes a particular assignment of node states, and $\mathbf{t} = (t_1, .., t_m)$ denotes a particular outcome of m probes. We assume that the probe outcome is affected by *all nodes on its path*, and that node failures are marginally independent. These assumptions yield a causal structure depicted by a two-layer Bayesian network, such as one in Figure 1b. The joint probability $P(\mathbf{x}, \mathbf{t})$ for such network can be then written as follows:

$$P(\mathbf{x}, \mathbf{t}) = \prod_{i=1}^{n} P(x_i) \prod_{j=1}^{m} P(t_j | \mathbf{pa}(t_j)), \tag{1}$$

where $P(t_j|\mathbf{pa}(t_j))$ is the *conditional probability distribution (CPD)* of node T_i given the set of its *parents* $\mathbf{Pa_i}$, i.e. the nodes pointing to T_i in the directed graph, and $P(x_i)$ is the prior probability that $X_i = x_i$.

We now specify the quantitative part of those network, i.e. the CPDs $P(t_j|\mathbf{pa}(t_j))$. In general, a CPD defined on binary variables is represented as a k-dimensional table where $k = |Pa(t_j)|$. Thus, just the specification complexity is $O(2^k)$ which is very inefficient, if not intractable, in large networks with long probe path (i.e. large parent set). It seems reasonable to assume that each element on the probe's path affects the probe's outcome independently, so that there is no need to specify the probability of T_i for all possible value combinations of $X_{i_1}, ..., X_{i_k}$ (the assumption known as *causal independence* [6]). For example, in the absence of uncertainty, a probe fails if and only if at least one node on its path fails, i.e. $T_i = X_{i_1} \wedge ... \wedge X_{i_k}$, where $\wedge$ denotes logical AND, and $X_{i_1}, ..., X_{i_k}$ are all the nodes probe T_i goes through; therefore, once it is known that some $X_{i_j} = 0$, the probe fails independently of the values of other components. In practice, however, this relationship may be disturbed by "noise". For example, a probe can fail even though all nodes it goes through are OK (e.g., if network performance degradation leads to high response times interpreted as a failure). Vice versa, there is a chance the probe succeeds even if a node on its path is failed, e.g. due to routing change. Such uncertainties yield a *noisy-AND* model which implies that several causes (e.g., node failures) contribute independently to a common effect (probe failure)and is formally defined as follows:

$$P(t = 1|x_1, \ldots, x_k) = (1 - l) \prod_{x_i=0}^{n} q_i, \text{ and } P(t = 1|x_1 = 1, ..., x_k = 1) = 1 - l, \quad (2)$$

where l is the *leak probability* which accounts for the cases of probe failing even when all the nodes on its path are OK, and the *link probabilities*, q_i, account for the second kind of "noise" in the noisy-AND relationship, namely, for cases when probe succeeds with a small probability q_i even if node X_i on its path fails[2].

Once a Bayesian network is specified, the diagnosis task can be formulated as finding the *maximum probable explanation (MPE)*, i.e. a most-likely assignment to all X_i nodes given the probe outcomes, i.e.

$$\mathbf{x}^* = \arg\max_{\mathbf{x}} P(\mathbf{x}|\mathbf{t}) = \arg\max_{\mathbf{x}} P(\mathbf{x}, \mathbf{t}) = \max_{\mathbf{x}} \prod_{j=1}^{n} P(x_j) \prod_{i=1}^{m} P(t_i|\mathbf{pa}(t_i)). \quad (3)$$

An alternative approach is to look for the most likely value x_i^* of each node X_i separately, namely, to construct a diagnosis $\mathbf{x}' = (x_1', ..., x_n')$ where $x_i' = \arg\max_{x_i} P(x_i|\mathbf{t})$, $i = 1, ..., n$. Computing $\mathbf{x}'$ is sometimes easier than finding MPE, but generally, $x^* \neq \mathbf{x}'$.

When there is no noise in noisy-AND (i.e. leak and link probabilities are zero), the CPDs become deterministic, i.e. each probe $T_i = t_i$ imposes a constraint $t_i = x_{i_1} \wedge ... \wedge x_{i_k}$ on the values of its parent nodes $X_{i_1}, ..., X_{i_k}$. Now, finding an MPE can be viewed as a constrained optimization problem of finding $\mathbf{x}^* = \arg\max_{x_1,...,x_n} \prod_{j=1}^{n} P(x_j)$ subject to those constraints. In a particular case of uniform priors $P(x_j)$, diagnosis is reduced to solving a constraint satisfaction problem. Clearly, the quality of diagnosis depends on the set of probes: the only way to guarantee the correct diagnosis is to have a constraint set with a unique solution. This guarantee can only be achieved for $m \geq n$, since 2^m probe outcomes must "code" uniquely for 2^n node state assignments.

3 Accuracy of diagnosis

In this section, we derive a lower bound on the MPE diagnosis error. The *MPE error*, or *loss*, L_M, is the probability that diagnosis X^* differs from the true state X (by at least one

[2]Note that this noisy-AND definition is equivalent to the *noisy-OR* definition in [7] if we replace every value by its logical negation (all 0's will be replaced by 1's and vice versa). We also note that instead of considering the leak probability separately, we may assume there is an additional "leak node" always set to 0 that affects an outcome of a probe T_i according to its link probability $(1 - l_i)$.

bit). Given particular values $\mathbf{T} = \mathbf{t}$, $\mathbf{X} = \mathbf{x}$, and diagnosis $\mathbf{X}^* = \mathbf{x}^*$, we get $P(\mathbf{x} \neq \mathbf{x}*|\mathbf{t}) = I_{\mathbf{x}\neq\mathbf{x}*}$ where I_s is the *indicator function*, $I_s = 1$ if $s = true$ and $I_s = 0$ otherwise. Then the MPE error is

$$L_M = P(\mathbf{X} \neq \mathbf{X}^*|\mathbf{T}) = E_{\mathbf{x},\mathbf{t}} I_{\mathbf{x}\neq\mathbf{x}^*|\mathbf{t}} = E_{\mathbf{t}}(1 - P(\mathbf{x}^*|\mathbf{t})) = 1 - \sum_{\mathbf{t}} P(\mathbf{x}^*, \mathbf{t}), \quad (4)$$

where $\mathbf{x}^*$ is an MPE assignment, and E_z denotes expectation over z. Similarly, we can define the *bit error*, or bit loss, $L_b = P(X_i \neq X_i'|\mathbf{T})$ (clearly, the MPE error is generally higher than the bit error).

In the following, we assume equal fault priors $P(X_i = 0) = p$ for all nodes $i = 1, ..., n$; we also assume $p \leq 0.5$. Then

$$P(\mathbf{x}^*, \mathbf{t}) = \max_{\mathbf{x}} \prod_{j=1}^{n} P(x_j) \prod_{i=1}^{m} P(t_i|\mathbf{pa}(t_i)) \leq (1-p)^n \prod_{i=1}^{m} \max_{\mathbf{x}} P(t_i|\mathbf{pa}(t_i)). \quad (5)$$

The noisy-AND definition (expression 2) and the fact that $0 \leq q_i \leq 1$ yield $\max_{\mathbf{x}} P(t_i = 1|\mathbf{pa}(t_i)) = 1 - l_i$, and $\max_{\mathbf{x}} P(t_i = 0|\mathbf{pa}(t_i)) = 1 - \min_{\mathbf{x}} P(t_i = 1|\mathbf{pa}(t_i)) = 1 - (1 - l_i)\prod^{n}_{x_j \in \mathbf{pa}(t_i)} q_j$. Substitution of these two expressions in the expression 5 yields

$$P(\mathbf{x}^*, \mathbf{t}) \leq (1-p)^n \prod_{t_i=1} (1 - l_i) \prod_{t_i=0} [1 - (1 - l_i) \prod_{x_j \in \mathbf{pa}(t_i)}^{n} q_j]. \quad (6)$$

In order to simplify further derivation, we assume equal leak probabilities, $l_i = l$ for $i = 1, ..., m$, equal link probabilities $q_j = q$ for $j = 1, ..., n$, and equal parent set size (or probe route length) $|Pa(t_i)| = r$. Using the notation $P_k(\mathbf{x}^*, \mathbf{t})$ instead of $P(\mathbf{x}^*, \mathbf{t})$ where k is the number of $t_i = 1$ in $\mathbf{t}$, we get. Then expression 6 can be written as $P_k(\mathbf{x}^*, \mathbf{t}) \leq (1-p)^n \alpha^k (1 - \alpha q^r)^{m-k} = (1-p)^n \alpha^k \beta^{m-k}$, where $\alpha = 1 - l$ and $\beta = 1 - \alpha q^r$. Since there are $\binom{m}{k}$ vectors $\mathbf{t}$ having exactly k positive components t_i, we obtain

$$\sum_{\mathbf{t}} P(\mathbf{x}^*, \mathbf{t}) \leq (1-p)^n \sum_{k=0}^{m} \binom{m}{k} \alpha^k \beta^{m-k} = (1-p)^n (\alpha + \beta)^m. \quad (7)$$

Since $\alpha + \beta = (1 - l)(1 - q^r) + 1$ we finally get a *lower bound on MPE error* $\underline{L_M}$:

$$L_M = 1 - \sum_{\mathbf{t}} P(\mathbf{x}^*, \mathbf{t}) \geq 1 - (1-p)^n((1-l)(1-q^r) + 1)^m = \underline{L_M}. \quad (8)$$

Note that in the absence of noise (l=0 and q=0) we get $L_M \geq 1 - (1-p)^n 2^m$, thus, for uniform fault priors, $p = 0.5$, an error-free MPE diagnosis is only possible if $n = m$, as we noted before; however, for smaller p, zero-error can be achieved with smaller number of probes. Namely, solving $\underline{L_M} \leq 0$ for m yields the necessary condition for zero lower bound, $m \geq -n \frac{log(1-p)}{log(1+(1-l)(1-q^r))}$, plotted in Figure 2a as a function of p. Generally, solving $\underline{L_M} \leq 0$ for m provides a way of specifying the minimum necessary number of probes that yield zero lower bound for a specified values of other parameters[3].

Also, from the expression 8 we can see that the lower bound on the MPE diagnosis error is a monotone function of each parameter, n, m, p, l, q or r, given that other parameters are fixed. Namely, the error (bound) increases with increasing number of nodes n, fault probability p, leak probability l, and link probability q, but decreases with increasing number of probes m and probe route length r, which agrees with ones intuition that having more nodes on probe's path, as well as a larger number of probes, provides more information about the true node states. For example, the sensitivity of the error bound to noise is illustrated in Figure 2b: note a relatively sharp transition form 0 to 100%-error with increasing noise; it sharpness increases with increasing m and r.

[3]Clearly, finding a set of probes that may actually *achieve* the bound, if such set of probes exists, is a much harder task.

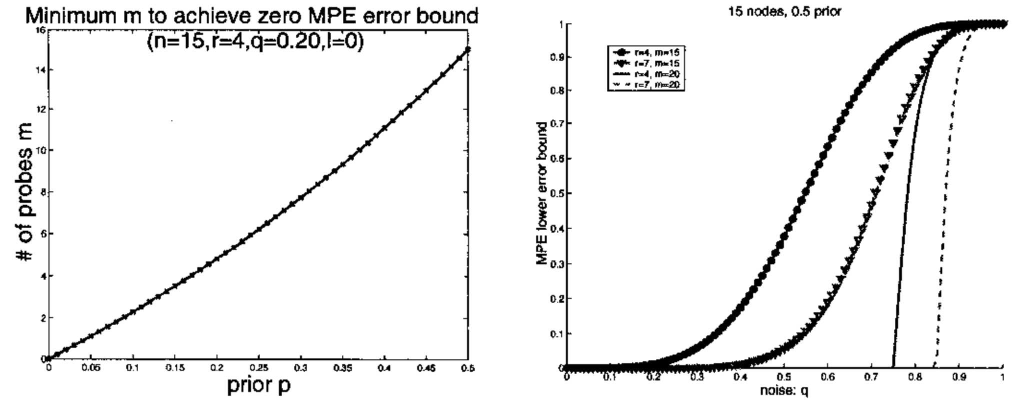

(a) (b)

Figure 2: (b) Minimum number of probes m to guarantee zero error bound, versus fault prior p: low prior yields lower than $n = 15$ number of probes; (b) lower bound on MPE error versus link probability q ("noise"): the longer the probes (higher r), and the more of them (higher m), the lower the MPE error bound for fixed noise q; also, the sharper the transition from 0 to 100% error.

4 Computational complexity of diagnosis and MPE approximations

Let us first consider the complexity of diagnosis in the absence of noise. Finding the most-likely diagnosis is reduced to constraint satisfaction in the following two cases. The first case is when the probe constraints allow exactly one solution (an assignment **x** simultaneously satisfying all constraints). The second case corresponds to the uniform priors $P(x_i)$ which also yield the uniform posterior probability $P(\mathbf{x}|\mathbf{t})$; therefore, any assignment **x** consistent with probe constraints is an MPE solution. Although constraint satisfaction is generally NP-hard, the particular problem induced by probing constraints can be solved in $O(n)$ time as follows.

Each successful probe yields a constraint $x_{i_1} \wedge ... \wedge x_{i_k} = 1$ which implies $x_i = 1$ for any node X_i on its path; the rest of the nodes are only included in constraints of the form $x_{i_1} \wedge ... \wedge x_{i_k} = 0$, or equivalently, $\neg x_{i_1} \wedge ... \wedge \neg x_{i_k} = 1$ imposed by failed probes. Thus, a $O(n)$-time algorithm assigns 1 to every node appearing on the path of a successful probe, and 0 to the rest of nodes. This is equivalent to *unit propagation* in *Horn theories*, which are propositional theories defined as a conjunction of clauses, or disjuncts, where each disjunct includes no more than one positive literal. It is easy to see that probe constraints yield a Horn theory and thus can be solved by unit propagation in linear time. Thus, finding MPE diagnosis is $O(n)$ time when it is equivalent to a constraint satisfaction in the absence of noise, as in cases of either uniform priors, or in case of unique diagnosis. In general, however, even in the absence of noise finding MPE is an NP-hard constrained optimization problem, with worst-case complexity $O(exp(n))$.

Similarly, in the presence of noise, finding MPE solution in a Bayesian network yields the complexity $O(expw^*)$ where w^* is the induced width of the network [2], i.e. the size of largest clique created by an exact inference algorithm, such as *variable elimination*. It is easy to show that $w^* \geq k$ where k is the maximum number parents of a probe node, and $w^* = n$ in the worst case[4].

[4]Algorithm *Quickscore*[5], specifically derived for noisy-OR networks, has the complexity $O(2^p)$ where p is the number of "positive findings" (failed probes in our case). However, the algorithm is tailored to belief updating and cannot be used for finding MPE.

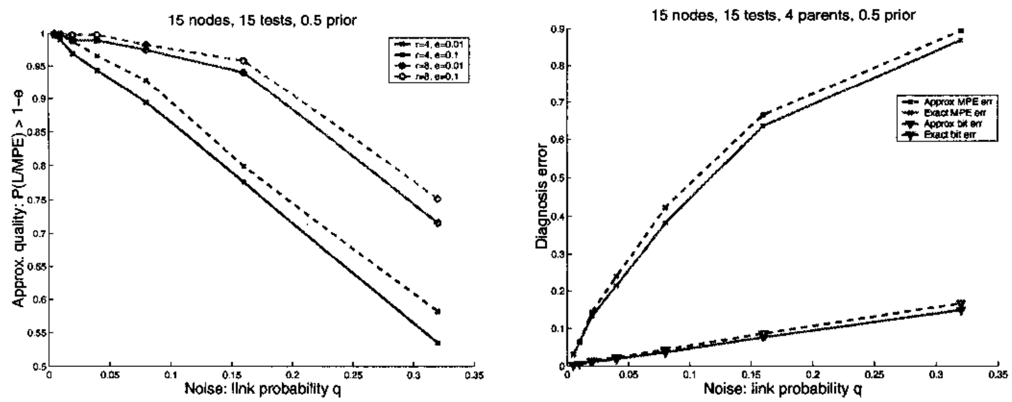

(a) (b)

Figure 3: (a) "Graceful degradation' of MPE approximation quality with noise, where the approximation quality is measured as $P(L/MPE) > 1 - e$ for $e = 0.01$ and $e = 0.1$: the quality is higher when the noise is smaller, and when the probe path is longer ($r = 8$ vs. $r = 4$); (b) MPE diagnosis error and bit diagnosis error for both exact and approximate diagnosis: for both MPE and bit errors, approximate results are very close to the exact ones.

Thus, we focused on approximating MPE, and studied empirically the algorithm $approx - mpe(i)$ (with $i = 1$, to be precise), which belongs to a family of the *mini-bucket* approximations for general constrained optimization, and particularly, for finding MPE [3, 10]. The idea of the mini-bucket approximation is to compute an upper bound on $MPE = \max_{\mathbf{x}} \prod_i P(x_i|\mathbf{pa}_i) \leq \prod_i \max_{x_i, \mathbf{pa}_i} P(x_i|\mathbf{pa}_i) = U$ and a lower bound L as a probability of an assignment $\mathbf{x}$ computed in a particular way, similarly to finding a solution to a constraint satisfaction problem after partial constraint propagation. Indeed, in deterministic case, the $approx - mpe(1)$ scheme is equivalent to arc-consistency in a constraint network, or unit propagation in propositional satisfiability [10] (increasing the parameter i corresponds to a more "coarse" partitioning of $P(x_i|\mathbf{pa}_i)$ into subproducts before maximization; e.g., $i = n$ yields the exact MPE computation).

We tested approx-mpe(1) on the networks constructed in a way that guarantees the unique diagnosis in the absence of noise (particularly, besides m probes each having r randomly selected parents, we also generated n additional probes each having exactly one parent node, so that all X_i nodes are tested directly). Since approx-mpe(1) is equivalent to unit propagation in the absence of noise, its diagnosis coincides must coincide with the MPE. Adding noise in a form of link probability q caused "graceful degradation' of the approximation quality, as shown in Figure 3a that plots the fraction of cases when the rationL/MPE was within the interval [1-e,1] for small values of e. This resulted into a diagnosis error very close to the error obtained by exact diagnosis, both for MPE error and bit error (Figure 3b). Also, as demonstrated in Figures 4a and b, there is a clear positive correlation between MPE value and approximation quality measured both as L/MPE (Figure 4a) and U/L (Figure 4a); there is also an interesting threshold phenomena: the approximation quality suddenly increases to practically perfect (L/MPE=1) once the MPE reaches a certain threshold value determined by the network parameters m, n, and r.

References

[1] G.F. Cooper. The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42(2–3):393–405, 1990.

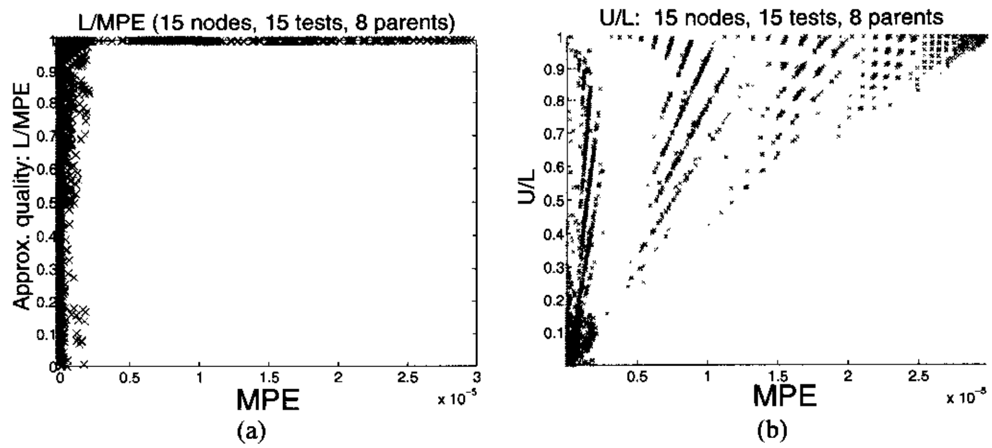

Figure 4: Approximation quality of algorithm approx-mpe(1) tends to be higher for higher MPE, i.e. for more likely diagnosis: a) L/MPE vs. MPE and b) U/L vs. MPE. The results are summarized for 30 "signals" per network, 30 random networks with $n = 15$ nodes, $n + n$ probes, $r = 8$ parents per probe, leak $l = 0$ and varying link q ("noise") from 0.005 to 0.64. Note: a) a sharp transition in approximation quality for $MPE \approx 2e - 6$; similar results observed for other networks, where the "transition point" is determined by parameters n, m and p; b) lower bound L is often more accurate than the upper bound U (U/L is far from 1 when L/MPE is near 1).

[2] R. Dechter and J. Pearl. Network-based heuristics for constraint satisfaction problems. *Artificial Intelligence*, 34:1–38, 1987.

[3] R. Dechter and I. Rish. A scheme for approximating probabilistic inference. In *Proc. Thirteenth Conf. on Uncertainty in Artificial Intelligence (UAI97)*, 1997.

[4] B.J. Frey and D.J.C. MacKay. A revolution: Belief propagation in graphs with cycles. *Advances in Neural Information Processing Systems*, 10, 1998.

[5] D. Heckerman. A tractable inference algorithm for diagnosing multiple diseases. In *Proc. Fifth Conf. on Uncertainty in Artificial Intelligence*, pages 174–181, 1989.

[6] D. Heckerman and J. Breese. Causal independence for probability assessment and inference using Bayesian networks. Technical Report MSR-TR-94-08, Microsoft Research, 1995.

[7] M. Henrion, M. Pradhan, B. Del Favero, K. Huang, G. Provan, and P. O'Rorke. Why is diagnosis using belief networks insensitive to imprecision in probabilities? In *Proc. Twelfth Conf. on Uncertainty in Artificial Intelligence*, 1996.

[8] K. Kask I. Rish and R. Dechter. Approximation algorithms for probabilistic decoding. In *Uncertainty in Artificial Intelligence (UAI-98)*, 1998.

[9] J. Pearl. *Probabilistic Reasoning in Intelligent Systems.* Morgan Kaufmann, 1988.

[10] I. Rish. *Efficient reasoning in graphical models.* PhD thesis, 1999.