

Accuracy vs. Efficiency Trade-offs in Probabilistic Diagnosis

Irina Rish and Mark Brodie and Sheng Ma

IBM T.J. Watson Research Center

30 Saw Mill River Road

Hawthorne, NY 10532

{rish,brodie,shengma}@us.ibm.com

Abstract

This paper studies the accuracy/efficiency trade-off in probabilistic diagnosis formulated as finding the *most-likely explanation (MPE)* in a Bayesian network. Our work is motivated by a practical problem of efficient real-time fault diagnosis in computer networks using test transactions, or *probes*, sent through the network. The key efficiency issues include both the cost of probing (e.g., the number of probes), and the computational complexity of diagnosis, while the diagnostic accuracy is crucial for maintaining high levels of network performance. Herein, we derive a lower bound on the diagnostic accuracy that provides necessary conditions for the number of probes needed to achieve an asymptotically error-free diagnosis as the network size increases, given prior fault probabilities and a certain level of noise in probe outcomes. Since the exact MPE diagnosis is generally intractable in large networks, we investigate next the accuracy/efficiency trade-offs for very simple and efficient *local* approximation techniques, based on variable-elimination (the mini-bucket scheme). Our empirical studies show that these approximations “degrade gracefully” with noise and often yield an optimal solution when noise is low enough, and our initial theoretical analysis explains this behavior for the simplest (greedy) approximation. These encouraging results suggest the applicability of such approximations to certain almost-deterministic diagnostic problems that often arise in practical applications.

Introduction

Accurate diagnosis of some unobserved states of the world from the outcomes of some measurements, or tests, is one of the most common problems occurring in practice. Numerous examples include medical diagnosis, computer troubleshooting, airplane failure isolation, noisy channel coding, and speech recognition, just to name a few. However, an accurate diagnosis often comes at a cost of large number of tests or a computationally expensive inference. Thus, our objective is to develop cost-efficient approaches to diagnosis that yield a good trade-off between the solution accuracy and the computational efficiency.

A practical motivation for our work is the problem of fault diagnosis in distributed computer systems by using a selected set of *probes*, or test transactions, that can be sent through the network in order to provide information about

its components (e.g., a probe can be *ping* or *traceroute* command, test email message, or web-page access request). A set of probe outcomes (e.g., response times or return codes) can be used to diagnose the states of system components. The rapid growth of distributed computer systems and networks in size and complexity makes fault diagnosis an increasingly challenging task, and requires extremely efficient inference techniques. The key efficiency issues include minimizing both the number of probes (tests) and the computational complexity of diagnosis while maximizing its accuracy. From a theoretical prospective, we wish to investigate the achievable limits of diagnostic accuracy for given levels of noise and fault probabilities.

We approach the problem from a probabilistic prospective and use the graphical probabilistic framework called *Bayesian (belief) networks* (Pearl 1988) that provides a compact representation for multivariate probabilistic distributions and allows for efficient inference techniques. Given a set of observations (e.g., probe outcomes), the diagnosis problem can be formulated as finding the most-likely vector of states of all unobserved nodes (e.g., network components), called the *most-likely explanation (MPE)*. Clearly, there are alternative formulations of the diagnosis problem, such as, for example, finding the posterior probability distribution for each node and selecting the k most-likely faults. However, in this paper, we focus on the MPE formulation, which is in our view a more general approach that does not make any additional assumptions (e.g., about the number of faults). This approach is sometimes criticized as being too complex and computationally intractable; note, however, that updating the probability of a single node is also an NP-hard problem (Cooper 1990) that can be (in the worst case) as time consuming as finding an MPE. This is exactly why the main focus of this paper is on *approximation* techniques and their accuracy/efficiency trade-offs.

The complexity of inference is usually associated with large probabilistic dependencies recorded during inference (clique size, or *induced width*)(Dechter 1996). Thus, a popular approximation approach is to restrict the complexity by focusing only on local interactions. We investigated the performance of two *local inference* techniques, *greedy-mpe* and *approx-mpe(1)*, which are the simplest members of the parametric family of variable-elimination algorithms known as *mini-bucket* approximations (Dechter & Rish 1997; 2002).

The mini-bucket scheme is closely related to other local approximations, such as *iterative belief propagation (IBP)* and *generalized belief propagation (GBP)* algorithms (Yedidia, Freeman, & Weiss 2001), that recently became the state-of-the-art approximation techniques successfully used in practical applications such as probabilistic error-correcting coding (Frey & MacKay 1998). The mini-bucket algorithms can be viewed as simplified, non-iterative versions of those approaches. The reason for focusing on the mini-bucket scheme instead of those more popular techniques was our desire to gain a theoretical understanding of the empirical success of the mini-bucket scheme on problems with low-noise (nearly-deterministic) dependencies, reported both in this paper and in the existing literature (Dechter & Rish 1997; 2002). Also, we hope that a theoretical understanding of very simple local approximations can later provide us with insights for the analysis of more complicated iterative approximations¹.

In this paper, we (1) propose a Bayesian network formulation for the problem of fault diagnosis in distributed systems using probes; (2) derive a bound on the diagnosis accuracy and analyze it with respect to the problem parameters such as noise level and the number of tests, suggesting feasible regions when an asymptotic (with problem size) error-free diagnosis can be achieved; (3) evaluate empirically the performance of two efficient local approximation schemes with respect to network parameters, such as the level of noise; and (4) provide some theoretical explanation of the accuracy vs. noise relation. In summary, 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 is higher for higher-quality (i.e., higher-MPE) models, and that it "degrades gracefully" with increasing noise. On the other hand, the computational complexity of the approximation used here is linear in the number of nodes, instead of exponential for exact inference. We conclude that simple local inference algorithms provide a promising approach to handling large-scale diagnosis, especially for low-noise problems often encountered in various practical application areas.

The rest of the paper is structured as follows. The next section provides background information on Bayesian networks and defines the fault diagnosis problem. Then, in the subsequent section, a theoretical bound on the diagnosis accuracy is derived. Next, we analyze the complexity of diagnosis and study the approximation schemes. The paper concludes with a summary and discussion of related work.

Background and problem formulation

As a motivating application, we consider a particular problem of network fault diagnosis using probes, although most of our results can be applied in a more general setting. Let us assume simplified model of a computer network where

¹To the best of our knowledge, existing theoretical analysis of IBP and GBP properties is mostly focused on their convergence but does not address directly their behavior on low-noise problems (Weiss 2000; Yedidia, Freeman, & Weiss 2001).




Figure 1: 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 $\mathbf{T} = (T_1, T_2)$.

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 $\mathbf{X} = (X_1, \dots, X_n)$ 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 $\mathbf{T} = (T_1, \dots, T_m)$ 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. $\mathbf{x} = (x_1, \dots, x_n)$ denotes a particular assignment of node states, and $\mathbf{t} = (t_1, \dots, t_m)$ 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 shown in Figure 1. The network represent a joint probability $P(\mathbf{x}, \mathbf{t})$:

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

where $P(t_j | \mathbf{pa}(t_j))$ is the *conditional probability distribution (CPD)* of node T_j given the set of its *parents* \mathbf{Pa}_j , i.e. the nodes pointing to T_j in the directed graph, and $P(x_i)$ is the prior probability that $X_i = x_i$. Formally, a Bayesian network BN over a set of variables X_1, \dots, X_k is a tuple (G, P) where G is the directed acyclic graph encoding the independence assumptions of the joint distribution $P(\mathbf{X})$, and where $P = \{P(x_i | \mathbf{pa}(\mathbf{x}_i))\}$ is the set of all CPDs.

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 $k = |\mathbf{Pa}(t_j)|$. 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_j for all possible value combinations of X_{i_1}, \dots, X_{i_k} (the assumption known as *causal independence* (Pearl 1988; Heckerman & Breese 1995)). For example, in the absence of uncertainty, a probe fails if and only if at least one node on its path fails, i.e. $T_j = X_{i_1} \wedge \dots \wedge X_{i_k}$, where \wedge denotes logical AND, and X_{i_1}, \dots, X_{i_k} are all the nodes probe T_j 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:

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

$$P(t = 1 | x_1 = 1, \dots, x_k = 1) = 1 - l, \quad (3)$$

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 fails².

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. $\mathbf{x}^* = \arg \max_{\mathbf{x}} P(\mathbf{x} | \mathbf{t})$. Since $P(\mathbf{x} | \mathbf{t}) = \frac{P(\mathbf{x}, \mathbf{t})}{P(\mathbf{t})}$, where $P(\mathbf{t})$ does not depend on \mathbf{x} , we get $\mathbf{x}^* = \arg \max_{\mathbf{x}} P(\mathbf{x}, \mathbf{t})$.

Accuracy of diagnosis

In this section, we will derive a lower bound on the diagnosis error. Note that the bound is not tight in many cases, and defines only necessary conditions for the asymptotically error-free diagnosis. Identifying tighter error bounds and thus more accurate conditions for error-free MPE diagnosis appears to be a much harder problem and remains a direction for future work³.

The error of the MPE diagnosis, denoted Err_{MPE} , is defined (similarly to 0/1 classification error) as the probability of making a mistake, i.e. the probability $P(\mathbf{X} \neq \mathbf{X}^*(\mathbf{T}))$ that the diagnosis vector $\mathbf{X}^*(\mathbf{T})$ differs from the true state of unknown variables, a random vector \mathbf{X} (by its definition

²Note that this noisy-AND definition is equivalent to the *noisy-OR* definition in (Pearl 1988; Henrion *et al.* 1996) 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 $(1 - l_i)$.

³This problem can be viewed as an asymptotic error analysis of a *constrained code*, where the probe outcomes "encode" the node-state vector, but the encoding is restricted by the nature of probes (logical-AND functions of node states) and by the network topology constraints on the possible set of probes. Note that well-known information-theoretic result, the Shannon's limit (Shannon 1948; Cover & Thomas 1991), provides the asymptotic error for MAP decoding in case of unconstrained codes; extending this result to a particular type of a constrained code may not be straightforward.

as the most-likely explanation, $\mathbf{X}^*(\mathbf{T})$ is the deterministic function of the random probe vector \mathbf{T} , assuming a deterministic tie-breaking rule given multiple MPEs). Then, by the rule of total probability, we get

$$Err_{MPE} = P(\mathbf{X} \neq \mathbf{X}^*(\mathbf{T})) = \sum_{\mathbf{x}, \mathbf{t}} P(\mathbf{x} \neq \mathbf{x}^*(\mathbf{t}), \mathbf{x}, \mathbf{t}) = \sum_{\mathbf{x}, \mathbf{t}} P(\mathbf{x}, \mathbf{t}) P(\mathbf{x} \neq \mathbf{x}^*(\mathbf{t}) | \mathbf{t}).$$

Since $P(\mathbf{x} \neq \mathbf{x}^*(\mathbf{t}) | \mathbf{t}) = I_{\mathbf{x} \neq \mathbf{x}^*(\mathbf{t})}$, where I_s is the *indicator function* ($I_s = 1$ if $s = \text{true}$ and $I_s = 0$ otherwise), we obtain

$$Err_{MPE} = \sum_{\mathbf{x}, \mathbf{t}} P(\mathbf{x}, \mathbf{t}) I_{\mathbf{x} \neq \mathbf{x}^*(\mathbf{t})} = \quad (4)$$

$$\sum_{\mathbf{t}} (1 - P(\mathbf{x}^*(\mathbf{t}) | \mathbf{t})) = 1 - \sum_{\mathbf{t}} P(\mathbf{x}^*(\mathbf{t}), \mathbf{t}). \quad (5)$$

From now on, we will use \mathbf{x}^* as a shorthand for $\mathbf{x}^*(\mathbf{t})$.

Assumptions. In order to simplify our further analysis, we will use the following notation. We denote as p the maximum prior probability over all nodes and their possible states, i.e. $p = \max_i \max\{P(X_i = 0), 1 - P(X_i = 0)\}$. For example, we may assume that all nodes have same prior probabilities, and that $p > 0.5$ is the probability of being in the OK state ($p = P(X_i = 1)$). Also, we denote as α_0 the maximum conditional probability value (over all test variables and their corresponding parent assignments) attained when a test outcome is 0, namely, $\alpha_0 = \max_i \max_{\mathbf{pa}(t_i)} P(t_i = 0 | \mathbf{pa}(t_i))$. Similarly, we define $\alpha_1 = \max_i \max_{\mathbf{pa}(t_i)} P(t_i = 1 | \mathbf{pa}(t_i))$. Then we get $P(\mathbf{x}^*, \mathbf{t}) = \max_{\mathbf{x}} \prod_{j=1}^n P(x_j) \prod_{i=1}^m P(t_i | \mathbf{pa}(t_i)) \leq \leq p^n \prod_{t_i=0} \alpha_0 \prod_{t_i=1} \alpha_1 = p^n \alpha_0^r \alpha_1^{m-r}$, where r is the number of $t_i = 0$ in \mathbf{t} . Since there are $\binom{m}{r}$ vectors \mathbf{t} having exactly r variables T_i assigned $t_i = 0$, we obtain

$$\sum_{\mathbf{t}} P(\mathbf{x}^*, \mathbf{t}) \leq p^n \sum_{r=0}^m \binom{m}{r} \alpha_0^r \alpha_1^{m-r} = p^n (\alpha_0 + \alpha_1)^m \quad (6)$$

and therefore, we get $Err_{MPE} = 1 - \sum_{\mathbf{t}} P(\mathbf{x}^*, \mathbf{t}) \geq L_{MPE}$, where L_{MPE} is the *lower bound* on the MPE diagnosis error. Thus, we just proved the following

Lemma 1 *Given Bayesian network $BN=(G,P)$ defining a joint distribution $P(\mathbf{x}, \mathbf{t})$ as specified by the equation 1, the MPE diagnosis error is given by $Err_{MPE} = 1 - \sum_{\mathbf{t}} P(\mathbf{x}^*, \mathbf{t}) \geq L_{MPE}$, and*

$$L_{MPE} = 1 - p^n (\alpha_0 + \alpha_1)^m = 1 - [p(\alpha_0 + \alpha_1)^{m/n}]^n, \quad (7)$$

where $p = \max_i \max\{P(X_i = 0), 1 - P(X_i = 0)\}$, and $\alpha_k = \max_i \max_{\mathbf{pa}(t_i)} P(t_i = k | \mathbf{pa}(t_i))$.

Our next question is *whether an error-free diagnosis is achievable asymptotically with $n \rightarrow \infty$* . We will assume a constant *test rate* $k = m/n$ (inspired by similar notion of code rate measuring the redundancy added to the input signal in order to decrease the decoding error in noisy channel coding). From the lower bound on the diagnosis error

(expression 7), we get the following *necessary condition* of (asymptotic) error-free diagnosis:

$$\lim_{n \rightarrow \infty} [p(\alpha_0 + \alpha_1)^k]^n \rightarrow z, \quad (8)$$

where the limit z must satisfy $z \geq 1$. This is equivalent to

$$p(\alpha_0 + \alpha_1)^k \geq 1 \Leftrightarrow (\alpha_0 + \alpha_1)^k \geq 1/p. \quad (9)$$

Note that $1 \leq 1/p \leq 2$ and $1 \leq \alpha_0 + \alpha_1 \leq 2$. The largest $1/p = 2$ (the "worst case" in terms of diagnostic error) corresponds to uniform priors $P(X_i)$. Then, in order to achieve an error-free diagnosis, it is necessary to have $k \geq 1$, i.e. $m \geq n$. In general, from 9 we get the following condition on the "amount of redundancy", or "code rate" $k = m/n$, necessary for the error-free diagnosis (note that the condition may not be sufficient for actually *achieving* zero error since we still have to investigate when the lower bound L_{MPE} is achievable):

$$L_{MPE} = 0 \Leftrightarrow k = m/n \geq \frac{\log(1/p)}{\log(\alpha_0 + \alpha_1)} \quad (10)$$

Applying the lemma 1 to the particular case of noisy-AND diagnosis with probing, we get the following

Corollary 2 *Given Bayesian network $BN=(G,P)$ defining a joint distribution $P(\mathbf{x},\mathbf{t})$ as specified by the equation 1, where all nodes X_i have same prior probability $p = P(X_i = 0) \leq 0.5$, and where all $P(t_j|\mathbf{pa}(t_j))$ are noisy-AND CPDs having same link probability q , leak probability l , and the number of parents r , the MPE diagnosis error is at least*

$$L_M = 1 - (1 - p)^n((1 - l)(1 - q^r) + 1)^m. \quad (11)$$

Note that in the absence of noise ($l=0$ and $q=0$) we get $L_M = 1 - (1 - p)^{n2^m}$, 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 $L_M \leq 0$ for m yields the necessary condition for zero lower bound, $m \geq -n \frac{\log(1-p)}{\log((1-l)(1-q^r) + 1)}$, plotted in Figure 2 as a function of p . Generally, solving $L_M \leq 0$ for m provides a way of specifying the minimum necessary number of probes that yield zero lower bound for a specified values of other parameters⁴. Also, the expression 11 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.

Diagnosis complexity and approximations

We focus first on the MPE diagnosis in the absence of noise (i.e., for deterministic test outcomes). The deterministic

⁴Clearly, finding a set of probes that may actually *achieve* the bound, if such set of probes exists, is a much harder task.

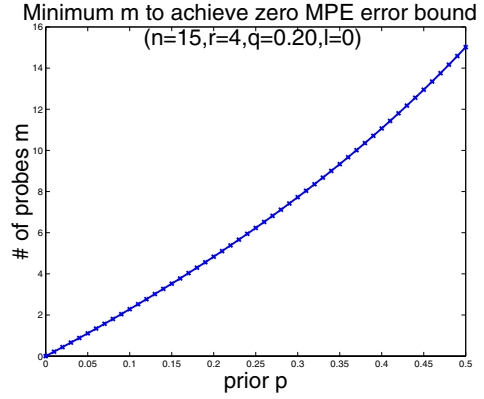


Figure 2: Minimum number of probes m to guarantee zero error bound, versus fault prior p : low prior yields lower than $n = 15$ number of probes.

CPDs reduce to a set of constraints imposed by the test outcomes on the values of X_1, \dots, X_n . For example, in the fault diagnosis domain, each probe outcome $T_i = t_i$ imposes a logical-AND constraint $t_i = x_{i_1} \wedge \dots \wedge x_{i_k}$ on the values of its parent nodes X_{i_1}, \dots, X_{i_k} . The MPE diagnosis becomes a constrained optimization problem of finding $\mathbf{x}^* = \arg \max_{x_1, \dots, x_n} \prod_{j=1}^n P(x_j)$ subject to those constraints. In a particular case of uniform priors $P(x_j)$, diagnosis is reduced to solving a constraint satisfaction problem. The problem can also be cast as a constraint satisfaction rather than optimization if there exist a unique solution satisfying the constraints (see (Brodie, Rish, & Ma 2001) for more details on how to construct such probe sets).

Although constrained optimization and constraint satisfaction problems (CSPs) are generally NP-hard, it is interesting to note that the probing domain yields a tractable set of constraints.

Proposition 3 *A set of constraints $t_j = x_{j_1} \wedge \dots \wedge x_{j_k}$, $j = 1, \dots, m$ over a set of variables X_1, \dots, X_n , where $x_i \in \{0, 1\}$, $t_j \in \{0, 1\}$ for $i = 1, \dots, n$ and $j = 1, \dots, m$, defines a propositional Horn theory, and can be, therefore, solved in $O(n)$ time by the unit-propagation algorithm*

Indeed, each successful probe yields a constraint $x_{i_1} \wedge \dots \wedge x_{i_k} = 1$ 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 $x_{i_1} \wedge \dots \wedge x_{i_k} = 0$, or equivalently, $\neg x_{i_1} \wedge \dots \wedge \neg x_{i_k} = 1$ imposed by failed probes which yields a *Horn theory* (i.e. a conjunction of clauses, or disjuncts, where each disjunct includes no more than one positive literal. 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 applying *unit propagation* to our Horn theory.

In the presence of noise, the MPE diagnosis task can be written as finding $\mathbf{x}^* = \arg \max_{x_1} \dots \max_{x_n} \prod_i P(x_i | \mathbf{pa}_i) = \arg \max_{x_1} F_1(x_1) \dots \max_{x_n} F_n(x_n, S_n), \quad (12)$

where each $F_i(x_i, \mathbf{S}_i) = \prod_{\mathbf{x}_k} \mathbf{P}(\mathbf{x}_k | \mathbf{pa}(\mathbf{x}_k))$ is the product of all probabilistic components involving X_i and a subset of

lower-index variables $S_i \subseteq \{X_1, \dots, X_{i-1}\}$, but *not* involving any X_j for $j > i$. The set of all such components is also called the *bucket* of X_i (Dechter 1996). An exact algorithm for finding MPE solution, called *elim-mpe* (Dechter 1996), uses *variable-elimination* (also called *bucket-elimination*) as a preprocessing: it computes the product of functions in the bucket of each variable X_i , from $i = n$ to $i = 1$ (i.e., from right to left in the equation 12), maximizes it over X_i , and *propagates* the resulting function $f(\cdot)$ to the bucket of its highest-order variable. Once variable-elimination is completed, the algorithm finds an optimal solution by a backtrack-free greedy procedure that, going from $i = 1$ to $i = n$ (i.e., in the opposite direction to elimination), assigns $X_i = \arg \max_{x_i} F_i(x_i, S_i = s_i)$ where $S_i = s_i$ is the current assignment to S_i . It is shown that *elim-mpe* is guaranteed to find an optimal solution and that the complexity of the variable-elimination step is $O(n \cdot \exp(w^*))$ where w^* , called the *induced width*, is the largest number of arguments among the functions (old and newly recorded) in all buckets (Dechter 1996). For the probing domain, 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.

Since the exact MPE diagnosis is intractable for large-scale networks, we focused on *local* approximation techniques. Particularly, we used a simple ($O(n)$ time) backtrack-free greedy algorithm, called here *greedy-mpe*, which performs no variable-elimination preprocessing, and the simplest and fastest member of the *mini-bucket* approximation family, algorithm *approx-mpe(1)* (Dechter & Rish 1997; 2002), that performs a very limited preprocessing similar to *relational arc-consistency* (Dechter & Rish 2002) in constraint networks.

The greedy algorithm *greedy-mpe* does no preprocessing (except for replacing observed variables with their values in all related function prior to algorithm's execution). It computes a suboptimal solution

$$\mathbf{x}' = (\arg \max_{x_1} F_1(x_1), \dots, \arg \max_{x_n} F_n(x_n, S_n = \mathbf{s}_n)), \quad (13)$$

where $S_i = \mathbf{s}_i$, as before, denotes the current assignment to the variables in S_j computed during the previous $i - 1$ maximization steps.

Generally, the mini-bucket algorithms *approx-mpe(i)* perform a limited level of variable-elimination, similar to enforcing directional i -consistency, prior to the greedy assignment. The preprocessing allows to find an upper bound U on $M = \max_{\mathbf{x}} P(\mathbf{x}, \mathbf{t})$, where \mathbf{t} is the evidence (clearly, $MPE = M/P(\mathbf{t})$), while the probability $L = P(\mathbf{x}', \mathbf{e})$ of their suboptimal solution provides an lower bound on M . Generally, L increases with the level of preprocessing controlled by i , thus allowing a flexible accuracy vs. efficiency trade-off. The algorithm returns the suboptimal solution \mathbf{x}' and the upper and lower bounds, U and L , on M ; ratio U/L is a measure of the approximation error.

We tested *greedy-mpe* and *approx-mpe(1)* on the networks constructed in a way that guarantees the unique diagnosis in the absence of noise. Particularly, besides m tests each having r randomly selected parents, we also generated n *direct* tests \hat{T}_i , $i = 1, \dots, n$, each having exactly one parent node

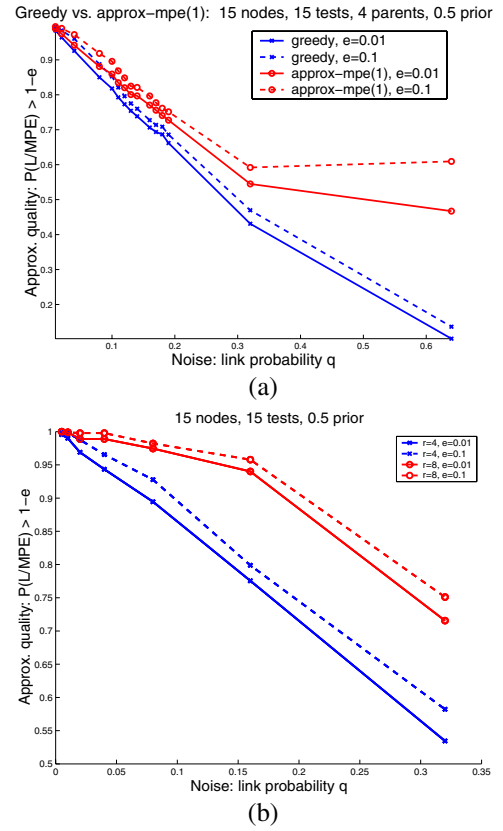


Figure 3: (a) Graceful degradation of the approximation quality of both greedy solution and an *approx-mpe(1)* solution with noise, where the approximation quality is measured as $P(L/M) > 1 - e$ for $e = 0.01$ and $e = 0.1$; the quality of *approx-mpe(1)* approximation degrades much slower than the quality of the greedy solution for larger noise (especially, for $q > 0.3$; (b) "zooming-in" on the quality of *approx-mpe(1)* for lower noise, $q \leq 0.32$; the accuracy gets higher for longer (more informative) probes (i.e., $r = 8$ vs. $r = 4$).

X_i . It is easy to see that, for such networks, both *greedy-mpe* and *approx-mpe(1)* find an exact diagnosis in the absence of noise: *approx-mpe(1)* reduces to unit-propagation, an equivalent of relational-arc-consistency, while *greedy-mpe*, applied along a *topological* order of variables in the network's directed acyclic graph (DAG)⁵, immediately finds the correct assignment which simply equals the outcomes of the direct tests.

Adding noise in a form of link probability q caused graceful degradation of the approximation quality, as shown in Figure 3. The figure summarizes the results for 50 randomly generated networks with $n = 15$ unobserved nodes (having uniform fault priors $p = P(x_i = 0) = 0.5$), $n = 15$ direct probes, one for each node, and $n = 15$ noisy-AND probes, each with $r = 4$ randomly selected parents among

⁵A *topological* (or *ancestral*) ordering of a DAG is an ordering where a child node never appears before its parent.




Figure 4: The accuracy of the solution \mathbf{x}' found by algorithm *greedy-mpe*, measured by L/M , where $L = P(\mathbf{x}', \mathbf{t})$ and $M = P(\mathbf{x}^*, \mathbf{t})$, versus M . The results obtained for

the unobserved nodes, zero leak $l = 0$ probability. The link probability (noise level) q varied from 0.01 to 0.64, taking 15 different values; the results are shown for all noise levels together. For each network, 100 instances of evidence (probe outcomes) were generated by Monte-Carlo simulation of \mathbf{x} and \mathbf{t} according to their conditional distributions. Thus, we get $50 \times 100 = 5000$ samples for each value of noise q .

Figure 3a plots both for *greedy-mpe* and for *approx-mpe(1)* the fraction of cases when the ratio L/MPE (where L is the solution probability) was within the interval $[1-e, 1]$ for small values of e . As expected, *approx-mpe(1)* yields a slower accuracy degradation with noise, especially for higher noise levels (i.e. even the simplest preprocessing pays off). Also, we observed that longer probes ($r = 8$ vs. $r = 4$) yield higher diagnosis accuracy, i.e. are more informative (Figure 3b).

Also, as demonstrated in Figure 4 for the same set of experiments as above (i.e., $n = 14$ and $r = 4$), the approximation accuracy of *greedy-mpe*, measured as L/M where $L = P(\mathbf{x}', \mathbf{t})$ and $M = P(\mathbf{x}^*, \mathbf{t})$, clearly increases with increasing value M , and therefore with the probability of the exact diagnosis, which also depends on the "diagnostic ability" of a probe set (for same probe set size, a better probe set yields a higher MPE diagnosis, and therefore, a better approximation quality). There is an interesting threshold phenomenon, observed both for *greedy-mpe* and for *approx-mpe(1)* solutions (the results for *approx-mpe(1)* are omitted due to space restrictions), and for various problem sizes n : the suboptimal solution \mathbf{x}' found by algorithm *greedy-mpe* suddenly becomes (almost always) an exact solution \mathbf{x}^* (i.e., $L/M = 1$, where $L = P(\mathbf{x}', \mathbf{t})$ and $M = P(\mathbf{x}^*, \mathbf{t})$) when $M > \theta$ where θ is some threshold value. For $n = 15$, the threshold is observed between $2e - 6$ and $3e - 6$. A theoretical analysis in the next section yields a quite accurate prediction of $\theta \approx 2.46e - 6$.

The effect of noise on approximation error

We will prove this claim formally for the simplest approximation algorithm *greedy-mpe*.

Let $BN = (G, P)$ be a Bayesian network, where $\mathbf{T} = \mathbf{t}$ is evidence, i.e. a value assignment \mathbf{t} to a subset of variables $\mathbf{T} \subset \mathbf{X}$. We will also make an important assumption that the all observed variables are replaced by their values in all CPD functions. Also, recall that $F_i(x_i, \mathbf{s}_i)$ is the product of functions in the bucket of X_i along the ordering o , given the assignment \mathbf{s}_i of some variables in the previous buckets. Then

Lemma 4 [greedy-mpe optimality.] *Given a Bayesian network $BN = (G, P)$, an evidence assignment $\mathbf{T} = \mathbf{t}$ applied to all relevant probability functions, and a topological ordering o of unobserved nodes in the graph G , the algorithm *greedy-mpe* applied along o is guaranteed to find an optimal MPE solution if $P(\mathbf{x}', \mathbf{t}) \geq F_i(x_i, \mathbf{s}'_i)$ for every $i = 1, \dots, n$ and for every $x_i \neq x'_i$, where $\mathbf{S}_i = \mathbf{s}'_i$ is a partial assignment already found by *greedy-mpe*.*

Proof. Clearly, the solution \mathbf{x}' found by *greedy-mpe* is optimal, i.e. $\mathbf{x}' = \mathbf{x}^* = \arg \max_{\mathbf{x}} P(\mathbf{x}, \mathbf{t})$ if $P(\mathbf{x}', \mathbf{t}) \geq P(\mathbf{x}, \mathbf{t})$ for every $\mathbf{x} \neq \mathbf{x}'$. Since $\mathbf{x} \neq \mathbf{x}'$ implies $x_i \neq x'_i$ for some i (let us choose the smallest of such i 's), by the condition of lemma we get $P(\mathbf{x}', \mathbf{t}) \geq F_i(x_i, \mathbf{s}'_i)$, and, therefore, $P(\mathbf{x}', \mathbf{t}) \geq \prod_{j=1}^n F_j(x_j, \mathbf{s}_j)$ since each $F_j(x_j, \mathbf{s}_j)$ is a product of probabilities, and therefore, $0 \leq F_j(x_j, \mathbf{s}_j) \leq 1$. But $\prod_{j=1}^n F_j(x_j, \mathbf{s}_j) = P(\mathbf{x}, \mathbf{t})$ by equation 12, which concludes the proof. ■

We now discuss some particular classes of Bayesian networks that satisfy the conditions of lemma 4.

Lemma 5 (nearly-deterministic CPDs, no observations.) *Given a Bayesian network $BN = (G, P)$ having no observed variables, and all conditional (and prior) probabilities being nearly-deterministic, i.e. satisfying the condition $\max_{x_i} P(x_i | pa(X_i)) > 1 - \delta$, where $0 \leq \delta \leq 0.5$, algorithm *greedy-mpe* applied along a topological ordering o of G is guaranteed to find an optimal MPE assignment if $(1 - \delta)^n \geq \delta$.*

Proof. Given a topological ordering and no evidence variables, the bucket of every node X_i contains a single function $P(x_i | pa(X_i))$. Thus, the greedy solution \mathbf{x}' yields $P(\mathbf{x}') = \prod_{i=1}^n \max_{x_i} P(x_i | pa(X_i)) = (1 - \delta)^n$, while any other \mathbf{x} has the probability $P(\mathbf{x}) = \prod_{i=1}^n P(x_i | pa(X_i)) < \delta$ since for the very first i such that $x_i \neq x'_i$ we get $P(x_i | pa(X_i)) < \delta$ and this value can only decrease when multiplied by other probabilities $0 \leq P(x_j | pa(X_j)) \leq 1$. ■

Let us consider a simulation that happened to select only the most-likely values for \hat{T}_i and T_i , i.e. $t'_i = \arg \max_{t_i} P(t_i | pa(T_i))$, which can be viewed as an error-free "transmission over a noisy channel". From 3 we get $\max_{t_i} P(t_i | pa(T_i)) \geq (1 - q)$; also, for any $t''_i \neq \arg \max_{t_i} P(t_i | pa(T_i))$, $P(t''_i | pa(T_i)) < q$. It is easy to show (similarly to lemma 4) that algorithm *greedy-mpe* will find an assignment that produced this most-likely evidence, thus yielding $P(\mathbf{x}', \hat{\mathbf{t}}, \mathbf{t}) = \prod_{i=1}^n P(x_i) \prod_{i=1}^n P(\hat{t}_i) \prod_{i=1}^n P(t_i) > \frac{1}{2^n} (1 - q)^{n+m}$. On the other hand, for any other \mathbf{x} there exists $T_j = t_j$ where

t_j is not the most-likely choice for T_j given \mathbf{x} , and thus $\hat{P}(t_i|pa(t_i)) < q$ as can be seen from the noisy-AND definition. Thus, the greedy solution \mathbf{x}' is guaranteed to be optimal once for any $\mathbf{x} \neq \mathbf{x}'$, $P(\mathbf{x}', \hat{\mathbf{t}}, \mathbf{t}) > P(\mathbf{x}, \hat{\mathbf{t}}, \mathbf{t})$, i.e. once $(1 - q)^{n+m} > q$ (the constant $\frac{1}{2^n}$ on both sides of the inequality was cancelled). Note that simulating an unlikely evidence yields a low joint probability $M = P(\mathbf{x}^*, \hat{\mathbf{t}}, \mathbf{t}) < q$ for the optimal diagnosis \mathbf{x}^* .

In our experiments, $n = m = 15$, thus resolving $(1 - q)^{30} = q$ gives a threshold value $q \approx 0.0806$, and therefore $M = P(\mathbf{x}', \hat{\mathbf{t}}, \mathbf{t}) = \frac{1}{2^{15}}(1 - q)^{30} > \frac{1}{2^{15}}q \approx 2.46e - 6$, which is surprisingly close to the empirical threshold observed in Figure 4 which separates suboptimal from the optimal behavior of algorithm *greedy-mpe*.

Discussion and conclusions

In this paper, we address both theoretically and empirically the problem of the most-likely diagnosis given the observations (MPE diagnosis), studying as an example the fault diagnosis in computer networks using probing technology. The key efficiency issues include minimizing both the number of tests and the computational complexity of diagnosis while maximizing its accuracy. Herein, we derive a lower bound on the diagnostic accuracy that provides necessary conditions for the number of probes needed to achieve an asymptotically error-free diagnosis as the network size increases, given a certain level of noise in probe outcomes and prior fault probabilities. Since the exact MPE diagnosis is generally intractable in large networks, we investigate next the accuracy/efficiency trade-offs for very simple and efficient approximation techniques, based on variable-elimination (the mini-bucket scheme), and provide both an empirical study on randomly generated networks and an initial theoretical explanation of the results. We show that even the most simple and inexpensive members of the mini-bucket algorithmic family (e.g., the greedy approximation) often provide an exact solution given sufficiently low levels of noise; as the noise increases, a "graceful degradation" of the accuracy is observed. Our results suggest the applicability of such approximations to nearly-deterministic diagnosis problems that are often encountered in practical applications.

Although there exists an extensive literature on fault diagnosis in computer networks (Kliger *et al.* 1997; Huard & Lazar 1996; I.Katzela & M.Schwartz 1995), we are not aware of any previous work that would both consider the problem of "active" diagnosis using probes, and provide theoretical and empirical analysis of diagnostic error, as well as a study of efficient approximation algorithms as presented in this paper. A closely related recent work in (Brodie, Rish, & Ma 2001) proposes efficient algorithms for the optimal probe set construction, although only in deterministic setting. Extending this approach to noisy environments is an interesting direction for future work. Further investigation, both theoretical and empirical, should focus on more accurate, but also more expensive local approximation techniques such as *approx-mpe(i)* and the related family of recently proposed *generalized belief propagation* techniques

(Yedidia, Freeman, & Weiss 2001), as well as on the variational approximation techniques that has been successfully used in two-layer noisy-OR networks for medical diagnosis (Jaakkola & Jordan 1999).

References

- Brodie, M.; Rish, I.; and Ma, S. 2001. Optimizing probe selection for fault localization. In *Distributed Systems Operation and Management*.
- Cooper, G. 1990. The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence* 42(2-3):393-405.
- Cover, T., and Thomas, J. 1991. *Elements of information theory*. New York: John Wiley & Sons.
- Dechter, R., and Rish, I. 1997. A scheme for approximating probabilistic inference. In *Proc. Thirteenth Conf. on Uncertainty in Artificial Intelligence (UAI97)*.
- Dechter, R., and Rish, I. 2002. Mini-buckets: A General Scheme for Approximating Inference. *To appear in J. of ACM*.
- Dechter, R. 1996. Bucket elimination: A unifying framework for probabilistic inference. In *Proc. Twelfth Conf. on Uncertainty in Artificial Intelligence*, 211-219.
- Frey, B., and MacKay, D. 1998. A revolution: Belief propagation in graphs with cycles. *Advances in Neural Information Processing Systems* 10.
- Heckerman, D., and Breese, J. 1995. Causal independence for probability assessment and inference using Bayesian networks. Technical Report MSR-TR-94-08, Microsoft Research.
- Henrion, M.; Pradhan, M.; Favero, B. D.; Huang, K.; Provan, G.; and O'Rourke, P. 1996. Why is diagnosis using belief networks insensitive to imprecision in probabilities? In *Proc. Twelfth Conf. on Uncertainty in Artificial Intelligence*.
- Huard, J., and Lazar, A. 1996. Fault isolation based on decision-theoretic troubleshooting. Technical Report 442-96-08, Center for Telecommunications Research, Columbia University, New York, NY.
- I.Katzela, and M.Schwartz. 1995. Fault identification schemes in communication networks. In *IEEE/ACM Transactions on Networking*.
- Jaakkola, T. S., and Jordan, M. I. 1999. Variational probabilistic inference and the qmr-dt network. *Journal of Artificial Intelligence Research* 10:291-322.
- Kliger, S.; Yemini, S.; Yemini, Y.; Ohsie, D.; and Stolfo, S. 1997. A coding approach to event correlation. In *Intelligent Network Management (IM)*.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.
- Shannon, C. 1948. A mathematical theory of communication. *Bell System Technical Journal* 27:379-423, 623-656.
- Weiss, Y. 2000. Correctness of local probability propagation in graphical models with loops. *Neural Computation* 12:1-41.
- Yedidia, J.; Freeman, W. T.; and Weiss, Y. 2001. Generalized belief propagation. In *NIPS 13*, 689-695. MIT Press.