

Information-theoretic approaches to cost-efficient diagnosis

Irina Rish

IBM T.J. Watson Research
Hawthorne, NY 10532, USA
Email: rish@us.ibm.com

Abstract—This paper provides a summary of our recent work on cost-efficient probabilistic diagnosis in Bayesian networks with applications to fault diagnosis in distributed computer systems. We focus on achieving good trade-offs between the diagnostic accuracy versus the cost of testing and computational complexity of diagnosis. We present (1) theoretical results characterizing these trade-offs, such as lower bound on the number of probes necessary to achieve asymptotically error-free diagnosis, (2) adaptive online approach to selecting most-informative tests, as well as (3) approximation techniques using "loopy" belief propagation for handling intractable inference problems involved in both diagnosis and most-informative test selection in large-scale problems. Empirical results on realistic systems demonstrating the effectiveness of our approaches can be found in [14], [16], [12], [15].

I. INTRODUCTION

The problem of diagnosing an unobserved "state of the world" from a set of available measurements and/or tests is quite common in practice. Examples include medical diagnosis, computer system troubleshooting, and decoding messages sent through a noisy channel. However, there is a trade-off between the quality of diagnosis and its cost, which involves both the cost of testing (e.g., the number of test if their costs are equal) and the computational cost of performing diagnosis.

One way to look at diagnostic problem is to view it as a combined source-channel coding, where the unknown state of the world described by a set of hidden variables $\mathbf{X} = (X_1, \dots, X_n)$ represents an input message, while the set of observed test outcomes $\mathbf{Y} = (T_1, \dots, T_m)$ corresponds to the output message that results from sending some "encoding" of \mathbf{X} , defined by the nature of tests, through a "noisy channel", determined by the nature of environment. The main difference from classical coding problem is that (1) source and channel coding are not always separable and (b) coding is constrained: we can only choose from a set of available tests rather than freely select arbitrary encoding functions.

Particularly, in this paper, we will focus on *disjunctive testing* motivated by fault diagnosis problem in distributed computer systems using *probes*. A probe is an end-to-end test transaction (e.g., ping, webpage access, database query, an e-commerce transaction, etc.) sent through the system for the purposes of performance monitoring. A probe can be viewed as a

disjunctive test over the components the probe depends (it returns OK if and only if all components on its path are OK). In case of noisy probe outcomes, we address diagnosis as a probabilistic inference in a Bayesian network that represents the dependencies between the unobserved states of system components and observed probe outcomes; conditional probabilities for probe outcomes given the corresponding components are defined by the *noisy-OR* model which generalizes disjunctive tests to the case of noisy environment.

We consider the test selection problem in both *non-adaptive* and *adaptive* settings. Nonadaptive setting assumes that a subset of tests must be selected offline prior to diagnosis, while in adaptive case the outcomes of the previous tests are known prior to selecting the next test. The nonadaptive probe selection problem is NP-hard [14], but the greedy approaches based on maximizing information gain (i.e. minimizing the conditional entropy about the unobserved nodes) work quite well in practice, particularly in cases when the number of faults is small and thus the state space of unobserved variables can be easily enumerated.

However, in a general multi-fault case the state space is exponential in the number of variables, and a compact representation such as Bayesian network must be used. Unfortunately, exact computation of conditional entropies in a general Bayesian network can be intractable. While much existing research has addressed the problem of efficient and accurate probabilistic inference, other probabilistic quantities, such as conditional entropy and information gain, have not received nearly as much attention. Most of the existing literature on value of information and most-informative test selection [8], [2], [7], [15] does not seem to focus on the computational complexity of most-informative test selection in a general Bayesian network setting, except for the most recent work by [10]. However, [10] focus on the nonadaptive ("non-myopic") problem and provide algorithms for efficient selection of a most-informative test subset given a bound on its size. Our problem is different as we consider adaptive ("myopic") test selection without a particular bound on the number of tests. We will describe our approximation algorithm for computing marginal conditional entropy [16]. The algorithm is based on loopy belief propagation, a successful approximate inference method. We illustrate the algo-

rithm at work in the setting of fault diagnosis for distributed computer networks. However, the method is general enough to be used in other applications of Bayesian networks that require the computation of information gain and conditional entropies of subsets of nodes.

Finally, we present some theoretical results for efficiency versus accuracy trade-off in diagnosis. Motivated by the Shannon's channel capacity result that provides conditions for asymptotically error-free decoding, one may ask whether similar conditions can be stated for certain classes of diagnostic problems (such as noisy disjunctive testing, or noisy-OR problems) as both the number of hidden and observed variables increase. While deriving achievable limit for such constrained classes of coding problems appears to be challenging, we are able to derive a lower bound on the diagnostic accuracy that provides *necessary* conditions for the number of probes needed to achieve asymptotically error-free diagnosis. Herein, we summarize our results from [12], providing lower bounds on the *bit-error rate* which assumes most-likely diagnosis for *each unobserved variable* ("bit-wise decoding"). See [13] for analysis of *block-error rate*, or Maximum-A-Posteriori (MAP) diagnosis.

II. BACKGROUND AND DEFINITIONS

Let $\mathbf{X} = \{X_1, X_2, \dots, X_N\}$ denote a set of N discrete random variables and \mathbf{x} a possible realization of \mathbf{X} . A *Bayesian network* is a directed acyclic graph (DAG) G with nodes corresponding to X_1, X_2, \dots, X_N and edges representing direct dependencies [11]. The dependencies are quantified by associating each node X_i with a local conditional probability distribution $P(x_i | \mathbf{pa}_i)$, where \mathbf{pa}_i is an assignment to the parents of X_i (nodes pointing to X_i in the Bayesian network). The set of nodes $\{x_i, \mathbf{pa}_i\}$ is called a *family*. The joint probability distribution function (PDF) over \mathbf{X} is given as product

$$P(\mathbf{x}) = \prod_{i=1}^N P(x_i | \mathbf{pa}_i). \quad (1)$$

We use $\mathbf{E} \subseteq \mathbf{X}$ to denote a possibly empty set of *evidence* nodes for which observation is available.

For ease of presentation, we will also use the terminology of *factor graphs* [6], which unifies directed and undirected graphical representations of joint PDFs. A factor graph is an undirected bipartite graph that contains factor nodes (usually shown as squares) and variable nodes (shown as circles). (See Fig. 1 for an example.) There is an edge between a variable node and a factor node if and only if the variable participates in the *potential function* of the corresponding factor. The joint distribution is assumed to be written in a factored form

$$P(\mathbf{x}) = \frac{1}{Z} \prod_a f_a(\mathbf{x}_a), \quad (2)$$




Fig. 1. Factor graph of the fault diagnostic Bayes net.

where Z is a normalization constant called the *partition function*, and the index a ranges over all factors $f_a(\mathbf{x}_a)$, defined on the corresponding subsets \mathbf{X}_a of \mathbf{X} .

Particularly, we will consider the following diagnostic Bayesian networks. Let $\mathbf{X} = \{X_1, X_2, \dots, X_N\}$ denote a set of unobserved random variables we wish to diagnose, and let $\mathbf{T} = \{T_1, T_2, \dots, T_M\}$ denote the available set of tests. We assume that test outcomes are independent given the states of components, and that component failures are marginally independent. These assumptions are captured by a bipartite Bayesian network that represents the above independence assumptions about the joint probability $P(\mathbf{x}, \mathbf{t})$:

$$P(\mathbf{x}, \mathbf{t}) = \prod_{i=1}^n P(x_i) \prod_{j=1}^m P(t_j | \mathbf{pa}(t_j)). \quad (3)$$

Fig. 1 shows a factor graph representation of our model.

Given the probe outcomes, we wish to find the most-likely assignment (called *maximum a posteriori probability*, or *MAP*) 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 $\mathbf{x}^* = \arg \max_{\mathbf{x}} P(\mathbf{x}, \mathbf{t})$. An alternative approach is to find the most likely value x_i^* of each node X_i separately, i.e. to find an assignment $\mathbf{x}' = (x'_1, \dots, x'_n)$ where $x'_i = \arg \max_{x_i} P(x_i | \mathbf{t})$, $i = 1, \dots, n$. We refer to the latter approach as *bit-wise diagnosis* (*bit-wise decoding*), while the MAP approach can be viewed as a *block-wise diagnosis* (*block-wise decoding*). Bit-wise diagnosis is more suited when using belief updating algorithms that compute posterior probability $P(X_i | \mathbf{T})$ for each X_i , rather than perform global optimization to find MAP, using either search or dynamic programming [4].

Unfortunately, both MAP inference and belief updating are known to be NP-hard [1], and the complexity of best-known inference techniques is exponential in the graph parameter known as *treewidth*, or *induced width* [3], which reflects the size of a largest clique in the graph (and thus the largest dependency) created by an inference algorithm. However, there exists a simple linear-time approximate inference algorithm known as *belief propagation* (*BP*) [11]. BP is provably correct on polytrees (i.e. Bayesian networks with no undirected cycles), and can be used as an approximation on

general networks. In belief propagation, probabilistic messages are iterated between the nodes. The process could diverge; convergence is guaranteed only for polytrees.

III. DIAGNOSIS WITH TEST SELECTION

In many diagnosis problems, the user has an opportunity to select tests in order to improve the accuracy of diagnosis. For example, in medical diagnosis, doctors face the *experiment design* problem of choosing which medical tests to perform next.

Our objective is to maximize diagnostic quality while minimizing the cost of testing. The diagnostic quality of a subset of tests \mathbf{T}^* can be measured by the amount of uncertainty about \mathbf{X} that remains after observing \mathbf{T}^* . From the information-theoretic perspective, a natural measurement of uncertainty is the conditional entropy $H(\mathbf{X} | \mathbf{T}^*)$. Clearly, $H(\mathbf{X} | \mathbf{T}) \leq H(\mathbf{X} | \mathbf{T}^*)$ for all $\mathbf{T}^* \subseteq \mathbf{T}$. Thus the problem is to find $\mathbf{T}^* \subseteq \mathbf{T}$ which minimizes both $H(\mathbf{X} | \mathbf{T}^*)$ and the cost of testing. When all tests have equal cost, this is equivalent to minimizing the number of tests.

This problem is known to be NP-hard [14]. A simple greedy approximation is to choose the next test to be $T^* = \arg \min_T H(\mathbf{X} | T, \mathbf{T}')$, where \mathbf{T}' is the currently selected test set. The expected number of tests produced by the greedy strategy is known to be within a $O(\log N)$ factor from optimal (see [16]). The same result holds for approximations (within a constant multiplicative factor) to the greedy approach. Furthermore, our empirical results show that the approach works well in practice [14].

We make a distinction between nonadaptive (off-line) test selection and adaptive (online) test selection. In online selection, previous test outcomes are available when selecting the next test. Off-line test selection attempts to plan a suite of tests before any observations have been made. We will focus on the online approach, sometimes called *active diagnosis*, which is typically much more efficient in practice than its off-line counterpart [14].

Adaptive Test Selection Problem: Given the observed outcome \mathbf{t}' of previously selected sequence of tests \mathbf{T}' , select the next test to be $\arg \min_T H(\mathbf{X} | T, \mathbf{t}')$.

In a Bayesian network, the joint entropy $H(\mathbf{X})$ can be decomposed into sum of entropies over the families and thus can be easily computed using the input potential functions. Conditional marginal entropies, on the other hand, do not generally have this property - see the lemmas below (and [16] for the proofs). Under certain independence conditions they decompose into functions over the families. But computing those functions will require inference

Lemma 1: Given a Bayesian network representing a joint PDF $P(\mathbf{X})$, the joint entropy $H(\mathbf{X})$ can be decomposed into the sum of entropies over the families: $H(\mathbf{X}) = \sum_{i=1}^N H(X_i | \mathbf{Pa}_i)$.

Lemma 2: Given a Bayesian network representing a joint PDF $P(\mathbf{X}, \mathbf{T})$, where $\forall i : \mathbf{pa}_{T_i} \subseteq \mathbf{X}$ (i.e. tests T_i and T_j are independent given a subset of \mathbf{X}), the observation \mathbf{t}' of previously selected test set, and a candidate test T , the conditional marginal entropy $H(\mathbf{X} | T, \mathbf{t}')$ can be written as

$$H(\mathbf{X} | T, \mathbf{t}') = - \sum_{t, \mathbf{x}_{\mathbf{pa}_T}} P(\mathbf{x}_{\mathbf{pa}_T}, t | \mathbf{t}') \log P(t | \mathbf{x}_{\mathbf{pa}_T}) + \sum_t P(t | \mathbf{t}') \log P(t | \mathbf{t}') + \text{const}, \quad (4)$$

where *const* is a constant expression.

Minimizing conditional entropy is a particular instance of *value-of-information* (VOI) analysis [7], where tests are selected to minimize the expected value of a certain *cost function* $c(\mathbf{x}, t, \mathbf{t}')$. The result of Lemma 2 can be generalized to this case if the cost function is decomposable over the families (see [16]).

Since observations of test outcome correlate the parent nodes, the exact computation of all the posterior probabilities in Eqn. (4) is intractable. We can certainly use an existing approximation method to compute $P(\mathbf{s}_{\mathbf{pa}_T}, t | \mathbf{t}')$ and $P(t | \mathbf{t}')$. But a more efficient approach is possible if we exploit the belief propagation infrastructure.

IV. BP FOR ENTROPY APPROXIMATION

Let us consider the problem of computing the conditional marginal entropy

$$H(\mathbf{X}_a | \mathbf{e}) = - \sum_{\mathbf{x}_a} P(\mathbf{x}_a | \mathbf{e}) \log P(\mathbf{x}_a | \mathbf{e}), \quad (5)$$

where $P(\mathbf{x}_a | \mathbf{e}) = \sum_{\mathbf{x} \setminus \mathbf{x}_a} P(\mathbf{x} | \mathbf{e})$, $\mathbf{x} \setminus \mathbf{x}_a$ representing variable nodes not in \mathbf{x}_a . The trick is to replace the marginal posterior $P(\mathbf{x}_a | \mathbf{e})$ with its factorized BP approximation, and make use of the BP message passing mechanism to perform the summation over \mathbf{x}_a . We call this process Belief Propagation for Entropy Approximation (BPEA).

Pick any node X_0 from \mathbf{X}_a and designate it as the root node. We modify the final message passed to X_0 as follows:

$$m'_{a \rightarrow 0}(x_0) := - \sum_{\mathbf{x}_a \setminus x_0} \tilde{b}_a(\mathbf{x}_a) \log \tilde{b}_a(\mathbf{x}_a). \quad (6)$$

Here, $\tilde{b}_a(\mathbf{x}_a)$ is the unnormalized belief of X_a (i.e., $\tilde{b}_a(\mathbf{x}_a) = \sigma b_a(\mathbf{x}_a)$, where $\sigma = \sum_{\mathbf{x}_a} \tilde{b}_a(\mathbf{x}_a)$).

Plugging in $\tilde{b}_a(\mathbf{x}_a)$ in place of $P(\mathbf{x}_a | \mathbf{e})$ in Eqn. 5, we see that it only remains to sum over the root node X_0 and normalize properly.

$$\tilde{h}(\mathbf{X}_a | \mathbf{e}) := \sum_{x_0} m'_{a \rightarrow 0}(x_0), \quad (7)$$

$$h(\mathbf{X}_a | \mathbf{e}) := \frac{\tilde{h}(\mathbf{X}_a | \mathbf{e})}{\sigma} + \log \sigma. \quad (8)$$

It follows immediately that BPEA is exact whenever BP is exact.

The normalization constant σ is already computed during normal BP iterations. The computation of $b_a(\cdot)$,

$m'_{a \rightarrow i}$, and $\tilde{h}(\cdot)$ can all be piggy-backed onto the same BP infrastructure, and therefore does not impact its overall complexity. Furthermore, due to the local and parallel message update procedure in BP, we can compute the marginal posterior entropies of multiple families in one single sweep. This is an important advantage for the adaptive testing setup.

It is also easy to show that the approach is extendible beyond the entropy computation, to an arbitrary cost function decomposable over families (see [16]). The cost function replaces the negative logarithm in Eqns. (5) and (6).

V. DIAGNOSIS WITH DISJUNCTIVE TESTS

Suppose we wish to monitor a system of networked computers. Let \mathbf{X} represent the binary state of N network elements. $X_i = 0$ indicates that the element is in normal operation mode, and $X_i = 1$ indicates that the element is faulty. We can take X_i to be any system component whose state can be measured using a suite of tests. If the system is large, it is often impossible to test each individual component directly. A common solution is to test a subset of components with a single *test probe*. If all the test components are okay, the test would return a 0. Otherwise the test would return 1, but it does not reveal which components are faulty.

We assume there are machines designated as *probe stations*, which are instrumented to send out *probes* to test the response of the network elements represented by \mathbf{X} . Let \mathbf{T} denote the available set of probes. A probe can be as simple as a *ping* request, which detects network availability. A more sophisticated probe might be an e-mail message or a webpage-access request.

In the absence of noise a probe is a disjunctive test: it fails if and only if there is at least one failed node on its path. More generally, it is a noisy-OR test [11]. The joint PDF of all tests and network nodes forms the well-known QMR-DT model [9]:

$$P(x_j) = (\alpha_j)^{x_j} (1 - \alpha_j)^{(1-x_j)}, \quad (9)$$

$$P(t_i = 0 \mid \mathbf{s}_{\mathbf{pa}_i}) = \rho_{i0} \prod_{j \in \mathbf{pa}_i} \rho_{ij}^{x_j}, \quad (10)$$

$$P(\mathbf{x}, \mathbf{t}) = \prod_i P(t_i \mid \mathbf{s}_{\mathbf{pa}_i}) \prod_j P(x_j). \quad (11)$$

Here, $\alpha_j := P(x_j = 1)$ is the prior fault probability, ρ_{ij} is the so-called inhibition probability, and $(1 - \rho_{i0})$ is the leak probability of an omitted faulty element. The inhibition probability is a measurement of the amount of noise in the network.

As discussed in Section III, we adopt the adaptive testing framework for fault diagnosis, sequentially selecting probes to minimize the conditional entropy. Our previous work [13] makes the single-fault assumption, which effectively reduces \mathbf{S} to one random variable with $N+1$ possible states. In general, however, multiple faults could exist in the system simultaneously, which requires the more complicated conditional entropy given in Eqn. (4).

Let $A(T, \mathbf{X}_{\mathbf{pa}_T} \mid \mathbf{t}')$ denote the first term in Eqn. (4). This is the cross entropy between the posterior probability of T and its parents, and the conditional probability of T given its parents. The second term in Eqn. (4) is simply the negative conditional entropy $-H(T \mid \mathbf{t}')$.

We deal with the two entropy terms separately. For $H(T \mid \mathbf{t}')$, we may use approximation methods such as BP or GBP to calculate the belief $b(t \mid \mathbf{t}')$, which can then be used to directly compute $H(T \mid \mathbf{t}')$. (Note that the summation over values of T is simple since T is binary-valued.) To calculate $A(T, \mathbf{X}_{\mathbf{pa}_T} \mid \mathbf{t}')$, we use the entropy approximation method BPEA, as described in Section IV. Because BP message updates are done locally, we can compute $A(T, \mathbf{X}_{\mathbf{pa}_T} \mid \mathbf{t}')$ for all unobserved T nodes during a single application of BP. Thus, picking the next probe requires only one run of the BPEA approximation algorithm.

For each candidate probe, we designate the probe node T itself as the root node. The unnormalized belief is $\tilde{b}_t(t, \mathbf{x}_{\mathbf{pa}_T}) := P(t \mid \mathbf{x}_{\mathbf{pa}_T}) \prod_{j \in \mathbf{pa}_T} n_{j \rightarrow t}(x_j)$. This is used to calculate the modified message $m'_{a \rightarrow t}(t)$ (cf. Eqn. (6)). However, since $A(T, \mathbf{S}_{\mathbf{pa}_T} \mid \mathbf{t}')$ is a cross entropy term, we do not take the log of \tilde{b} , but rather take the logarithm of the known probabilities $P(t \mid \mathbf{x}_{\mathbf{pa}_T})$. This simplifies the normalization step described in Eqn. (8) to $A(T, \mathbf{X}_{\mathbf{pa}_T} \mid \mathbf{t}') = \tilde{A}(T, \mathbf{X}_{\mathbf{pa}_T} \mid \mathbf{t}')/\sigma$, where $\sigma = \sum_{t, \mathbf{x}_{\mathbf{pa}_T}} \tilde{b}_t(t, \mathbf{x}_{\mathbf{pa}_T})$.

VI. DIAGNOSTIC ERROR BOUNDS

An interesting question one may ask is how many tests might be needed to guarantee accurate diagnostic results, assuming an ideal situation when the tests (probes) can be constructed rather than selected from a predefined set of available probes. We can ask for Shannon-limit type of a result, i.e. what is the minimal redundancy given by the ratio between the number of probes versus the number of unobserved nodes, that can guarantee that diagnostic error will approach zero as the number of components and probes goes to infinity? While deriving achievable limit is hard, we show below a lower bound on diagnostic error when using bit-wise most-likely diagnosis (measured as the *bit error rate (BER)*), for general bipartite Bayesian networks and particularly for noisy-OR bipartite networks.

The bit-error rate (BER) of diagnosis can be defined as $BER = \frac{\sum_{i=1}^n P(X_i \neq X'_i(\mathbf{T}))}{n}$, where $X'_i(\mathbf{T}) = \arg \max_x P(X_i = x \mid \mathbf{T})$ is the most-likely assignment to X_i given observed vector \mathbf{T} . Note that $X'_i(\mathbf{T})$ is a deterministic function if a deterministic tie-breaking rule is used for most-likely assignment (e.g., $X'_i = 0$ if $P(X_i = 0 \mid \mathbf{T}) = 0.5$).

Theorem 3: Given a bipartite Bayesian network that defines a joint distribution $P(\mathbf{x}, \mathbf{t})$ as specified by the equation 3, the bit error rate (BER) of bit-wise most-likely diagnosis is bounded from below as follows

$$BER \geq L_{BER} = 1 - p_{max}(\alpha_0 + \alpha_1)^c, \quad (12)$$

where $c = \max_i |ch_i|$, $|ch_i|$ being the number of X_i 's children, $p_{max} = \max_i \max_{j \in \{0,1\}} P(X_i = j)$ is the maximum prior probability over all nodes, and $\alpha_k = \max_{j \in \{1, \dots, m\}} \max_{\mathbf{pa}_j(t_j)} P(t_j = k | \mathbf{pa}_j(t_j))$ is the maximum conditional probability of the test outcome $k \in \{0, 1\}$, over all test variables and over all assignments to their corresponding parent nodes. See [12] for the proof.

Particularly, this can be applied to diagnosis in noisy-OR networks. To simplify our analysis, let us assume a particular structure that we will call a (k, c) -regular bipartite graph, where each node in the lower layer has exactly k parents in the upper layer, and each node in the upper layer has $c = km/n$ children in the lower layer (recall that there are n nodes in the upper layer and m nodes in the lower layer). Then the following results follow straightforwardly from the previous theorem:

Corollary 4: Given a Bayesian network having the (k, c) -regular bipartite graph structure, where n is the number of hidden nodes, m is the number of tests, and where all conditional probabilities $P(t_j | \mathbf{pa}(t_j))$ are noisy-OR functions having the link probability at least q and the leak probability at most q_{leak} , the bit error rate (BER) of bit-wise most-likely diagnosis is bounded from below as follows: $BER \geq$

$$L_{BER}^{NOR} = 1 - p_{max}(1 + q_{leak}(1 - q^k))^{km/n}. \quad (13)$$

Corollary 5: Given a bipartite Bayesian network that defines a joint distribution $P(\mathbf{x}, \mathbf{t})$ as specified by the equation 3, a necessary condition for achieving error-free bit-wise diagnosis is


$$L_{BER} \leq 0 \leftrightarrow c \geq \frac{\log 1/p_{max}}{\log(\alpha_0 + \alpha_1)}, \quad (14)$$

where c , α_0 and α_1 are defined as in Theorem 2. Particularly, for noisy-OR networks defined in Corollary 13, the necessary condition is

$$L_{BER}^{NOR} \leq 0 \leftrightarrow \frac{m}{n} \geq \frac{\log 1/p_{max}}{k \log(1 + q_{leak}(1 - q^k))}. \quad (15)$$

Assuming equal prior fault probabilities $p = P(X_i = 1)$, where $p < 0.5$ (typically, system's components are unlikely to be faulty), we get $\frac{m}{n} \geq \frac{\log(1/(1-p))}{k \log(1 + q_{leak}(1 - q^k))}$. In Figure 2a, we illustrate the growth of the lower bound on rate m/n with the increasing prior fault probability p , for different probe sizes k , and for a fixed noise parameters. As expected, higher probe to node ratio is necessary for higher fault probability p . Also, somewhat intuitively, longer probes (larger k) allow to reduce the required number of probes per node. However, this does not always happen in practice, which indicates that the bound is not tight, and indeed provides only necessary, but not sufficient, conditions for error-free diagnosis.

One direction for future work would be to provide achievable bounds, similar to Shannon limit, for the above constrained coding problem that only permits disjunctive codes, and a particular type of channel defined by noisy-OR model. Namely, one would like



(a) (b)

Fig. 2. A lower bound on rate m/n necessary for achieving zero-error diagnosis, plotted the versus fault prior p , for different probe length k .

to know if asymptotically error-free diagnosis is actually achievable at finite rate m/n , and under what conditions on prior p , noise parameters, and probe set construction. While there is a large amount of related work in the area of group testing (e.g., see [5]), this particular setting does not seem to be studied before. Moreover, taking into account constraints on probe construction (e.g., due to the network topology restrictions) makes the analysis much more complicated.

REFERENCES

- [1] G.F. Cooper. The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42(2-3):393-405, 1990.
- [2] J. de Kleer and B.C. Williams. Diagnosing Multiple Faults. *Artificial Intelligence*, 32(1), 1987.
- [3] R. Dechter. Bucket elimination: A unifying framework for probabilistic inference. In *Proc. Twelfth Conf. on Uncertainty in Artificial Intelligence*, pages 211-219, 1996.
- [4] R. Dechter and I. Rish. Mini-buckets: A General Scheme for Approximating Inference. *J. of ACM*, 50(2):107-153, 2003.
- [5] D.-Z. Du and F.K. Hwang. *Combinatorial Group Testing and Its Applications (2nd edition)*. World Scientific, 2000.
- [6] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger. Factor graphs and the sum-product algorithm. *IEEE Trans. Info. Theory*, pages 498-519, 2001.
- [7] D. E. Heckerman, E. J. Horvitz, and B. Middleton. An approximate nonmyopic computation for value of information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15:292-298, 1993.
- [8] R. Howard. Information value theory. *IEEE Trans Syst Sci Cybern.* 2(1):22-26, 1966.
- [9] T. Jaakkola and M. Jordan. Variational probabilistic inference and the qmr-dt database. *Journal of Artificial Intelligence Research*, pages 291-322, 1999.
- [10] A. Krause and C. Guestrin. Near-optimal Nonmyopic Value of Information in Graphical Models. In *Proc. of Uncertainty in Artificial Intelligence (UAI)*, 2005.
- [11] J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, San Mateo, California, 1988.
- [12] I. Rish. Distributed Systems Diagnosis Using Belief Propagation.
- [13] I. Rish, M. Brodie, and S. Ma. Accuracy vs. Efficiency Trade-offs in Probabilistic Diagnosis. In *Proceedings of AAAI-2002, Edmonton, Alberta, Canada*, pages 560-566, 2002.
- [14] I. Rish, M. Brodie, S. Ma, N. Odintsova, A. Beygelzimer, G. Grabarnik, and K. Hernandez. Adaptive diagnosis in distributed systems. *IEEE Transactions on Neural Networks (special issue on Adaptive Learning Systems in Communication Networks)*, 16(5):1088-1109, 2005.
- [15] I. Rish, M. Brodie, N. Odintsova, S. Ma, and G. Grabarnik. Real-time Problem Determination in Distributed Systems using Active Probing. In *Proceedings of 2004 IEEE/IFIP Network Operations and Management Symposium (NOMS 2004)*, Seoul, Korea, pages 133-146, 2004.
- [16] A. Zheng, I. Rish, and A. Beygelzimer. Efficient Test Selection in Active Diagnosis via Entropy Approximation. In *Proceedings of UAI-05*, 2005.