APPLICATION OF BAYESIAN REASONING TECHNIQUES TO FAULT LOCALIZATION IN FCS NETWORKS

Małgorzata Steinder
Adarshpal S. Sethi

Computer and Information Sciences
University of Delaware, Newark, DE

ABSTRACT

FCS networks are aimed at providing a highly automated, secure, and survivable paradigm of battlefield operations. This goal cannot be achieved without an ability to rapidly isolate and correct network faults. A fault management system for FCS networks, which are ad-hoc and mobile, should be characterized by high accuracy and efficiency as well as the ability to deal with uncertainty, unreliability, and dynamism – inherent properties of such networks. We propose an application of Bayesian reasoning techniques to fault localization in FCS networks and present a fault localization algorithm capable of identifying multiple simultaneous faults in an efficient and event-driven manner. The algorithm provides an accurate fault hypothesis in the presence of uncertain information about the system structure and is resilient to noise in observed symptoms. We evaluate the algorithm through simulation in which its accuracy and performance are assessed in identifying root causes of end-to-end connectivity problems. [1]

1 INTRODUCTION

Automatic fault localization [5, 7, 14], is known to be a difficult problem in any circumstances. In a wireless environment, the well known issues of environment-related device failures, inherently unreliable communication medium, power supply problems, and restricted bandwidth make the fault management function even more challenging. They result in new failures to consider, higher fault frequencies, higher symptom loss rates, increased number of transient faults, less computing resources available, and severely restricted amount of management information that may be exchanged between network nodes. In ad hoc and mobile networks, the problem is further complicated by the dynamically changing topology. Yet, without the ability to rapidly isolate and correct network faults, the highly automated, secure, and survivable paradigm of battlefield operations expected of FCS networks may not be achieved.

While investigating fault localization techniques for FCS networks, we set the following objectives:

- *Dealing with uncertainty*, which is indispensable when mutual dependencies between system components are transient or difficult to obtain or when so called *soft* (byzantine) problems have to be diagnosed.
- *Event-driven fault localization* as opposed to window-based analysis [1, 7, 14]. In a mobile environment, postponing symptom analysis until the end of time-window may cause the symptoms to be analyzed using an outdated information about the system structure.
- *Resilience to observation noise*, which is due to lost or spurious symptoms and can dramatically reduce the accuracy of the fault localization process [10, 12].
- *High accuracy and low computational complexity* – Fault localization in nondeterministic systems has been shown an NP-hard problem [7]. Since an optimal solution is likely to be prohibitively time-consuming, it is important to design efficient yet accurate approximate techniques.

This paper investigates an application of probabilistic reasoning to fault localization in FCS networks. We address the objectives of our research by adapting known algorithms of probabilistic reasoning in belief networks [3, 8] to create an efficient and accurate probabilistic fault localization technique. In Section 2, we present a technique of modeling FCS networks using belief networks [8]. In Section 3, we propose a heuristic that applies the belief-updating algorithm for polytrees [8] to perform event-driven diagnosis in non-polytree belief networks. In Section 4, we extend the system model to incorporate lost and spurious symptoms. In Section 5, we present the results of the simulation study.

2 MODELING FAULT PROPAGATION WITH BELIEF NETWORKS

Similar to many fault localization techniques [1, 4, 14], the technique presented in this paper relies on a graphical fault propagation model (FPM), which represents causal relationships among events [1, 4, 14]. This causality graph may be interpreted as a belief network with binary-valued nodes.

A *belief network* [3, 8] (BN) is a directed acyclic graph

[1] Prepared through collaborative participation in the Communications and Networks Consortium sponsored by the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-2-0011. The U. S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

(DAG), in which each node V_i represents a random variable over a multivalued domain. The set of all nodes is denoted by V. The set of directed edges E denotes an existence of causal relationships between the variables and the strengths of these influences are specified by conditional probabilities. Formally, a belief network is a pair (G, P), where G is a DAG, $P{=}\{P_i\}$, and P_i is the conditional probability matrix associated with a random variable V_i. We focus on a class of BNs representing a simplified model of conditional probabilities called *noisy-OR gate* [8], which contains binary-valued random variables and associates an inhibitory factor with every cause of an effect. The effect is absent only if all inhibitors corresponding to the present causes are activated. Thus, instead of conditional probability matrices associated with BN nodes, the noisy-OR BN assigns conditional probability values to the BN edges. The model assumes that all inhibitory mechanisms are independent [8]. This assumption of independence is ubiquitous in probabilistic fault localization [7, 14]. It is known that quering a BN is, in general, NP-hard [2]. A belief updating algorithm, polynomial with respect to $|V|$, is available for *polytrees* [8].

In a BN used as an FPM, each 0,1-valued variable represents a failure of a particular system entity. An assignment of 1 or 0 indicates that the entity experiences or does not experience the represented failure, respectively. Several distinct variables may be associated with the same entity to represent its various failure conditions [11]. The fact that a failure of one entity may cause a failure of another entity is represented by a causal edge between the corresponding BN nodes, which is weighted with the probability of the causal implication.

A symptom is defined as an observation that an entity experiences a particular failure (*negative* symptom), or that it does not experience this failure (*positive* symptom). We will denote by $\mathcal{S}$ the set of all possible symptoms. If V_i is a BN node corresponding to a failure of a system entity, then the negative and positive symptoms are interpreted as an instantiation of V_i with value 1 and 0, respectively. The sets of all observed negative and positive symptoms will be denoted by $\mathcal{S}_N$ and $\mathcal{S}_P$, respectively. The set of all observed symptoms, $\mathcal{S}_O{=}\mathcal{S}_N{\cup}\mathcal{S}_P{\subseteq}\mathcal{S}$, becomes the evidence. In general, $\mathcal{S}_N{\cup}\mathcal{S}_P{\neq}\mathcal{S}$, as some symptoms may be unobservable. The ratio $|\mathcal{S}_O|/|\mathcal{S}|$ will be called an *observability ratio* (OR). A BN in which $\mathcal{S}_O$ represents the set of all observable symptoms will be denoted by $\mathcal{BN}(\mathcal{S}_O)$.

A fault is a failure of a system entity that may not be further explained with a given FPM. It is represented by the assignment of 1 to the corresponding BN node. The set of all possible faults is denoted by $\mathcal{F}$. The problem of finding the set of faults, $\mathcal{F}_c \subseteq \mathcal{F}$ that best explains the set of observed symptoms $\mathcal{S}_O$ may be solved by computing the most probable explanation (MPE) query in $\mathcal{BN}(\mathcal{S}_O)$.

3 FAULT LOCALIZATION TECHNIQUE

Recall from Section 2 that in singly-connected BNs representing the noisy-or model of conditional probability distribution, Bayesian belief updating may be computed in polynomial time [8]. BNs used as FPMs typically are not polytrees because they contain undirected loops [11]. We apply Pearl's belief updating algorithm [8] as an approximation scheme to perform fault localization in an FPM with loops.

Pearl's belief updating [8] utilizes a message passing schema in which BN nodes exchange λ and π messages. Message $\lambda_X(v_j)$ that node X sends to its parent V_j for every valid V_j's value v_j, denotes a posterior probability of the entire body of evidence in the sub-graph obtained by removing link $V_j{\rightarrow}X$ that contains X, given that $V_j{=}v_j$. Message $\pi_{U_i}(x)$ that node X sends to its child U_i for every valid value of X, x, denotes a probability that $X{=}x$ given the entire body of evidence in the subgraph containing X created by removing edge $X{\rightarrow}U_i$. Based on the messages received from its parents and children, node X computes $bel(x)$, a probability that $X = x$ given the entire body of observed evidence, and messages λ and π for the node's parents and children, respectively. The calculation of messages λ and π, and belief metric bel is explained in [8].

In the initialization phase, for all observed nodes X, $\lambda(x)$ is set to 1 if x is the observed value of X. For other values of x, $\lambda(x)$ is set to 0. For all unobserved nodes $\lambda(x)$ is set to 1 for all values of x. Parentless nodes have their $\pi(x)$ set to the prior probabilities. The belief propagation algorithm in polytrees starts from the evidence node and propagates the changed belief along BN edges by computing $bel(x)$, $\lambda_X(v_i)$s and $\pi_X(u_i)$s in every visited node.

For the purpose of event-driven fault localization this paper adapts the iterative belief updating as follows. The adapted algorithm, **APPROX-FL**, starts with a BN all of whose evidence nodes corresponding to observable symptoms are assigned to 0, and all other nodes are unassigned, i.e., their $\lambda(0) = \lambda(1) = 1$. Then, the algorithm proceeds in an event-driven manner, after every symptom observation applying one iteration of belief updating traversing the graph according to some order. For every symptom we define a different ordering that is equivalent to the breadth-first order started from the node representing the observed symptom. The initialization phase considers all observable symptoms positive and calculates fault probability distribution in the presence of no negative observations. When a negative symptom is observed the propagation of negative evidence reverses the results of the corresponding positive symptom analysis performed in the initialization phase.

The iterative belief propagation described above allows us to obtain the marginal posterior distribution resulting from the observation of the evidence (symptoms). We use this distribution to estimate the set of faults that are the most probable causes of the observed symptoms. For this purpose, we introduce the following heuristic: (1) choose a fault node with the highest posterior probability, (2) place the corresponding fault in the MPE hypothesis, (3) mark the node as observed with value 1, and (4) perform one iteration of the belief propagation starting from the chosen fault node. Steps (1)-(4) are repeated as long as (1) the posterior distribution contains fault nodes whose probability is greater than 0.5, and (2) unexplained negative symptoms remain. An inherent property of the adapted algorithm is the capability to isolate multiple simultaneous faults even if their symptoms overlap.

Algorithm APPROX-FL

```
Inference iteration starting from node $Y_i$:
    let $o$ be the breadth-first order starting from $Y_i$
    for all nodes $X$ along ordering $o$ do
        if $X$ is not an unobserved or positive symptom node then
           compute $\lambda_X(v_j)$ for all $X$'s parents, $V_j$,
                       and for all $v_j \in \{0,1\}$
           compute $\pi_{U_i}(x)$ for all $X$'s children, $U_i$,
                       and for all $x \in \{0,1\}$
Initialization phase:
    for every symptom $S_i \in \mathcal{S}_O$ do
        mark $S_i$ as observed to have value of 0
        run inference iteration starting from $S_i$
Symptom analysis phase:
    for every symptom $S_i \in \mathcal{S}_N$ do
        mark $S_i$ as observed to have value of 1
        run inference iteration starting from $S_i$
    compute $bel(v_i)$ for every node $V_i$, $v_i \in \{0,1\}$
Fault selection phase:
    while $\exists$ fault node $V_j$ for which $bel(1)>0.5$ and $\mathcal{S}_N \neq \emptyset$ do
        take $V_j$ with the greatest $bel(1)$ and
           place it in the set of detected faults $\mathcal{F}_D$
        mark $V_j$ as observed to have value of 1
        remove all symptoms explained by $V_j$ from $\mathcal{S}_N$
        run inference iteration starting from $V_j$
        compute $bel(v_i)$ for every node $V_i$, $v_i \in \{0,1\}$
```

Unlike other approaches to fault localization [7, 14], which delay symptom analysis until all symptoms are collected, algorithm APPROX-FL does not require all symptoms to be observed before their analysis may be started. On the contrary, it analyzes a symptom independently of other symptom observations. The knowledge resulting from analyzing a symptom is stored for the next iterations in the BN nodes in the form of λ and π messages, allowing algorithm APPROX-FL to utilize time more efficiently. The complexity of the entire algorithm is $\mathcal{O}(|\mathcal{S}_O||E|)$ [13].

4 DEALING WITH NOISY OBSERVATIONS

In FCS networks, an observation of network state is likely to be disturbed by the presence of lost or spurious symptoms (referred to as observation noise). Alarms may be lost as a result of using an unreliable communication mechanism to transfer alarms from their origin to the management node. Too liberal threshold values may also prevent an existing problem from being reported, thereby causing alarm loss. When the fault localization algorithm relies on positive information to create the most likely fault hypothesis, alarm loss, if ignored by the algorithm, could lead to an incorrect solution. Spurious alarms are caused by intermittent network faults or by overly restrictive threshold values.

We address the problem of lost and spurious alarms by augmenting the BN model described in Section 2 using the technique we introduced in [10]. Let $V_S{=}\{V_{S_1},\ldots,V_{S_m}\}{\subset}V$, where $m{=}|\mathcal{S}|$, be the set of all BN nodes which correspond to symptoms $\{S_1,\ldots,S_m\}{=}\mathcal{S}$. We introduce a set of nodes $\{V'_{S_1},\ldots,V'_{S_m}\}$, which represent unobservable failures. Then for every node $V_{S_i}{\in}V_S$ and for every $V_j{\in}V{-}V_S$ such that $(V_j,V_{S_i}){\in}E$ we (1) remove (V_j,V_{S_i}) from E and (2) add (V_j,V'_{S_i}) to E. Then, we add directed edges $V'_{S_i}{\rightarrow}V_{S_i}$, $i{=}1\ldots m$. With every directed edge $V'_{S_i}{\rightarrow}V_{S_i}$ we associate the probability of causal relationship between V'_{S_i} and V_{S_i} equal to $1{-}p_{\text{loss}}(S_i)$, where $p_{\text{loss}}(S_i)$ is the probability that alarm S_i is lost. The values of $p_{\text{loss}}(S_i)$ may be obtained, for example, by analyzing packet loss rate in the network used to transport symptom S_i from its origin to the management station.

To model spurious symptoms, we introduce nodes $V^*_{S_i}$, $i{=}1\ldots m$. Then, we add directed edges $V^*_{S_i}{\rightarrow}V_{S_i}$. With every $V^*_{S_i}$ we associate prior belief $p_{\text{spurious}}(S_i)$ that represents the cumulative probability of events (other than persistent faults) that trigger alarm S_i. The value of $p_{\text{spurious}}(S_i)$ may be learned by analyzing historical alarm log files. Every directed edge $V^*_{S_i}{\rightarrow}V_{S_i}$ is labeled with $1{-}p_{\text{loss}}(S_i)$.

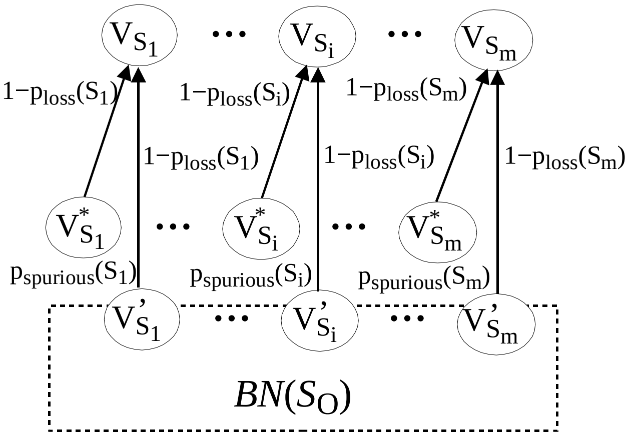

Figure 1: A belief network modeling lost and spurious symptoms ($\mathcal{BN}(\mathcal{S}_O, p_{\text{loss}}, p_{\text{spurious}})$)

The resultant BN, $\mathcal{BN}(\mathrm{S}_O, p_{\text{loss}}, p_{\text{spurious}})$, is presented in Fig. 1. Observe that, $\mathcal{BN}(\mathcal{S}_O, 0, 0)$ is equivalent to the original BN, $\mathcal{BN}(\mathcal{S}_O)$. When the existence of either lost or spurious (but not both) alarms may be neglected, one should use $\mathcal{BN}(\mathcal{S}_O, 0, p_{\rm spurious})$ or $\mathcal{BN}(\mathcal{S}_O, p_{\rm loss}, 0)$, respectively.

To perform fault localization that includes lost and spurious symptoms in the analysis using $\mathcal{BN}(\mathcal{S}_O,\, p_{\rm loss}, p_{\rm spurious})$ as an FPM, the algorithm APPROX-FL may be applied.

5 SIMULATION STUDY

Network connectivity is frequently achieved through a sequence of intermediate nodes. A failure of such a node may cause availability or performance problems on one or more end-to-end paths established using the malfunctioning node. Identification of a failing node, which is a critical to restoring network operation, is particularly difficult when a failure is byzantine or when the number of possible suspects is large. In this fault localization problem, an FPM is a bipartite causality graph with path failures at the heads and link failures at the tails of the edges, respectively.

To evaluate algorithm APPROX-FL in the application to the diagnosis of end-to-end connectivity problems, we randomly generate tree-shaped network topologies. (The choice of tree-shaped network structure makes scenario generation faster and easier, while not affecting the validity of the results [9].) We calculate two parameters that describe the algorithm's accuracy: (1) detection rate, DR, defined as the ratio of existing faults that are correctly identified, and (2) false positive rate, FPR – the ratio of proposed faults that did not exist in the system in the considered scenario.

In the first set of experiments, we compare the accuracy achievable with algorithm APPROX-FL and an optimal algorithm [3]. The results of this study, which are presented in Figs. 2(a)-2(b), show that algorithm APPROX-FL offers a close-to-optimal accuracy while being applicable to networks of much bigger size that the optimal algorithm, which is not suitable networks bigger that 10 intermediate nodes because of the excessive fault localization time.

To assess the impact of symptom loss on the accuracy of algorithm APPROX-FL, we compare DR and FPR achieved while using BNs $\mathcal{BN}(\mathcal{S}_O, LR, 0)$ (taking symptom loss into account) and $\mathcal{BN}(\mathcal{S}_O, 0, 0)$ (disregarding the possibility of symptom loss), varying OR between 0.2 and 0.5. We set SSR=0 in the simulation model. Loss rate is either 0.05 or 0.1. Figs. 3(a) and 3(b) show that, by including loss rate in the analysis, the detection (false positive) rate may be increased (decreased) by up to 10%.

The impact of including spurious symptoms in the fault localization process is evaluated by applying algorithm APPROX-FL to FPMs $\mathcal{BN}(\mathcal{S}_O, 0, SSR)$ and $\mathcal{BN}(\mathcal{S}_O, 0, 0)$ using LR=0, and OR=0.5. We vary SSR between 0.01 and 0.1. As shown in Fig. 4(a), the inclusion of spurious symptoms in the fault localization process in small networks decreases DR. This is explained by the fact that in small networks (in particular, with small observability ratios), only a few symptoms are available to the fault localization process. Consequently, the algorithm refuses to identify a fault thereby achieving very low DR. One can conclude that small networks need to be better instrumented (i.e., have higher OR) to allow fault localization to benefit from the analysis of spurious symptoms. In larger networks, the inclusion of spurious symptoms does not cause a decrease in the DR; in fact, as shown in Fig. 4(a), it allows the DR to be improved. Fig. 4(b) presents the impact of including spurious symptoms on the FPR. It shows that regardless of the network size, by taking spurious symptoms into account, the FPR of the fault localization process may be significantly decreased.

6 CONCLUSION

This paper investigates an application of Bayesian reasoning using belief networks [8] to non-deterministic fault diagnosis in FCS networks. We propose a BN as a representation of causal relationships among system events and show that the fault localization problem may be solved by calculating the MPE query in BNs. We show that the approximate technique proposed in this paper has close-to-optimal accuracy and is resilient to observation noise.

Future work will involve assessing the algorithm's efficiency using wireless network simulator, investigating other (more efficient) fault localization techniques, improving the fault localization efficiency through distributed diagnosis, designing efficient approximate methods of building an FPM, and integrating fault localization with self-healing and network restoration mechanisms [6] to provide a comprehensive fault management solution for FCS networks. [2]

REFERENCES

[1] C. S. Chao, D. L. Yang, and A. C. Liu. An automated fault diagnosis system using hierarchical reasoning and alarm correlation. *Journal of Network and Systems Management*, 9(2):183–202, 2001.

[2] G. F. Cooper. Probabilistic inference using belief networks is NP-Hard. Technical Report KSL-87-27, Stanford University, 1988.

[3] R. Dechter. Bucket Elimination: A unifying framework for probabilistic inference. In *Uncertainty in AI*, 1996, pp. 211–219.

[4] M. Hasan, B. Sugla, and R. Viswanathan. A conceptual framework for network management event correlation and filtering systems. In *Integrated Network Management VI*, May 1999, pp. 233–246.

[5] G. Jakobson and M. D. Weissman. Alarm correlation. *IEEE Network*, 7(6):52–59, Nov. 1993.

[6] L. Kant, A. S. Sethi, and M. Steinder. Fault localization and self-healing mechanisms for FCS networks. In *Proc. of Army Science Conference*, Orlando, FL, Dec. 2002.

[7] I. Katzela and M. Schwartz. Schemes for fault identification in communication networks. *IEEE Transactions on Networking*, 3(6):733–764, 1995.

[2] The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied of the Army Research Lab or the U.S. Government.

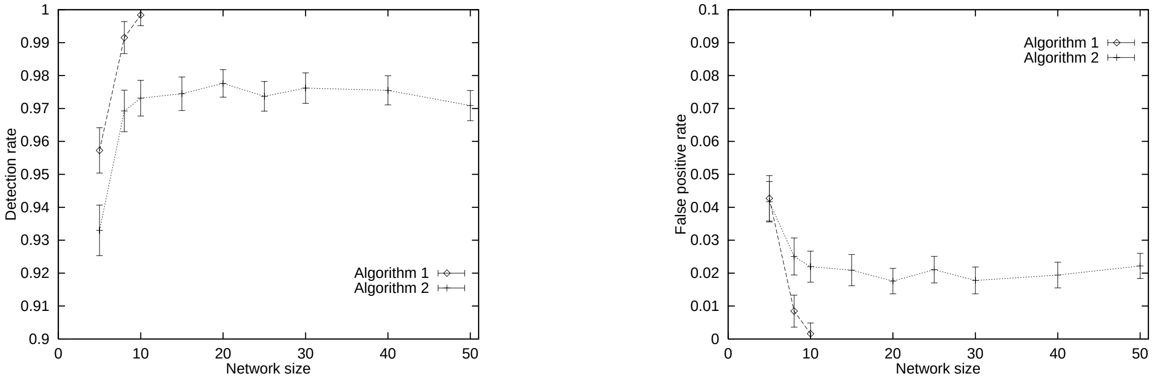

(a) Detection rate

(b) False positive rate

Figure 2: Comparison of accuracies achievable with the optimal algorithm (1) and algorithm APPROX-FL (2).

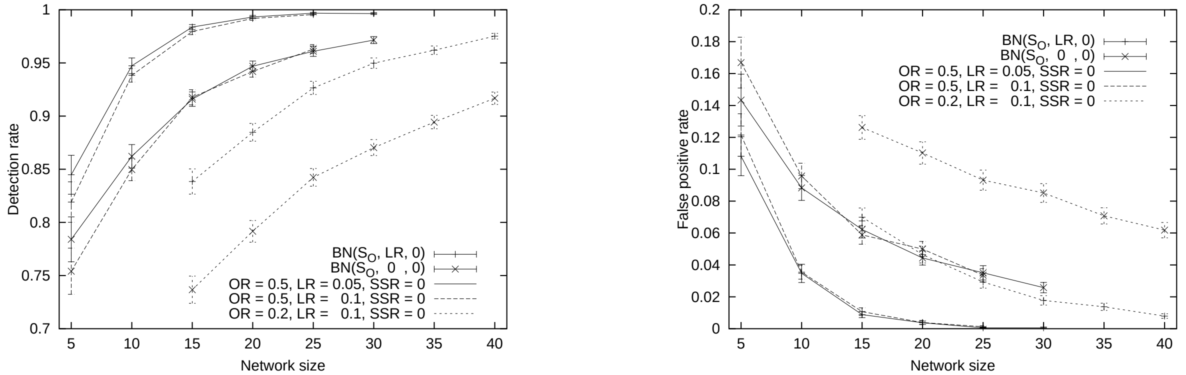

(a) Detection rate

(b) False positive rate

Figure 3: Accuracy of algorithm APPROX-FL using fault propagation models $\mathcal{BN}(\mathcal{S}_O, LR, 0)$ and $\mathcal{BN}(\mathcal{S}_O, 0, 0)$.

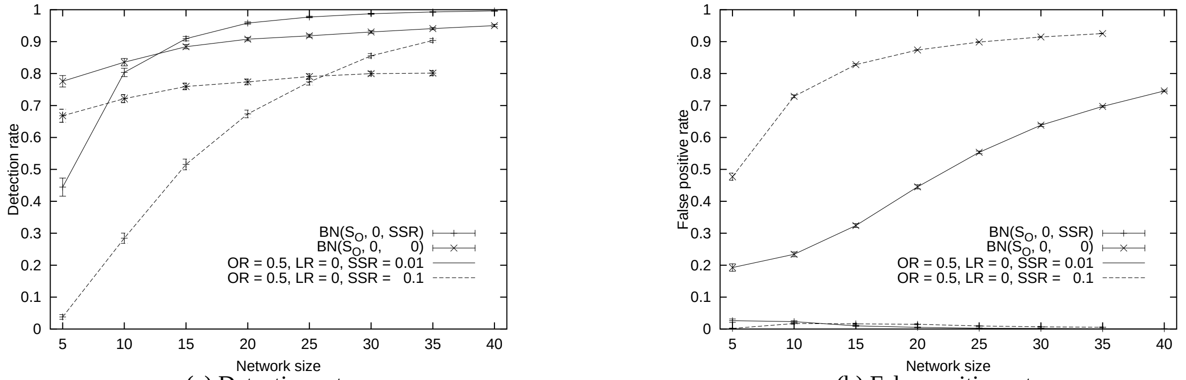

(a) Detection rate

(b) False positive rate

Figure 4: Accuracy of algorithm APPROX-FL using fault propagation models $\mathcal{BN}(\mathcal{S}_O, 0, SSR)$ and $\mathcal{BN}(\mathcal{S}_O, 0, 0)$.

[8] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, 1988.

[9] M. Steinder and A. S. Sethi. Distributed fault localization in hierarchically routed networks. In *Int'l Workshop on Distributed Systems: Operations and Management*, Montréal, Canada, Oct. 2002. Springer, pp. 195–207.

[10] M. Steinder and A. S. Sethi. Increasing robustness of fault localization through analysis of lost, spurious, and positive symptoms. In *Proc. of IEEE INFOCOM*, New York, NY, 2002.

[11] M. Steinder and A. S. Sethi. Non-deterministic fault localization in communication systems using belief networks. Technical Report 2003-03, CIS Dept., Univ. of Delaware, Sep. 2002. www.cis.udel.edu/∼steinder/PAPERS/TR-2003-03.pdf.

[12] M. Steinder and A. S. Sethi. Non-deterministic event-driven fault diagnosis through incremental hypothesis updating. In *Integrated Network Management VIII*, Colorado Springs, CO, Mar. 2003. (to appear).

[13] M. Steinder and A.S. Sethi. End-to-end service failure diagnosis using belief networks. In *Network Operation and Management Symposium*, Florence, Italy, Apr. 2002. pp. 375–390.

[14] S. A. Yemini, S. Kliger, E. Mozes, Y. Yemini, and D. Ohsie. High speed and robust event correlation. *IEEE Communications Magazine*, 34(5):82–90, 1996.