This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the ICC 2008 proceedings.

Cost-Efficient Transmitter/Receiver Deployment for Proactive Fault Diagnosis in All-Optical Networks

Yonggang Wen, *Student Member, IEEE*, Vincent W.S. Chan, *Fellow, IEEE*,
and Eric Swanson, Senior Member, IEEE
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
77 Massachusetts Avenue, Cambridge, MA 02139

***Abstract*— A scalable fault management system, including fault detection and localization capability, is crucial for future all-optical networks. In our previous work [5-7], we have proposed adaptive fault diagnosis schemes that deploy proactive lightpath probes to identify network failures, and have developed an asymptotically optimal run-length probing scheme to minimize the diagnostic effort (i.e., the number of lightpath probes). In this research, we aim to investigate the diagnostic hardware cost, i.e., the cost resulted from transmitter/receiver (Tx/Rx) pairs for probe transmission and detection. As a benchmark, we first show that, in order to identify all possible network failures, all the network nodes have to be equipped with diagnostic Tx/Rx pairs. We then develop a probabilistic analysis framework to characterize the trade-off between hardware cost (i.e., the number of nodes equipped with Tx/Rx pairs) and diagnosis capability (i.e., the probability of successful failure detection and localization). Our results suggest that, for practical situations, the hardware cost can be reduced significantly by accepting a small amount of uncertainty about the failure status.**

I. INTRODUCTION

All-optical networks [1-2], where data traverses lightpaths without any optical-to-electrical conversion at intermediate nodes, promise significant cost benefits that will enable broadband network services to be delivered to large populations at a much lower cost (per bit) than with today's technologies. The cost savings are mainly due to the replacement of electronic switching with optical switching of high date-rate lightpaths at intermediate nodes, thus eliminating the expensive process of optical-to-electrical-to-optical (OEO). However, all-optical networks are susceptible to various physical failures, e.g., fiber cuts, switch node failures, transmitter/receiver breakdowns, and optical amplifier breakdowns [3]. These failures can lead to costly disruption in communication, and their detection and localization can constitute a significant fraction of recurring networking operating costs [3]. Therefore, to ensure specified levels of quality of service at an affordable cost, an efficient network management system – including an efficient fault diagnosis function – should be in place when all-optical networks are fully deployed. In this work, we focus on developing cost-efficient fault diagnosis schemes, which detect and localize failures in the optical layer, for all-optical networks.

The significant difference between envisioned all-optical networks and current optical networks suggests that new fault diagnosis mechanisms are required. Presently, Synchronous Optical Network (SONET) and G.709 networks infer the health of each link by verifying the parity check bits embedded in the overhead of data frames at intermediate OEO nodes [3]. In some cases, optical spectrum analyzers embedded in nodes may try to infer the health of an optical link. These approaches are based on passive monitoring mechanisms that obtain network state information from existing traffics. However, such an approach cannot be extended to all-optical networks efficiently. If the passive monitoring scheme is to be deployed in all-optical networks, optical signals need to be tapped out at intermediate nodes for fault identification. This process of tapping out optical signals for fault detection and localization would diminish much of the cost benefits of all-optical networks. Moreover, because all of these measurements are piggybacked onto existing lightpaths, the states of infrequently used links might be obsolete when they are accessed. This will cause serious problems in some real-time applications with critical time deadlines [4].

Motivated by these shortcomings of the passive monitoring approach, we have proposed adaptive fault diagnosis schemes in [5-7], based on a proactive lightpath probing technique. In an all-optical network, optical signals traverse a lightpath without being detected by intermediate nodes. This property of all-optical networks permits lightpath probes to test the health of several links/nodes simultaneously. Our approach exploits this fact to develop cost-effective fault diagnosis schemes. Specifically, optical probe signals are sequentially sent along a set of lightpaths over an all-optical network to test their states of health; and the network state (i.e., the failure pattern) is then inferred from this set of end-to-end measurements. Each successive probe is dynamically chosen among the set of permissible probes according to the results of previous probe signals (i.e., probe syndromes), with the objective of minimizing the total number of probes.

We are interested in two design metrics for these proactive fault diagnosis schemes: the diagnostic effort (i.e., the number of lightpath probes) and the diagnostic transmitter/receiver (Tx/Rx) cost. Previously in [5-7], we have established a theoretical framework to minimize the number of probes and have developed asymptotically optimal fault diagnosis schemes (i.e., the run-length probing scheme). At the same time, the hardware cost, i.e., Tx/Rx pairs needed to transmit and detect optical probe signals, contributes a significant portion of the fault-diagnosis cost. In this paper, we aim to

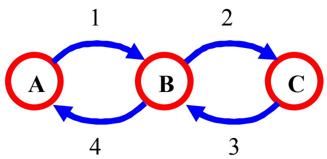

Fig. 1. A motivational example: the trade-off between the cumulative diagnosability probability and the number of nodes equipped with diagnostic Tx/Rx pairs.

characterize this Tx/Rx hardware cost and understand its implications on practical network design.

We develop a probabilistic framework to investigate the Tx/Rx cost for the proactive fault diagnosis paradigm. As a benchmark, we first show that all the network nodes should be equipped with diagnostic Tx/Rx pairs in order to identify all possible network states. This result prompts us to investigate the impact on the diagnosis performance when only a small fraction of nodes is equipped with diagnostic Tx/Rx pairs. Our analytical results suggest a trade-off between the number of nodes equipped with diagnostic Tx//Rx pairs and the diagnosis capability. The metric we employ for the diagnosis capability is the probability of all identifiable network states, defined as *the cumulative diagnosability probability*. This trade-off can be illustrated via the following example.

Consider a linear network with 3 nodes and 4 directed arcs, where nodes do not fail and arcs fail independently with probability p, as illustrated in Fig. 1. First, if only one node is equipped with a Tx/Rx pair (either A, B or C), one can only diagnose the network state with zero arc failure, and the cumulative diagnosability probability is $(1-p)^4$. Second, if two nodes are equipped with Tx/Rx pairs (e.g., node A and C), the identifiable network state set is {Φ, {1}, {2}, {3}, {4}, {1,2}, {3,4}, {1,4}, {2,3}}, where Φ denotes the network state with zero arc failures and {1} denotes the network state with arc 1 failure, and so on. In this arrangement, only a subset of the network states with two arc failures can be uniquely identified (e.g., network states {1,3} and {2,4} cannot be uniquely diagnosed). The cumulative diagnosability probability thus is $(1-p)^4+4p(1-p)^3+4p^2(1-p)^2$. Finally, if all the nodes are equipped with diagnostic Tx/Rx pairs, any network state can be identified and thus the cumulative diagnosability probability is 1. The cumulative diagnosability probability increases as the number of nodes equipped with Tx/Rx pairs increases.

This example suggests an opportunity to reduce the diagnostic Tx/Rx hardware cost, by accepting a reduced cumulative diagnosability probability. In particular, when the network is relatively reliable, only a small fraction of nodes equipped with Tx/Rx pairs is needed to provide high diagnosis fidelity. It follows that a significant portion of the worst-case fault-diagnosis hardware cost can be saved in exchange for an acceptable amount of uncertainty about the network's state.

This paper is organized as follows. In Section II, we present the proactive fault-diagnosis architecture for all-optical networks including the network model, the fault-diagnosis design metrics, and the probabilistic analysis framework. In Section III, we derive the cumulative diagnosability probability for any ring network by decomposing the network into a set of canonical linear networks with Tx/Rx pairs at both end nodes, and characterize the trade-off between the number of nodes equipped with Tx/Rx pairs and the cumulative diagnosability probability for ring networks. In Section IV, the trade-off for mesh networks is characterized via two alternative approaches: the cutset-based approach and the Euler-Trail-based approach.

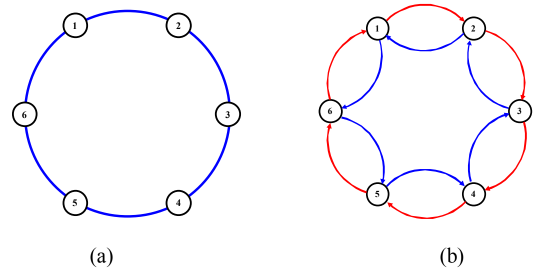

Fig. 2: Network topology for all-optical networks: (a) undirected graph, and (b) directed graph. Each undirected link in the undirected graph is replaced by two directed arcs in opposite directions, to illustrate bidirectional connections between adjacent nodes.

II. Proactive Fault Diagnosis Architecture for All-Optical Networks

A. Network Model

All-optical networks can be abstracted as undirected graphs. An *undirected graph* G is a pair of sets (V,E), where V is the set of network nodes of size n, and E is the set of optical links of size m. For example, Fig. 2(a) illustrates an optical network with 6 nodes arranged in a ring structure. However, in practice, connections between adjacent nodes are bidirectional and are usually achieved via two parallel optical fibers transmitting optical signals in opposite directions. To capture this practical constraint, we replace each undirected edge in the undirected graph with two directed arcs in opposite directions. It follows that the original undirected graph is transformed into a directed graph, as illustrated in Fig. 2(b). The number of arcs in the directed graph is $2m$.

We assume in this paper that nodes are invulnerable (the node failure case can be investigated via similar approaches as in [6]), and that arcs fail independently with probability p $(0\le p\le 0.5)$. Moreover, we assume that the state of an individual arc does not change over the duration of the fault-diagnosis process. Therefore, each arc state can be modeled by a Bernoulli random variable, taking value 1 with probability p for arc failure, and value 0 with probability $1-p$ for no failure. A network state $s\in S$ is referred to as a realization of all arc states, where $S=\{0,1\}^{2m}$ denotes the set of all possible network states.

To detect and localize possible arc failures, we adopt the adaptive fault-diagnosis paradigm, based on the proactive lightpath-probe mechanism developed in [5-7]. In particular, optical probing signals are sequentially sent along a set of permissible lightpaths in the network and network failures are identified through the set of probe results. The result of each probe is called the *probe syndrome*, denoted as $r_t = 0$ if all the arcs along the probed lightpath are UP (no failure) and the probe signal arrives successfully; and $r_t = 1$ if any of the arcs along the probed lightpath is DOWN (at least one failure) and the probe signal never reaches the destination. Moreover, each successive probe is determined according to previous probe syndromes, with the objective of minimizing some preset cost function.

B. Design Metrics for Fault Diagnosis Schemes

We are interested in two design metrics for fault diagnosis schemes: the diagnostic effort (i.e., the number of lightpath probes) and the diagnostic Tx/Rx hardware cost. Optical Tx/Rx pairs are used in the data plane for probe transmission and detection. This part of the diagnostic cost is a one-time cost and is proportional to the number of nodes equipped with diagnostic Tx/Rx pairs. The diagnostic effort indicates the effort expanded to scheduling, transmitting and detecting optical probes and reporting probe syndromes. The diagnostic effort is recurring and is proportional to the number of lightpath probes deployed to identify the network state.

For each design metric, there is a trade-off associated with it. When the diagnostic effort [5-7] is of interest, a trade-off exists between the diagnostic effort (i.e., the number of lightpath probes) and the diagnostic delay (i.e., the number of diagnostic steps), via three alternative diagnosis paradigms: adaptive diagnosis, non-adaptive diagnosis and multi-stage diagnosis. In this paper, our concern is to minimize the diagnostic Tx/Rx cost (i.e., the number of nodes equipped with diagnostic Tx/Rx pairs). Specifically, we investigate a trade-off between the fraction of nodes equipped with diagnostic Tx/Rx pairs and the cumulative diagnosability probability (i.e., the probability of successful diagnosis).

C. Probabilistic Analysis Framework

To identity all possible network states, any fault diagnosis scheme has to diagnose the network state with all the arcs failing simultaneously. This, in turn, requires the diagnosis scheme to be able to probe each directed arc individually, which can be achieved only if each node in the network is equipped with a pair of diagnostic transmitter and receiver. It follows that, for a network of n nodes and m links (or equivalently $2m$ arcs), the number of nodes equipped with diagnostic Tx/Rx pairs is

$$n_d = n, \tag{1}$$

in order to identify all possible network states. However, the cost of such a worst-case approach could be prohibitively high and limits its application for future all-optical networks.

The fact that the probability mass is not evenly distributed among all network states provides us an opportunity to reduce the diagnosis hardware cost, with little loss in diagnosis capability. Due to the probabilistic arc failure model, some network states can occur with extremely small probability. However, in the worst-case analysis, the diagnosis scheme has to identify these network states by paying a high cost. Here, we propose a probabilistic analysis under which the objective of fault diagnosis is to identify the majority of network states by deploying less Tx/Rx pairs than the number of nodes in the network. This is similar to the lossy source coding problem in Information Theory [9] by encoding only the "typical sets".

The probabilistic analysis works as follows. We assume that n_d nodes are equipped with diagnostic Tx/Rx pairs. The fraction of network nodes equipped with Tx/Rx pairs is then defined as

$$\eta = \frac{n_d}{n}, \tag{2}$$

where $0 < \eta < 1$. For a given subset of nodes equipped with diagnostic Tx/Rx pairs, the set of all network states, denoted as S, is partitioned into two mutually exclusive and collectively exhaustive subsets: the set of identifiable network states (S_I), and the set of unidentifiable network states (S_U), with $S = S_I \cup S_U$. We define the *cumulative diagnosability probability* as the sum probability of all the network states in the set of identifiable network states, i.e.,

$$\beta_D(n, n_d, p) = \sum\nolimits_{s \in S_I} \Pr(s), \tag{3}$$

where $\Pr(s) = p^i (1-p)^{2m-i}$ is the probability of any network state with $0 \le i \le 2m$ arc failures. Similarly, we define the *cumulative undiagnosability probability* as the sum probability of all the networks in the set of unidentifiable network states,

$$\alpha_F(n, n_d, p) = \sum\nolimits_{s \in S_U} \Pr(s). \tag{4}$$

The example of a 3-node linear network in Section I suggests a trade-off between the cumulative diagnosability probability (or the cumulative undiagnosability probability) and the number of nodes equipped with diagnostic Tx/Rx pairs. That is, the cumulative diagnosability probability increases as the number of nodes equipped with diagnostic Tx/Rx pairs increases and more network states can be identified uniquely. In the rest of this paper, we characterized this trade-off for ring networks and mesh networks, and develop useful insights for practical network design.

III. Efficient Tx/Rx Deployment for Ring Networks

In this section, we present a systematic approach to calculate the cumulative diagnosability probability for any ring network with a fraction of nodes equipped with Tx/Rx pairs, by decomposing the network into a set of canonical linear networks, both end nodes of which are equipped with diagnostic Tx/Rx pairs. For example, in Fig. 2(b), if node 1 and node 4 are equipped with diagnostic Tx/Rx pairs, the network can be decoupled into two canonical linear networks, i.e., 1-2-3-4 and 4-5-6-1. In both canonical linear works, only end nodes are equipped with diagnostic Tx/Rx pairs. Therefore, we can first derive the cumulative diagnosability probability of canonical linear networks, and then synthesize

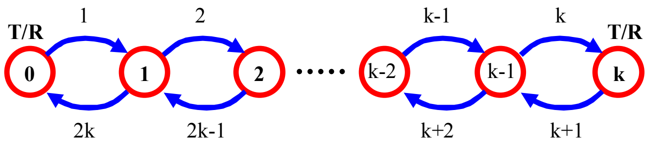

Fig. 3. The canonical linear network with k+1 nodes and 2k arcs: nodes at both ends are equipped with diagnostic Tx/Rx pairs.

the cumulative diagnosability probability for any ring network with a subset of nodes equipped with Tx/Rx pairs. Using this result, we then characterize the trade-off between the target cumulative diagnosability probability and the required fraction of nodes equipped with diagnostic Tx/Rx pairs.

A. Canonical Network Analysis: Linear Network with Diagnostic Tx/Rx Pairs at Both End Nodes

In this subsection, we consider a canonical linear network with $k+1$ nodes and $2k$ unidirectional arcs. As illustrated in Fig. 3, only the two end nodes (i.e., node *0* and node *k*) are equipped with diagnostic Tx/Rx pairs.

Only a subset of network states can be identified uniquely as a result of only two diagnostic Tx/Rx pairs. We have looked at the case of $k=2$ in the motivational example. Generalized from this simple case, for the canonical linear network, any network state with more than three arc failures can not be uniquely identified. Among the set of network states with two or less arc failures, three types of failure patterns can be identified with adaptive fault diagnosis schemes:

1. The first type of identifiable failure patterns contains network states with zero arc failure. The number of network states in the first category is 1 and the probability of that network state is $(1-p)^{2k}$. As illustrated in Fig. 4(a), this network state can be uniquely identified by two probes from node *0* to node *k* and from node *k* to node *0*.
2. The second type of identifiable failure patterns contains network states with a single arc failure. The number of network states in the second category is $2k$ and the probability of such network state is $p(1-p)^{2k-1}$. As illustrated in Fig. 4(b), any network state in this category can be uniquely identified by two probes from node *0* to node *k* and from node *k* to node *0* to detect, followed by a binary searching algorithm to localize.
3. The third type of identifiable patterns contains a subset of the network states with two arc failures. In particular, among all the $k(k-1)/2$ network states with two arc failures, the following two classes of failure patterns are diagnosable, i.e., failure patterns with two arc failures in both directions of one bidirectional link (i.e., arc failures at {1, 2k}, {2,2k-1,…, {k-1,k+2} or {k,k+1}}, and failure patterns with two arc failures in two consecutive arcs in the same direction (i.e., arc failures at {1,2}, {2,3}, …, {k-1,k}, {k+1, k+2}, …, {2k-1,2k}. As illustrated in Fig. 4(c) and (d), any network state in this category can be uniquely identified by two probes from node *0* to node *k* and from node *k* to node *0* to detect, followed by two binary searching procedures from both ends to localize. The total number of network states in the third category is $3k-2$ and the probability of such network state is $p^2(1-p)^{2k-2}$.

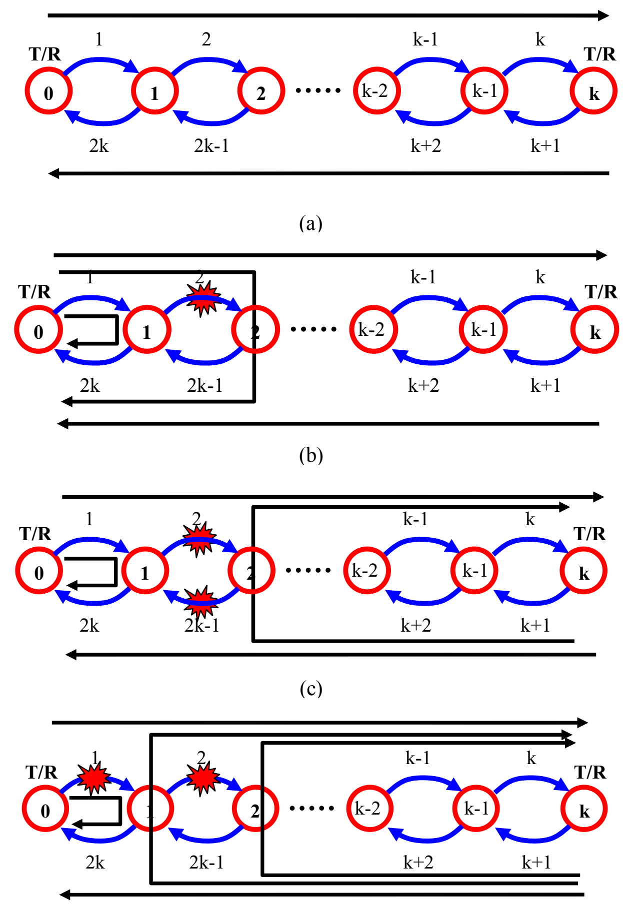

(d)

Fig. 4. Diagnosis schemes for canonical linear networks under three types of diagnosable failure patterns: (a) failure pattern with zero arc failure, (b) failure patterns with single arc failure, (c) failure patterns with two arc failures in both directions of one bidirectional link and (d) failure patterns with two arc failures in two consecutive arcs in the same direction.

The cumulative diagnosability probability for the canonical linear network is thus equal to the sum probability of all the identifiable network states,

$$\beta_D^{\dagger}(k,p) = (1-p)^{2k} + 2kp(1-p)^{2k-1} + (3k-2)p^2(1-p)^{2k-2}, \quad (5)$$

for $k \geq 1$ and $0<p<1$. Notice that the ratio between the number of identifiable network states with two arc failures and the number of network states with two arc failures is on the order of $1/k$ (i.e., $(3k-2)/(k(k-1)/2)$). When the arc failure probability is small, the contribution of the subset of identifiable network states with two arc failures is negligible. However, when the arc failure probability is high, we need to keep the length of the canonical linear network small so that

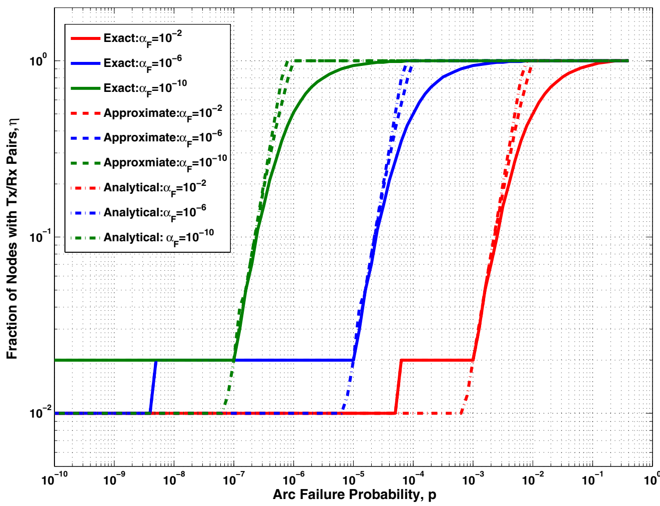

Fig. 5. The required fraction of nodes with Tx/Rx pairs for different target cumulative diagnosability probabilities is plot against the arc failure probability. They share similar "S" shapes.

the contribution of this subset of identifiable network states with two arc failures is kept insignificant.

B. Cumulative Diagnosability Probability for Ring Networks

In this subsection, the cumulative diagnosability probability for a ring network is derived by decomposing it into a set of canonical linear networks.

Consider a ring network with n nodes, among which a subset of n_d nodes are equipped with diagnostic Tx/Rx pairs. The ring network is decoupled into n_d canonical linear networks, both end nodes of which are equipped with diagnostic Tx/Rx pairs. We denote the length of each canonical linear network as k_i, $i=1,2,\cdots,n_d$. Using the cumulative diagnosability probability for the canonical linear network, we can synthesize the cumulative diagnosability probability for the ring network as

$$\beta_D(n,n_d,p)=\prod_{i=1}^{n_d}\beta_D^{\dagger}(k_i,p). \tag{6}$$

For a given number of nodes equipped with Tx/Rx pairs, it is natural to maximize the cumulative diagnosability probability by optimally distributing them among all the network nodes. We have not yet derived the optimum distribution, but have assumed that the set of n_d diagnostic Tx/Rx pairs are evenly distributed among all the network nodes and derive the cumulative diagnosability probability under such a deployment policy. Although the uniform distribution policy may not be optimal, it is a reasonable starting point, especially for symmetric graphs.

Under the uniform Tx/Rx deployment policy, the length of each decomposed canonical linear networks is made as equal as possible and the length of each canonical linear network could be k^* and k^*+1, where $k^*=\lfloor n/n_d \rfloor$. Moreover, the number[1] of decomposed canonical linear networks with length k^* is $(k^*+1)n_d-n$, and the number of decomposed canonical linear networks with length k^*+1 is $n-k^*n_d$. Notice that, when n/n_d is an integer, all the decomposed canonical linear networks have the same length of k^*. It follows that the cumulative diagnosability probability is given by

$$\beta_D(n,n_d,p)=\left\{\beta_D^{\dagger}(k^*,p)\right\}^{(k^*+1)n_d-n}\cdot\left\{\beta_D^{\dagger}(k^*+1,p)\right\}^{n-k^*n_d}, \tag{7}$$

where the first term results from the decomposed canonical linear networks of length k^* and the second term is due to the decomposed canonical linear networks of length k^*+1.

In practice, the cumulative diagnosability probability of (7) can be further approximated as a function of the fraction of nodes equipped with Tx/Rx pairs. For the special case that n/n_d is an integer, the cumulative diagnosability probability of (7) would be reduced to be $\beta_D(n,\eta,p)=\left\{\beta_D^{\dagger}(\eta^{-1},p)\right\}^{n\eta}$. In general, using the approximation $\beta_D^{\dagger}(\hat{k},p)\approx\beta_D^{\dagger}(\hat{k}+1,p)$, we can approximate the cumulative diagnosability probability as

$$\beta_D(n,\eta,p)\approx\left\{\beta_D^{\dagger}(\eta^{-1},p)\right\}^{n\eta}. \tag{8}$$

Therefore, for the rest of this paper, we use (8) to approximate the cumulative diagnosability probability for ring networks.

C. Diagnostic Cost-Performance Trade-Off

In this sub-section, we characterize the trade-off between the diagnostic hardware cost (i.e., the number of nodes equipped with diagnostic Tx/Rx pairs) and the diagnostic performance (i.e., the cumulative diagnosability probability). Our results demonstrate that the diagnostic hardware cost can be reduced significantly by accepting some reasonable amount of uncertainty about the network state.

For practical engineering design, we would like to calculate the fraction of nodes equipped with Tx/Rx pairs required to provide a target cumulative diagnosability probability (or a tolerable cumulative undiagnosability probability). Indeed, for a given cumulative diagnosability probability of β_D, we can identify the minimum fraction of nodes equipped with diagnostic Tx/Rx pairs by exhaustively searching over (8).

In Fig. 5, we plot the required faction of nodes equipped with diagnostic Tx/Rx pairs, for different target cumulative diagnosability probabilities, as a function of the arc failure probability, for a ring network with 100 nodes. Notice that all the curves share a similar "S" shape. In one extreme, when the arc failure probability is small, the number of nodes with

[1] Let x and y be the number of canonical linear networks with length k^* and k^*+1, respectively. First, because the total number of canonical linear network is n_d, we have $x+y=n_d$. Second, because the total number of segments is n, we have $xk^*+y(k^*+1)=n$. Solving these two equations, we obtain $x=(k^*+1)n_d-n$ and $y=n-k^*n_d$.

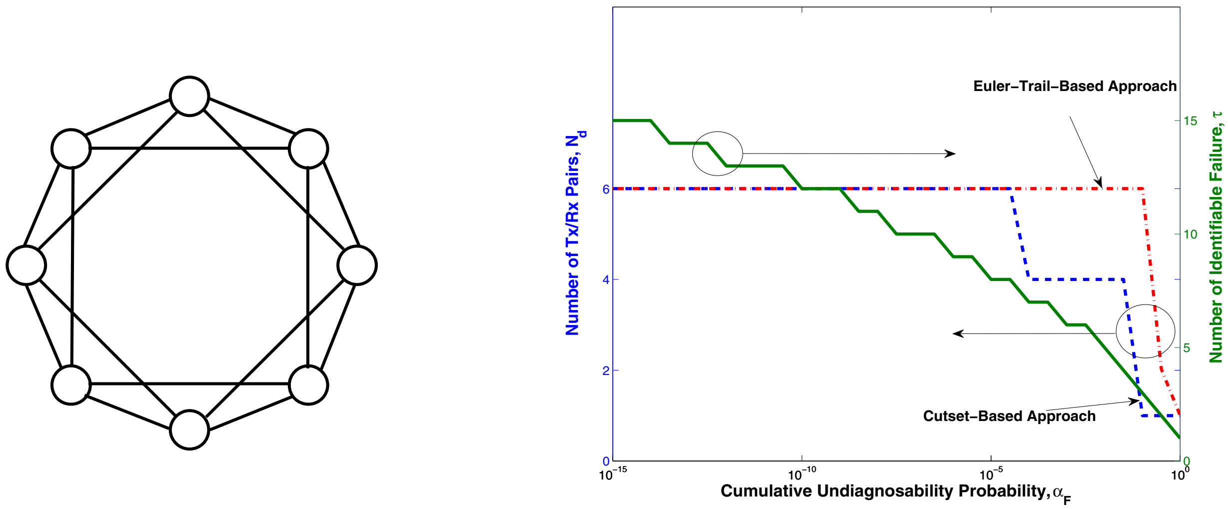

(a) (b)

Fig. 6. (a) Harary graph with 8 nodes and 16 links. (b) The number of identifiable link failures and the required number of nodes equipped with diagnostic Tx/Rx pairs as a function of cumulative undiagnosability probability.

Tx/Rx pairs is either 1 or 2. In the other extreme, when the arc failure probability is high, the required fraction of nodes equipped with Tx/Rx pairs is close to 1. Between these two extreme cases, there is a transition phase from a small fraction of nodes equipped with diagnostic Tx/Rx pairs to a large fraction of nodes equipped with diagnostic Tx/Rx pairs.

These observations can be understood as follows. The cumulative diagnosability probability in (8) can be expanded as

$$\beta_D(\eta) = (1-p)^{2n} + 2np(1-p)^{2n-1} + \mathbf{O}\left(p^2\right), \qquad (9)$$

where $\mathbf{O}\left(p^2\right)$ denotes a polynomial of p with an order of at least 2. Notice that each term in (9) corresponds to one class of identifiable network states. The first term of $(1-p)^{2n}$ corresponds to the subset of network states with zero arc failure. The second term of $2np(1-p)^{2n-1}$ corresponds to the subset of network states with a single arc failure. The third term corresponds to the subset of network states with two or more arc failures. The significance of these terms depends on the arc failure probability.

In one extreme, when the arc failure probability is small, the cumulative diagnosability probability is first dominated by the first term and then by the first two terms. In the former case, when the target cumulative diagnosability probability is less than $(1-p)^{2n}$, it is sufficient to diagnose the network state without any arc failure with one diagnostic Tx/Rx pair. In the latter case, when the target cumulative diagnosability probability is less than the sum of the first two terms, it is sufficient to diagnose the subset of network states containing zero or a single arc failure, achieved by two diagnostic Tx/Rx pairs. Therefore, there exist two thresholds as the arc failure probability increases, as shown in Fig. 5.

In the other extreme, when the arc failure probability is high, the probability mass of all network states is mostly contributed by networks states with multiple arc failures. In this case, almost all of the nodes have to be equipped with Tx/Rx pairs in order to identify the subset of network states with multiple arc failures.

Between these two extreme cases, for a target cumulative diagnosability probability, the required fraction of nodes equipped with Tx/Rx pairs increases as the arc failure probability increases. In this regime, we first hypothesize that, the cumulative diagnosability probability in each decomposed canonical linear network is dominated by the subset of network states with zero and a single arc failure. To verify this hypothesis, we can approximate the cumulative diagnosability probability in (8) as,

$$\beta_D(\eta) \approx \left\{(1-p)^{2\eta^{-1}} + 2\eta^{-1}p(1-p)^{2\eta^{-1}-1}\right\}^{n\eta}, \qquad (10)$$

where the contribution from the subset of identifiable network states with two arc failures in each decomposed linear network is suppressed. In Fig. 5, we also plot the fraction of nodes equipped with Tx/Rx pairs to provide a target cumulative diagnosability probability, obtained by exhaustively searching over (10). We observe that, the approximation is very close to the exact solution derived from (8), especially when the arc failure probability is small. With this approximation, the required fraction of nodes equipped with Tx/Rx pairs can be derived by solving the following equation,

$$\left\{(1-p)^{2\eta^{-1}} + 2\eta^{-1}p(1-p)^{2\eta^{-1}-1}\right\}^{n\eta} = 1-\alpha_F, \qquad (11)$$

for a tolerable cumulative undiagnosability probability of α_F. Solving (11) by using a second-order Taylor expansion, we can approximate the required fraction of nodes equipped with diagnostic Tx/Rx pairs as

$$\eta^*(\alpha_F) \approx \frac{2np^2}{\alpha_F}, \qquad (12)$$

for small α_F and $1/n \le 2np^2/\alpha_F \le 1$. In Fig. 5, we also plot the required fraction of nodes equipped with diagnostic Tx/Rx pairs to provide a tolerable cumulative undiagnosability probability, based on the analytical result in (12). Notice that the analytical result matches the exact solution from (8) closely, especially when the arc failure probability is small.

IV. Efficient Tx/Rx Deployment for Mesh Networks

In this section, we address the problem of efficient Tx/Rx deployment for mesh networks via two alternative approaches. One approach progressively identifies all the network states with up to κ link failures. The other approach extends our results for ring networks to Eulerian graphs.

A. Cutset-based Approach

For a given mesh network of n nodes and m links, we can order all the network states, based on the number of link failures contained in the network state, into a sequence of disjoint subsets, $S_0, S_1, \cdots, S_m$, where S_i denotes the set of network states containing i link failures. The sum probability of all the network states in state set S_i is

$$P_i = \binom{m}{i} p^i (1-p)^{m-i}, \tag{13}$$

for $i = 0, 1, \cdots, m$.

For a target cumulative diagnosability probability β_D, starting from the set S_0, we can progressively identify all the network states with up to κ link failures by deploying more diagnostic Tx/Rx pairs, so that the sum probability of all the identifiable sets of network states is larger than the target diagnosability probability, i.e.,

$$\sum_{i=0}^{\kappa} P_i \ge \beta_D. \tag{14}$$

Solving inequality (14) numerically, we can obtain the number κ so that all the networks states with up to κ link failures are identifiable.

The number of Tx/Rx pairs required can be determined by the following cutset approach. In order to identify the network state set S_i (i.e., all the network states with i link failures), we need to deploy one Tx/Rx pair in each nontrivial subgraph (i.e., containing at least one link), resulted from any edge cutset of order i. Otherwise, it is not possible to uniquely identify the states of some edges in the cutset. It follows that the efficient Tx/Rx deployment problem can be translated into the following combinatorial problem: *for an integer number κ, what is the minimum set of nodes in a graph such that a node from the minimum set exists in each nontrivial subgraph resulted from any cutset with an order up to κ?*

As an example, we consider a Harary graph with 8 nodes and node degree $d = 4$ as illustrated in Fig. 6(a). We plot the required number of nodes equipped with diagnostic Tx/Rx pair, n_d, and the number of identifiable link failures, κ, as a function of the cumulative undiagnosability probability α_F in Fig. 6(b). We can see that the required number of nodes equipped with diagnostic Tx/Rx pairs decreases as the cumulative undiagnosability probability increases.

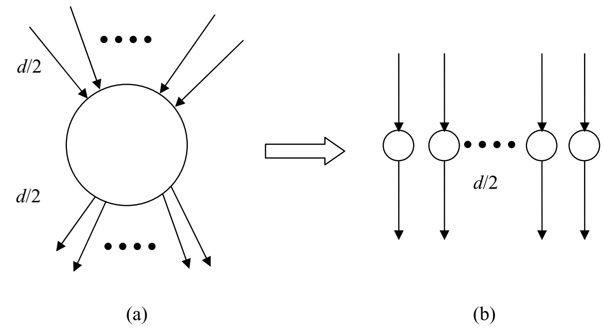

Fig. 7. Node replication approach: (a) a node of degree d has d/2 in degree and d/2 out degree, (b) the node is replicated with d/2 virtual nodes, each of which has 1 in degree and 1 our degree.

However, for a generalized mesh network, this problem is a NP-hard problem (in the worst case) in graph theory[2], and thus we will seek an alternative approach, based on ring network results, in the next sub-section.

B. Euler-Trail-based Approach

The analysis for ring networks can be extended to derive performance bounds for network topologies with an embedded ring structure, such as Eulerian graphs (an Eulerian graph contains a path that passes through all the links without repetition.) Non-Eulerian graphs can be approximated well with Eulerian graphs by a path augmentation approach [7].

In particular, as illustrated in Fig. 7, all the links in an Eulerian graph can be re-arranged into a ring network by replicating each node with $d/2$ virtual nodes, where d is the node degree. Under such a network transformation, our analysis for ring networks can be applied directly. However, the transformation suppresses a rich set of possible probing paths in the original network. It follows that the derived cumulative diagnosability probability is a lower bound, i.e.,

$$\beta_D(\eta) \ge \left\{(1-p)^{\eta^{-1}} + \eta^{-1} p (1-p)^{\eta^{-1}-1}\right\}^{\frac{1}{2}nd\eta}, \tag{15}$$

for the transition phase. Due to the bidirectional graph model used here, (15) is different from (10). This result, in turn, suggests that the resulting fraction of nodes equipped with Tx/Rx pairs for any target cumulative diagnosability probability is an upper bound on the required fraction of nodes equipped with diagnostic Tx/Rx pairs, i.e.,

$$\eta^*(\alpha_F) \le \frac{ndp^2}{4\alpha_F}, \tag{16}$$

where α_F is the tolerable cumulative undiagnosability probability. Notice that the required fraction of nodes

[2] When $\kappa \ge 4$, this problem can be converted into the vertex cover problem in graphs with maximum vertex degree of 3, since any edge can be turned into a connected component by deleting all other edges adjacent to its endpoints.

equipped with diagnostic Tx/Rx pairs decreases as the cumulative diagnosability probability increases. The result of (16) is plot in Fig. 6(b) for the Harary graph of 8 nodes and node degree 4. Notice that the number of nodes equipped with Tx/Rx pairs is larger than the result from the cutset approach, because rich connection in mesh networks is not exploited in the Euler-Trail-based approach.

The tightness of these performance bounds depends on both the arc failure probability and the node degree. When the arc failure probability is small and/or the node degree is small, these bounds are expected to be tight. When the arc failure probability is small, the cumulative diagnosability probability in each decomposed network is dominated by network states with zero and a single arc failure. When the node degree is small, the benefit of additional node degree is not significant enough to change the order of magnitude. However, when the node degree is large, these bounds could be loose. In this case, the rich set of connections in the mesh network of degree larger than 2 should be explored to identify failure patterns with multiple arc failures, and thus reduce the number of nodes equipped with Tx/Rx pairs.

V. Conclusion

In this paper, we built upon our previous research on proactive lightpath probing schemes to investigate the cost-effective Tx/Rx deployment for probe transmission and detection in all-optical networks. We developed a probabilistic framework to characterize the trade-off between the number of nodes equipped with diagnostic Tx/Rx pairs and the cumulative diagnosability probability. Our investigation suggested that the diagnostic hardware cost can be reduced significantly by accepting a reasonable amount of uncertainty about network failure status.

Acknowledgement

The research in this paper was supported in part by the Defense Advanced Research Projects Agency (DARPA) and the National Science Foundations (NSF). The authors would like to thank Prof. Muriel Medard, Prof. Lizhong Zheng and Guy Weichenberg at MIT for insightful discussions.

References

[1] V. W. S. Chan, et al, "A precompetitive consortium on wide-band all-optical networks," IEEE/LEOS Journal of Lightwave Technology, Vol. 11, No. 5/6, May/June, 1993, pp. 714-735.

[2] V. W. S. Chan, "All-Optical Networks," Scientific American 273, No. 3, September 1995, pp. 56-59.

[3] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*, 2nd edition, Morgan Kaufmann Publishers, 2002.

[4] A. Saleh, "Dynamic multi-terabit core optical networks: architecture, protocols, control and management (CORONET), Proposer Information Pamphlet for Defense Advanced Research Project Agency Strategic Technology Office, 2006.

[5] Y.G. Wen, V. W. S. Chan and L. Z. Zheng, "Efficient Fault Diagnosis Algorithms for All-Optical WDM Networks with Probabilistic Link Failures (*invited paper*)," *IEEE/OSA Journal of Lightwave Technology*, Vol. 23, No. 10, October 2005, pp.3358-3371.

[6] Yonggang Wen, Vincent W. S. Chan, and Lizhong Zheng, "Efficient Fault Detection and Localization for All-Optical Networks," GlobeCom'06, San Francisco, CA, December 2006.

[7] Yonggang Wen, Vincent W. S. Chan and Lizhong Zheng, "Efficient Fault Diagnosis Algorithms for All-Optical Networks: An Information Theoretic Approach," ISIT'06, Seattle, WA, July 2006.

[8] M. Subramanian, *Network Management: Principles and Practice*, 1st edition, Addison-Wesley Lognman, 2000.

[9] Thomas M. Cover and Joy A. Thomas, Elements of Information Theory, John Wiley & Sons Ltd., 1991.