

Efficient Fault-Diagnosis Algorithms for All-Optical WDM Networks With Probabilistic Link Failures

Yonggang Wen, *Student Member, IEEE*, Vincent W. S. Chan, *Fellow, IEEE, Fellow, OSA*,
and Lizhong Zheng, *Member, IEEE*

Invited Paper

Abstract—This paper investigates the fault-diagnosis problem for all-optical wavelength-division-multiplexing (WDM) networks. A family of failure-localization algorithms that exploit the unique properties of all-optical networks is proposed. Optical probe signals are sequentially sent along a set of designed lightpaths, and the network state is inferred from the result of this set of end-to-end measurements. The design objective is to minimize the diagnosis effort (e.g., the average number of probes) to locate failures. By establishing a mathematical equivalence between the fault-diagnosis problem and the source-coding problem in information theory, we obtain a tight lower bound for the minimum average number of probes per edge (of the network modeled as a graph) as $H_b(p)$, the entropy of the individual edges. Using the rich set of results from coding theory to solve the fault-diagnosis problem, it is shown that the “ 2^m -splitting” probing scheme is optimum for the special case of single failure over a linear network. A class of near-optimum run-length probing schemes that have low computation complexity is then developed. Analytical and numerical results suggest that the average number of probes per edge for the run-length probing scheme is uniformly bounded above by $(1 + \epsilon)H_b(p)$ and converges to the entropy lower bound as the failure probability decreases. From an information-theoretic perspective, it is shown that the run-length probing scheme outperforms the greedy probing scheme of the same computational complexity. The investigation reveals a guideline for efficient fault-diagnosis schemes: Each probe should provide approximately 1 bit of information, and the total number of probes required is approximately equal to the entropy of the state of the network. This result provides an insightful guideline to reduce the overhead cost of fault management for all-optical networks and can further the understanding of the relationship between information entropy and network management. Several practical issues are also addressed in the implementation of run-length probing schemes over all-optical WDM networks.

Index Terms—All-optical networks, fault diagnosis, fault management, network management, run-length code.

I. INTRODUCTION

ALL-OPTICAL networks [1], [2], where data traverses along lightpaths without any optical-to-electrical conversion, will be increasingly prevalent in future broadband networks due to its inherent large transmission bandwidth, lower cost, and transparency to different signal formats and

communication protocols. However, similar to other networks, all-optical networks are also vulnerable to failures [3], such as fiber cuts and transmitter/receiver breakdowns. Moreover, there are new types of failures that are unique to all-optical networks—failures related to subtle changes in signal power, optical signal-to-noise ratio, crosstalk, Kerr effects, or other nonlinear effects. These failures can result in the disruption of communication and can be difficult to detect, localize, and repair. Hence, when parts of a network are malfunctioning, it is critical to locate and identify these failures as soon as possible. At the same time, the cost to detect and locate failures must be small to keep the network cost low. This paper proposes a family of efficient fault-diagnosis algorithms that exploit the unique property of all-optical wavelength-division-multiplexing (WDM) networks, where optical signals are not usually detected at intermediate nodes along lightpaths (mostly for cost reasons).

According to the scale of their effect, failures in all-optical WDM networks can be classified into two categories. One category is the wavelength-level failure, which impacts the quality of transmission of each individual lightpath, e.g., transmitter/receiver failures in the case of one dedicated transmitter/receiver per wavelength, optical-filter failures, and individual-channel failures of a frequency-selective switch. The other category is the fiber-level failure, which affects all the lightpaths on an individual fiber, for example, fiber cuts, erbium-doped fiber amplifiers (EDFAs) breakdowns, and transmitter/receiver failures in the case of only one tunable transmitter/receiver per fiber. From a graph-theory perspective, we can attribute both categories to edge failures in a network (graph) topology. In this paper, we focus on the ON/OFF edge failure, which is modeled by a binary-value function of value 0 if the required quality of transmission is met and of value 1 otherwise. Besides, all the failures that do not belong to the same risk group are considered independent. In real life, there may be failure correlation among risk groups due to physical proximity or accessibility from the same malicious attack entry point. In those cases, the results in this paper can still provide very useful upper and lower bounds on the diagnostic effort required to localize the failures.

Since the fault-diagnosis problem [4] was first proposed in 1967, it has been investigated extensively in electrical networks under a system-diagnosis context [5]–[7]. In this context, most current research is focused on a “single-hop”-test model, i.e.,

Manuscript received March 2, 2005. This work was supported in part by the Defense Advanced Research Projects Agency (DARPA).

The authors are with the Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology (MIT), Cambridge, MA 02139 USA (e-mail: eewyg@mit.edu; chan@mit.edu; lizhong@mit.edu).

Digital Object Identifier 10.1109/JLT.2005.855695

signals are transmitted between adjacent nodes to determine whether failure occurs on the edge connecting them. The result of each test can be represented as one bit of diagnosis information: 1 or 0, corresponding to “failure” or “no failure.” Indeed, for a Synchronous Optical Network (SONET), each SONET link (single hop) checks the health of the link using parity checks within the SONET receiving chips. However, in all-optical networks, this single-hop-test assumption will usually not be applicable due to the unique property that optical signals are not typically detected at every (optically switched) node along the lightpath.

For SONET networks, the network-management system employs mechanisms such as BER measurement, optical trace, and alarm management to perform fault management at each regenerator. In particular, these functionalities may be carried over various types of optical-layer overhead [8], including pilot tone, subcarrier-modulated overhead, optical supervisory channel, rate-preserving overhead, and digital-wrapper overhead. To some degree, all these overheads are detected at some intermediate nodes along the lightpath. This, in fact, breaks the spirit of the transparency paradigm and adds to the complexity and cost of future all-optical networks that do not need signal detection along a lightpath.

Currently, to diagnose failures in all-optical WDM networks, researchers typically consider an optical (channel) performance-monitoring solution, where optical performance monitors are employed at a set of network nodes to watch for possible failures and report them to the network-management system [9]. However, little work has been done to quantify the overhead cost that this monitoring solution might incur. Instead, most research literatures [10], [11] follow essentially the same design approach as their electrical counterparts, implicitly assuming that each network node, or even each active optical component such as optical amplifiers and optical add-drop multiplexers (OADMs), is equipped with a performance-monitoring module that is active and reporting all the time. While this is an acceptable solution in the near term since signal detection comes for free at every regeneration point, it is desirable to develop more efficient and less costly methods when the all-optical network paradigm is fully implemented and the network size grows significantly. Reduced complexity is good for the following reasons. First, the total amount of monitored information and signaling grows quickly with the number of network elements (i.e., network nodes and edges). The huge amount of management information, together with faster switching speeds in the network, complicates the network-management system and stresses the limited capability of current network processing units, and a mechanism based on constant sensing and reporting of numerous individual active monitors does not scale well with the size and tuning agility of future all-optical networks. Second, since each monitor only tests one component without taking into consideration its failure statistics, the diagnosis overhead cost [e.g., the required number of tests per unit time with the interval between monitoring drawn from quality of service (QoS) specifications] of such a mechanism can be prohibitively high, limiting the efficacy and ultimately ubiquitous deployment of all-optical networks.

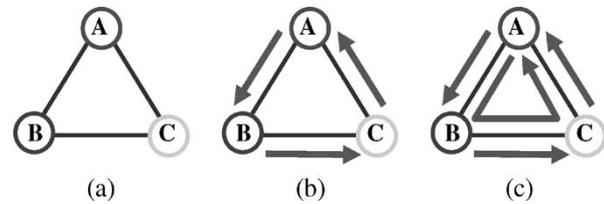


Fig. 1. Comparison between diagnosis paradigms of electrical networks and all-optical networks: (a) three-node ring network; (b) diagnosis with three single-hop tests; and (c) diagnosis with one three-hop test and three single-hop tests.

This paper seeks more efficient and elegant methods that greatly lower the hardware and computational complexity and cost of such functions for future all-optical WDM networks. Specifically, a family of failure-identification algorithms are developed that exploit the unique properties of all-optical networks to reduce the average number of diagnostic probes per unit time. In particular, optical signals will be sequentially sent along a set of lightpaths over an all-optical network to probe its state of health. The network state (i.e., failure pattern) is then inferred from the “syndromes” of this set of end-to-end measurements. To keep the required number of probes small, each successive probe is dynamically chosen among the set of permissible probes according to the results of the previous tests. Under this generalized model, the traditional diagnosis mechanism based on single-hop probes is then a special case and will be proven to be rather inefficient compared with the proposed designs.

In all-optical networks, the fact that optical signals can be carried over a lightpath of a number of interconnected edges without necessarily being detected by the intermediate nodes allows “multihop” tests to probe several edges simultaneously. This technique can be used to greatly reduce the amount of diagnosis effort, as illustrated with the three-node ring network in Fig. 1. In this example, we assume that each edge fails independently with probability of 0.1. If only single-hop tests are allowed, as in Fig. 1(b), the total number of tests to identify all edge states is three, by employing three single-hop tests (A-B, B-A, C-A). Note that the number of tests required is independent of the edge-failure probability. On the other hand, if multihop tests are allowed, we can first perform a three-hop test (A-B-C-A) as shown in Fig. 1(c). With a probability of $0.9^3 = 0.729$, it will be found that all edges are fault free and the diagnosis is concluded with only one test. We will resort to the single-hop tests only if we know there is at least one failure from the result of the first test, which has a probability of $1 - 0.729 = 0.271$. Thus, on the average, it requires only $0.729 \times 1 + 0.271 \times (1 + 3) = 1.813$ tests to fully diagnose this network. Intuitively, in most cases, the probability that a particular edge has failed is low when network diagnosis is performed; hence, it makes sense to test several edges together. Here, reducing the average number of tests required for network diagnosis, which is used in this paper as a measure of the cost, or efficiency, of the diagnosis process, is equivalent to reducing the network-diagnosis information bits [14] obtained in the probing sequence.

This example suggests that the fault-diagnosis problem can be better understood from an information-theoretic perspective. The network state can be viewed as a collection of binary-valued random variables, each associated with an edge in the network, indicating failure/no failure on that edge. The objective of a fault-diagnosis algorithm is to use a number of tests, whose results, also called the “syndromes,” can be used to uniquely identify the network state. To put it simply, we use probes to dig out all the information hidden in the unknown network state. In the above example, with the single-hop tests, the result of each test is “0” (for no failure) with a probability of 0.9 and “1” (for failure) with a probability of 0.1. Thus, the information about the network state contained in this test result is the entropy $H_b(0.1) = 0.469$ bit, where $H_b(x) = -x \log_2 x - (1-x) \log_2(1-x)$. In comparison, the three-hop test (A-B-C-A) contains $H_b(0.271) = 0.843$ bit of information. The information contained in a three-hop test is obviously larger than that of the single-hop test, indicating that multihop tests are more informative than single-hop tests for this case. As a result, in the second approach, the network state can be identified by a smaller number of probes, or equivalently, the network state is represented by the test syndromes in a more efficient way (this can be understood as encoding the network state with syndromes). Similar approaches can be used to determine the efficiency of a probe in a general network. More important, the design of an efficient fault-diagnosis algorithm is thus similar to the well-studied source-coding problem, whose goal is to use the minimum average number of bits to represent the source, which is also a collection of random variables.

Applying the above approach under a probabilistic failure model, where each edge is assumed to fail independently with a prior failure probability, we obtain the following main results. First, we establish the mathematical equivalence between the fault-diagnosis problem and the source-coding problem, which indicates that the minimum average number of probes required is lower bounded by the entropy of the network states. Second, since the sequential diagnosis problem for general network topology is NP-complete, we develop a family of novel near-optimal polynomial-time algorithms based on run-length codes [12], whose performance asymptotically approaches the theoretical entropy limit for large networks. Analytical and numerical investigation reveals a guideline for efficient probing schemes: Each probe should be designed to provide approximately one bit of information on the network state, and the number of probes required is approximately equal to the information entropy of the network states. Finally, we apply the family of near-optimum fault probing schemes to different WDM network scenarios. This general approach provides good insights and can be generalized to study other network failure models, including failures on both nodes and edges, correlated failures, and transient failures.

This paper is organized as follows. In Section II, we formulate the fault-diagnosis problem under a probabilistic edge-failure model. In Section III, we establish the source-coding/fault-diagnosis equivalence, which suggests a tight lower bound for the minimum average number of probes required, and show an optimal probing scheme for the special case of

finding a single failure in the network. In Section IV, we develop the near-optimum run-length probing scheme and analyze its performance in closed form. In Section V, we compare the performance of the run-length probing scheme to that of the greedy probing scheme within the same framework. In Section VI, we address some practical issues in the implementation of the run-length probing scheme.

II. FAULT-DIAGNOSIS-PROBLEM FORMULATION

A. Probabilistic Edge-Failure Model

In this paper, all-optical networks are abstracted as undirected graphs. An undirected graph G is an ordered pair of sets (V, E) , where V is the set of nodes of size n , and E is the set of edges of size m . In this paper, we first focus on Eulerian network topologies that have at least one Euler trail [13], which is a sequence of interconnected edges containing all the edges in the topology without repetition. Our results obtained in this paper can also be generalized to non-Eulerian topologies.

In this paper, we characterize the vulnerability of future all-optical WDM networks by the following probabilistic edge-failure model.

- 1) Nodes are invulnerable (the vulnerable node case will be treated later).
- 2) Edges are vulnerable and assumed to fail independently with a prior probability of $p(0 \leq p \leq 1)$.
- 3) We assume that the states of the edges do not change over the duration of the fault-diagnosis process (or only nonergodic failures occur) [14].

For a given network topology, we label each edge along an Euler trail with an index $\beta = 1, 2, \dots, m$. The state of the β th edge is represented by a Bernoulli random variable F_β , called the edge state. We assume for the moment that the edge states $F_\beta, \beta = 1, \dots, m$ are statistically independent and identically distributed with $\Pr(F_\beta = 1) = p$ for an edge failure and $\Pr(F_\beta = 0) = q (= 1 - p)$ for no failure.

We refer to the network state as a realization of the set of edge states $\{F_\beta\}_{\beta=1}^m$, written as $s = f_1 f_2 \cdots f_m \in S = \{0, 1\}^m$. The set of all possible network states is denoted as S . Using the fact that all edges fail independently, we obtain the prior probability of a particular realization of the network state $s = f_1 f_2 \cdots f_m$ as the product of prior probabilities of all edges, i.e.,

$$\Pr(s) = p^{\sum_{\beta=1}^m f_\beta} q^{m - \sum_{\beta=1}^m f_\beta}. \quad (1)$$

B. Sequential Probing Model

In this paper, we focus on diagnosing network states via the measurements of end-to-end probing signals. Specifically, each probe corresponds to sending an optical signal along some lightpath. We will illustrate the probing model in this section.

A permissible probe over an Eulerian network topology is a trail (a sequence of adjoined edges without repetition) over the graph. For a finite network, we label each probe with an index $t \in T = \{1, 2, \dots, |T|\}$ where $|T|$, the cardinality of the set T , is the number of distinct probes over the network. As

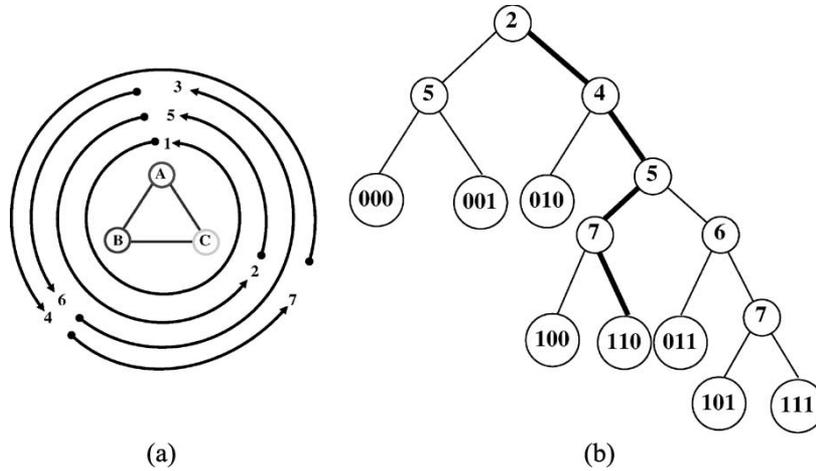


Fig. 2. (a) Set of permissible probes over the three-node ring topology. Total number of probes is 7. Each probe is indexed with a number near the arrow. (b) Probing scheme (decision tree) for the three-node ring topology.

an example, the three-node ring topology has seven permissible probes, as shown in Fig. 2(a).

When an optical signal is sent along a permissible probing trail, the signal either arrives at the destination when all the edges along the probe are ON or never reaches the destination (or the quality of the signal is unacceptable) when any of the edges along the probe is OFF. The result of each probe is called the probe syndrome, denoted as $r_t = 0$ if the probing signal arrives successfully; and $r_t = 1$ otherwise.

A probing scheme, denoted as π , is a sequential employment of probes such that any network state can be identified. The successive probe can be sequentially determined according to the syndromes of previous probes. Due to this sequential decision-making property, any probing scheme is equivalent to a binary decision tree, whose leaves are network states and inner nodes are probes. For example, a probing scheme for the three-node ring network is shown in Fig. 2(b), where each inner node is labeled with the probe employed. We adopt the convention that at any inner node, if the probe syndrome is 0 (no failure), the subsequent probe is given in the left child; otherwise, if the probe syndrome is 1, the probing process continues on the right child.

The set of all probing decision trees for the topology G is denoted as $\Pi(G)$. Without loss of optimality, we assume that any efficient probing scheme has the following properties.

- 1) A probe will not be employed if its syndrome can be inferred from previous syndromes. For example, if a probe returns no failure, it means that no edge in that probe has failed; hence, no probe that involves only a subset of these edges is performed thereafter.
- 2) When two probes are expected to reveal the same information, the probe with fewer hops is preferred. For example, if the state of an individual edge is known, then one should not start or end a probe with this edge, since dropping it loses no information.

C. Fault-Diagnosis Problem

To analyze the additional effort that different probing schemes incorporate into the network-management system, we need to associate a probing scheme with some cost metric.

In this paper, each probe $t \in T$, if employed, is assumed to cost one unit of diagnosis effort. Consequently, the probing cost of the state s , denoted by l_s^π , is equal to the number of probes from the root to the leaf node s in the probing decision tree π , called the probing depth of the state s . For example, as shown in Fig. 2(b), the probing depth of state 110 is 4.

Given a probing scheme $\pi \in \Pi(G)$, the average number of probes is

$$\mathcal{L}_\pi = \sum_{s \in S} \Pr(s) l_s^\pi \quad (2)$$

where l_s^π is the probing depth of network state s , and $\Pr(s)$ is the prior probability of this state given by (1). We observe that the average number of probes scales with the size of network topologies. In this research, to suppress the scaling effect, we focus on the average number of probes per edge, which is defined as

$$\bar{\mathcal{L}}_\pi = \frac{1}{m} \sum_{s \in S} \Pr(s) l_s^\pi \quad (3)$$

where m is the number of edges in the network topology. We use the average number of probes per edge as the cost metric to design optimum probing schemes.

For a given network topology $G(V, E)$, we want to find the optimal probing scheme that minimizes the average number of probes per edge, and thus, to minimize the network-monitoring overhead cost. Mathematically, it is formulated as the following optimization problem:

$$\begin{aligned} \min_{\pi} \bar{\mathcal{L}}_\pi &= \frac{1}{m} \sum_{s \in S} \Pr(s) l_s^\pi, \\ \text{s.t. } \pi &\in \Pi(G). \end{aligned} \quad (4)$$

The resulting minimum average number of probes per edge is written as

$$\bar{\mathcal{L}}^* = \min_{\pi \in \Pi(G)} \left\{ \frac{1}{m} \sum_{s \in S} \Pr(s) l_s^\pi \right\} = \frac{1}{m} \sum_{s \in S} \Pr(s) l_s^{\pi^*} \quad (5)$$

where π^* is the optimum probing decision tree.

III. OPTIMUM FAULT-DIAGNOSIS SCHEMES

In this section, we characterize some properties of the optimal probing schemes for Eulerian networks and the achievable performance. The insights developed in this section will provide guidance for designing near-optimum diagnosis schemes.

A. Source-Coding/Fault-Diagnosis Equivalence

In this section, we first build a connection between the fault-diagnosis problem and the source-coding problem. This connection inspires the use of the source-coding theory to design efficient network-diagnosis schemes.

Given a probing scheme $\pi \in \Pi(G)$, we denote the probe syndrome of network state s as $r(s) = r(t_1^s)r(t_2^s) \cdots r(t_{l_s^s}^s)$, where l_s^s is the probing depth of state s , and $\{t_1^s, t_2^s, \dots, t_{l_s^s}^s\}$ is the sequence of probes employed to identify state s . For example, the sequence of probes for state $s = 110$ in Fig. 2(b) is $\{2, 4, 5, 7\}$, and the probe syndrome is $r(s) = 1101$.

We observe that, given a probing scheme $\pi \in \Pi(G)$, there is a one-to-one correspondence between any network state s and its probe syndrome $r(s)$. Besides, we have the following conceptual mapping between the fault-diagnosis problem and the source-coding problem:

Network states \Leftrightarrow Source symbols

Prior probability of states \Leftrightarrow Prior probability of symbols

Probe syndromes \Leftrightarrow Coded symbols

Average # of probes \Leftrightarrow Average code length.

(6)

It follows that the set of probe syndromes constitutes a uniquely decodable code for the set of network states [15], summarized as follows.

Theorem 1: For any valid probing decision tree $\pi \in \Pi(G)$, the set of probe syndromes $R(\pi) = \{r(s), s \in S\}$ forms a uniquely decodable code.

This source-coding/network-probing equivalence has some important implications. First, our design objective, to minimize the average number of probes per edge, is the same as that of the source-coding problem: to use the minimum number of coded symbols to represent the source. It follows from the lossless source-coding theory [16] that the minimum average number of probes per edge is larger than the information entropy of the individual edge, i.e.,

$$\bar{L}^* \geq H_b(p). \quad (7)$$

Second, inspired by the source-coding problem, we observe that it is easier to focus on the diagnosis of large networks and develop simpler algorithms to find optimal probing schemes that minimize the average number of probes per edge as the network size, or the number of edges m , grows large. Note that this objective is weaker than the original optimization problem (4) because the result applies to large-size networks. However, from the engineering perspective, the difference between the two, say, a few more probes in the diagnosis of the

moderately sized network, is usually insignificant. On the other hand, the design of such near-optimal schemes can be much simpler.

Finally, this equivalence suggests an information-theoretic approach to translating existing source-coding algorithms into efficient network-diagnosis schemes. However, this translation is not trivial. In fact, not all source-coding algorithms can be transformed into fault-diagnosis algorithms. For instance, the well-known optimal source-coding algorithm, the Huffman code, is in general not applicable to fault-diagnosis problems. Huffman codes can be best understood as a sequence of yes/no questions, each of which corresponds to an inner node of the code tree. The optimality of the Huffman code comes from the optimal sequence of questions in the form of “Is the source realization in the set A?” Translated into the context of fault diagnosis, this corresponds to questions such as “Is edge 1 UP?” or “Is Edge 1 UP and edge 3 DOWN and edge 5 UP?” Clearly, not all of such questions are physically realizable probes, which can only test consecutive edges and ask questions such as “Are edges 2,3, and 4 all UP?” Thus, the nature of realizable probes posts an extra restriction that only a special class of questions can be asked. In the rest of this paper, we refer to this restriction as the “consecutive probing constraint” and study the fault-diagnosis problems, or the equivalent source-coding problems, under this constraint.

B. Edgewise Probing Schemes

As discussed in the Section I, the edgewise probing scheme, by probing each individual edge separately, is in general not optimal, especially when the failure probability of each edge p is small. On the other hand, if each edge fails with a high probability (which is unrealistic), the edgewise probing scheme becomes more efficient. Generalizing from the break-point theorem of the group-testing problem [23], we obtain the condition under which the edgewise probing scheme is optimal.

Theorem 2: For any nontrivial network topology with $m \geq 2$ edges and at least one path of more than one edge, the edgewise probing scheme is optimal if and only if the edge-failure probability is larger than $(3 - \sqrt{5})/2$.

The theorem suggests that the edgewise probing schemes based on single-hop tests, as used in the electrical networks and some of the current optical monitoring schemes, are strictly suboptimal in all-optical networks if $p < (3 - \sqrt{5})/2$ (≈ 0.382), which is the situation in most network-monitoring scenarios.

According to the theorem, the edgewise probing scheme is optimal for the case that $p > 1/2$. As p increases to 1, the lower bound $H_b(p)$ decreases, while the optimal approach is always one probe per edge. Intuitively, for large values of p , if all the edges fail with a high probability, we could reduce the number of probes required if there were a probe to test the scenario where a collection of edges are all in OFF states. Since such a probe cannot be implemented, our information-theoretical bound becomes less meaningful in the range of $p > 1/2$. Nevertheless, in almost all practical situations, the edge-failure probabilities are small; thus, in the remainder of this paper, we always assume $p \leq 1/2$.

C. Optimal Probing Schemes for Lightpath With Single Failure

In this section, as a special case to illustrate the technique, we focus on a linear network topology (i.e., bus) with h edges and only one faulty edge. In an all-optical network context, this can be understood as the case where there is only one failure along a particular lightpath. Conditioning on the fact that there is one and only one faulty edge, each edge has a uniform distribution of being the faulty one. For this case, the optimum probing scheme to minimize the average number of probes (per edge) has been found in [17].

The optimum probing scheme works as follows. Given that the linear network topology has h edges and the number of faulty edges is 1, we first split the path of length h into two subpaths of length h_1 and h_r according to the following criteria:

$$h_1 = g(h) = \begin{cases} 2^{\lfloor \log_2 h \rfloor - 1}, & \text{if } 2 \leq h \leq 3 \cdot 2^{\lfloor \log_2 h \rfloor - 1} \\ h - 2^{\lfloor \log_2 h \rfloor}, & \text{if } 3 \cdot 2^{\lfloor \log_2 h \rfloor - 1} < h \leq 2^{\lfloor \log_2 h \rfloor + 1} - 1 \end{cases} \quad (8)$$

and

$$h_r = h - h_1. \quad (9)$$

Now we probe the first subpath of length h_1 . If the syndrome is 1, meaning the faulty edge is in the first subpath, we continue to split the first subpath according to rule (8) and probe the resulting first subpath. If the syndrome is 0, meaning that the first subpath is fault free and the faulty edge is in the second subpath, we split the second subpath using (8) and probe the resulting first subpath. This process continues until the faulty edge is located.

Intuitively, when dividing the path into two subpaths of length h_1 and h_r , respectively, it is desirable to cut the path into equal halves; thus, the probability that the faulty edge is in the first subpath is as close to 1/2 as possible. However, this approach is in fact only locally optimal: it makes the probe over the current h_1 edges information efficient but may cause the subsequent probes to be inefficient. Equation (8) says that it is globally optimal to balance the two halves while making sure one of the subpaths has a length of an integer power of 2. The resulting probing scheme is called the “ 2^m -splitting” probing scheme. Note that, in both local and global optimums, we are trying to balance the probabilities of syndrome 0 and syndrome 1, indicating that each efficient probe should provide approximately one bit of network state information.

To illustrate the “ 2^m -splitting” probing scheme, let us consider a linear topology with seven edges. As shown in Fig. 3, we assume that the second edge (BC) fails. The probing algorithm outputs the syndrome 101 for the network state 0100000.

It is also important to observe that, if the problem is changed into the scenario where there is at least one faulty edge in the linear network, and our objective is to locate the first (leftmost) one, the optimal probing scheme is exactly the same as the one described above, since the algorithm never tests a subpath

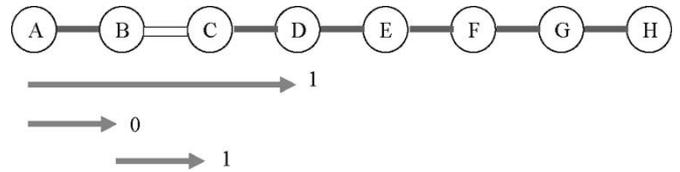


Fig. 3. Optimal probing scheme for the linear network with seven edges and one faulty edge of BC. The syndrome is 101.

without knowing that every edge to the left is fault free. It turns out that this is crucial in developing the run-length probing algorithm (RLPA) in the next section.

IV. RUN-LENGTH PROBING ALGORITHMS

The optimization problem (4) is, in fact, equivalent to designing an optimal binary decision tree for an equivalent decision problem. In general, the design of an optimal binary decision tree has been approached with well-established dynamic-programming algorithms [18], [19]. However, it has been shown that the sequential diagnosis problem is co-NP-complete [20], meaning that the computational complexity of probing algorithms grows exponentially with the network size. From a practical point of view, we are thus more interested in finding simpler algorithms to design near-optimum probing schemes.

In this section, as a tradeoff between complexity and performance, we seek a class of near-optimum probing schemes whose computational complexity is on the polynomial order of network size m . This class of near-optimum probing schemes have probe syndromes consisting of a series of cascaded run-length codes. We show that this probing scheme is asymptotically optimal in that it achieves the minimum average number of probes per edge for large-enough networks. Furthermore, the RLPA is easy to implement and its performance can be obtained in closed form.

A. Probing Schemes Based on Run-Length Codes

In Section III-C, we have shown the optimal probing scheme for the scenario where only one failure exists over a linear topology. We have also shown that the same probing scheme is optimal for the scenario where multiple failures exist, and the objective is to locate the first (leftmost) failure on the linear topology.

At the same time, given any network state $s = f_1 f_2 \cdots f_m$, it must have the format of $0^{i_1} 1 0^{i_2} 1 \cdots 0^{i_L}$ where i_1, i_2, \dots, i_L are nonnegative integers and 0^i means a run of i “0,” and each of the segments $0^{i_j} 1$ is called a substate. In order to encode the network state, it is natural to encode each of such substates separately, since any probe can locate at most one faulty edge at a time. This idea suggests that we can, instead of coding for binary input streams, code on the symbols in the set $Z_0 = \{0^i 1\}_{i=0}^{\infty}$, corresponding to always finding out the first (leftmost) faulty edge along the Euler trail.

The problem of such an approach is that the input alphabet has a size of infinity, which makes it difficult to encode. To this end, it is natural to have an upper limit on the number of

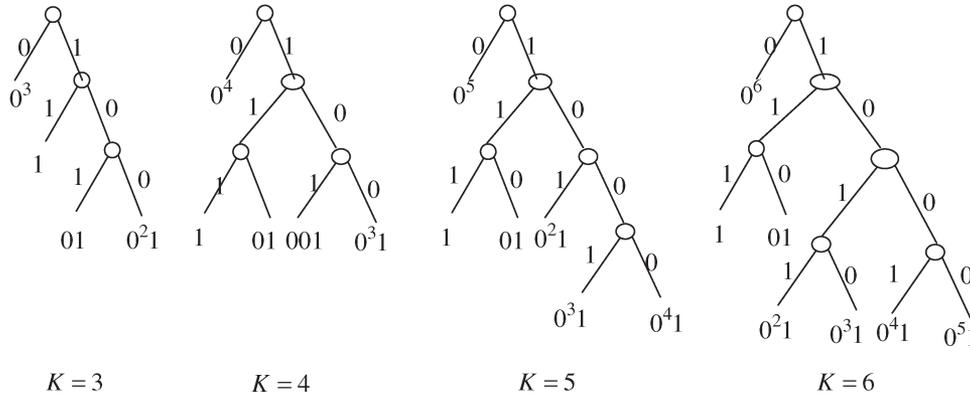


Fig. 4. Run-length probing decision trees for $K = 3, 4, 5, 6$, which are also the optimal Huffman code trees for the corresponding intermediate symbol sets Z .

edges being considered at a time, that is, to code over the finite alphabet set

$$Z = \{0^K, 0^i 1\}_{i=0}^{K-1}. \tag{10}$$

This corresponds to finding the first faulty edge, if any, in K consecutive edges. The prior probability distribution of the possible symbols in Z is given by

$$\Pr(0^i 1) = q^i p, \quad 0 \leq i \leq K - 1 \tag{11}$$

and

$$\Pr(0^K) = q^K. \tag{12}$$

Given this setup, a simple probing scheme follows. First, a probe is sent over a set of K consecutive edges along the Euler trail. If all the edges are fault free, we move onto the next set of K consecutive edges along the Euler trail. If, on the other hand, the first probe says that there is at least one faulty edge in this group, we can employ the “ 2^m -splitting” approach, described in Section III-C, to locate the first faulty edge. The process resumes with another group of K edges along the trail right after this faulty edge.

The parameter K is called the maximum probing length. Intuitively, for small values of p , or in other words if the edges are more reliable, one should probe more edges at a time, i.e., choose a larger value of K . The optimal value of K can be chosen such that the first probe of K consecutive edges returns ON or OFF with approximately equal probabilities so that this probe is more information efficient. This can be achieved by choosing K as the unique positive integer satisfying the inequality

$$q^K + q^{K+1} \leq 1 < q^K + q^{K-1}. \tag{13}$$

Solving this inequality, we obtain the maximum probing length as $K = \lceil -\log_q(1 + q) \rceil$.

It turns out that the probing algorithm developed here is equivalent to the run-length coding procedure for the source-coding problem. For a detailed description of the run-length code, please refer to [12] and [21]. For the rest of this paper, we refer to the algorithm above as the run-length probing scheme. A detailed description of this algorithm is given below.

Let P_i^j be a path (i.e., a permissible probe) that covers edge i to edge j over the Euler trail. When that path is being probed,

it is active. Let h_l denote the number of edges in the current active path, and h_r denote the number of edges in the subsequent active path that is to be probed if all the edges in the current active path are fault free or the current active path has only one edge and it fails, and let i be the start point of the active path. The run-length probing scheme is given by the following algorithm:

Run-Length Probing Algorithm (RLPA):

- Step 0:
 - Set $i = 0$.
- Step 1:
 - Set $h_l = h_r = K$.
- Step 2:
 - Probe the path $P_{i+1}^{i+h_l}$;
 - If the syndrome $r(P_{i+1}^{i+h_l}) = 0$
 - Set $i = i + h_l$, $h_l = h_r$ and $h_r = K$,
 - Go to Step 2;
 - Else if the syndrome $r(P_{i+1}^{i+h_l}) = 1$
 - Set $h = h_l$,
 - Set $h_l = g(h)$ [function $g(\cdot)$ is given by (8)],
 - Set $h_r = g(h - h_l)$,
 - If $h_l \geq 2$
 - Go to Step 2,
 - Else if $h_l = 1$
 - The edge $P_{i+1}^{i+h_l}$ fails,
 - Set $h_l = h_r = K$, $i = i + 1$,
 - Go to Step 2.

To understand the run-length probing scheme pictorially, we illustrate the corresponding probing decision trees for different K 's in Fig. 4. Note that, these trees are also the optimal Huffman code trees for the finite symbol set (10) for different K 's. It turns out that for the particular set $Z = \{0^K, 0^i 1\}_{i=0}^{K-1}$, the Huffman code can in fact be implemented under the consecutive probing constraint. This should not be very surprising since we have already known that 1) the Huffman code is always optimal for any given alphabet and 2) the algorithm above is optimal in locating the first faulty edge on a lightpath. The only missing logical step is that the symbol 0^K is always assigned to a length-1 codeword, corresponding to a single probe. It can be shown that this is optimal from a coding

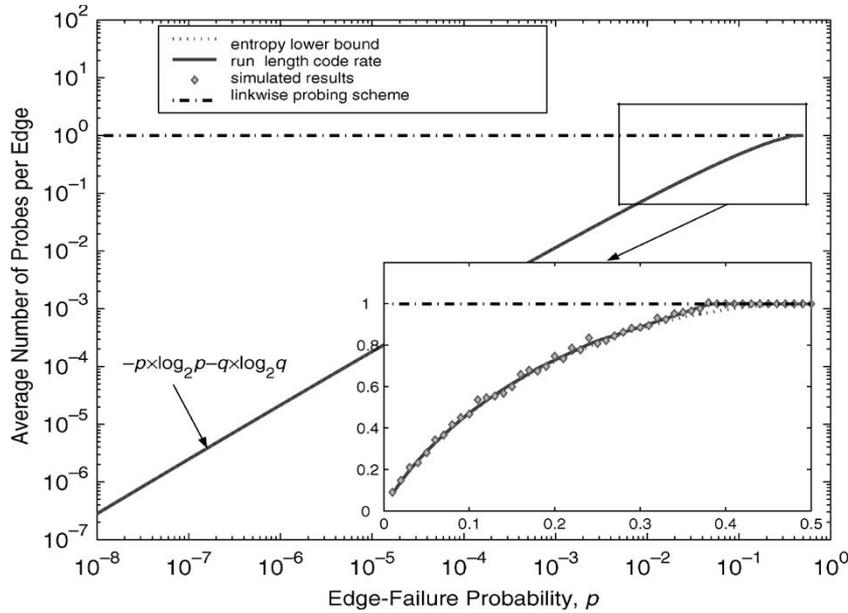


Fig. 5. Simulated average number of probes per edge for an Eulerian topology with 50 edges is compared with the run-length code rate and the entropy lower bound.

perspective since K is chosen such that the symbol 0^K is much more likely than the other symbols in Z . In fact, one can easily verify that the probability of 0^K is larger than $2/5$, which assures the optimality of assigning to it a length-1 codeword [16].

In summary, the RLPA is natural for the fault-diagnosis problem due to the following two reasons.

- 1) Each probe can at most locate one faulty edge; thus, it makes sense to split the network state into substates and locate the faulty edges one by one.
- 2) The probing algorithm can achieve the information-theoretical optimum in locating the individual faulty edges.

Note that the RLPA is restricted on an Euler trail of the network and ignores other connections. In general, this restriction may seriously reduce the set of admissible probes; thus, one cannot claim a general optimality of this algorithm over all possible sequential probing schemes. However, this probing scheme is optimal within the class of “nested” probing schemes where each successive probe includes only a subset of edges from the previous probe or a set of edges that are not tested in the previous probe. Wolf [22] has derived similar results under a totally different context of using group-testing approach to resolve the conflict in multiaccess communications and showed that a similar scheme is optimal within the class of nested group-testing algorithms [23].

One advantage of the run-length probing scheme is that it outperforms the greedy algorithm, although it has the same minimum computational complexity as the greedy probing scheme [19], which will be addressed in Section V. Another distinguished advantage of the run-length probing scheme is that, with information-theoretic insights, one can compare the performance of the run-length algorithm against the entropy lower bound of the global optimum as given in (7). To do that, we will start in the next section by deriving the average number of probes per edge of the run-length probing scheme, again

by taking advantage of the known results on the run-length source codes.

B. Average Number of Probes Required for Run-Length Probing Schemes

The following lemma characterizes the average number of probes per edge required for run-length probing schemes.

Lemma 1: The average number of probes per edge required by the run-length probing scheme to fully identify the network state $\bar{\mathcal{L}}_{\text{run-length}}$ satisfies

$$\bar{\mathcal{L}}_{\infty} \leq \bar{\mathcal{L}}_{\text{run-length}} \leq \bar{\mathcal{L}}_{\infty} + \frac{1}{m} \tag{14}$$

where m is the total number of edges in the network, and

$$\bar{\mathcal{L}}_{\infty} \triangleq p \cdot \left(\lceil \log_2 K \rceil + 1 + \frac{q^k}{1 - q^K} \right) \tag{15}$$

with $K = \lceil -\log_q(1 + q) \rceil$ and $k = 2^{\lceil \log_2 K \rceil + 1} - K$.

This lemma can be proved using the results from the run-length code [21].

We observe that for a reasonably large network, $\bar{\mathcal{L}}_{\infty}$ is independent of the network size (n and m) and is a good approximation of the actual performance. To verify this, the values of $\bar{\mathcal{L}}_{\infty}$, as a function of the edge-failure probability $p \in (0, 1/2)$, are plotted in Fig. 5, together with the simulated results of the actual number of probes required $\bar{\mathcal{L}}_{\text{run-length}}$, the performance of the edgewise probing scheme, and the entropy lower bound given in (7), for Eulerian networks with 50 edges. We have two observations from this plot. First, the plot indicates that $\bar{\mathcal{L}}_{\infty}$ is a good approximation of the actual performance of the run-length probing scheme over a broad range of reliability regime.

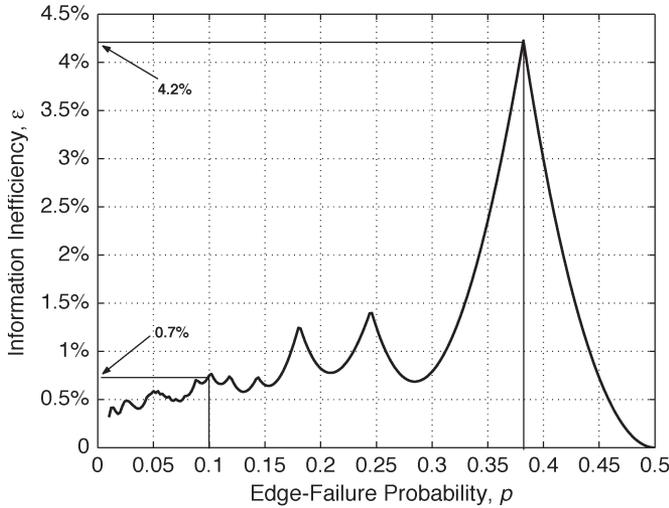


Fig. 6. Information inefficiency of run-length probing schemes for different edge-failure probabilities.

This suggests that for a large Eulerian network ($m \gg K$), we can approximate the average number of probes for the run-length probing scheme as $\mathcal{L}_{\text{run-length}} \approx m \cdot \bar{\mathcal{L}}_{\infty}$.

Second, when the edge-failure probability is small (of greater engineering interests), the average number of probes required is close to the entropy lower bound. For example, for an Eulerian network with 1000 edges and edge-failure probability $p = 0.01$, the run-length probing scheme requires only 81.05 probes on the average. Compared to the entropy lower bound of 80.79 probes, it requires only an additional number of 0.26 probes. This suggests that even though most of the information of the network topology is ignored, the run-length probing scheme still achieves a near-optimal performance.

To gain a clearer view of this performance, we define the information inefficiency of the probing scheme as the ratio between the extra number of probes per edge (compared to the entropy lower bound) and the entropy of each edge, i.e.,

$$\varepsilon(p) = \frac{\bar{\mathcal{L}}_{\pi} - H_b(p)}{H_b(p)} \tag{16}$$

where $\bar{\mathcal{L}}^{\pi}$ is the average number of probes per edge of any probing decision tree $\pi \in \Pi(G)$. In Fig. 6, we plot the information inefficiency of the run-length probing scheme, again with the actual value of $\bar{\mathcal{L}}_{\text{run-length}}$ approximated by $\bar{\mathcal{L}}_{\infty}$. As shown in Fig. 6, when the edge-failure probability p decreases, the run-length algorithm becomes more efficient, with the fluctuation due to the change of the choice of the maximum probing length K , which takes on only integer values. In particular, if the edge-failure probability is less than 0.1, the average number of probes per edge of the run-length probing scheme, for large networks, is upper bounded by

$$\bar{\mathcal{L}}_{\text{run-length}} \leq 1.007H_b(p) \tag{17}$$

which is only 0.7% higher than the entropy lower bound. Moreover, the difference between the achieved performance

and the entropy lower bound is uniformly bounded. In the range of $p \in (0, 0.5]$, the worst case occurs at $p = (3 - \sqrt{5})/2$, where

$$\bar{\mathcal{L}}_{\text{run-length}} \approx 1.0423H_b(p) \tag{18}$$

meaning that the actual performance of the run-length probing scheme is less than 5% larger than the lower bound.

Combing (7), (17), and (18), we conclude that the performance of the run-length probing scheme is bounded by the following inequalities:

$$H_b(p) \leq \bar{\mathcal{L}}_{\text{run-length}} \leq (1 + \varepsilon(p)) H_b(p) \tag{19}$$

where $\varepsilon(p)$ tends to decrease with smaller edge-failure probability, and we can approximate $\varepsilon(p) < 0.01$ for $p \leq 0.1$ and $\varepsilon(p, q) < 0.05$ for $0.1 < p \leq 0.5$.

Finally, the performance of the run-length probing scheme can be used as an upper bound for the minimum average number of probes per edge. In practical networks whose failure probability is usually small, both the upper and lower bounds are reduced to the entropy of the individual edge as indicated in Fig. 5, and thus are both tight. The convergence of both upper and lower bounds to the entropy and the scaling function (19) indicate that the minimum probing effort approximately equals the entropy of the network states. This actually follows from the fact that, in an efficient probing scheme (e.g., the run-length probing scheme), each probe is designed to provide approximately 1 bit of state information. Since the amount of unknown information in the network state is equal to the entropy, the number of probes required to identify the network state is, on the average, equal to the entropy.

C. Complexity of Run-Length Probing Schemes

In this section, we address the computational advantages of run-length probing schemes, including the constant storage-space requirement and the polynomial running time.

First, the run-length probing scheme is an online algorithm that chooses the successive probes dynamically and does not store the whole probing decision tree beforehand. In fact, the run-length probing scheme only stores three variables (i.e., i , h_l , and h_r) and reduces the requirement for the storage space to a constant of 3 [i.e., $O(1)$]. On the contrary, for a network with m edges, any offline algorithm, which the run-length probing scheme does not belong to, stores the whole probing decision tree, and thus, the requirement for the storage space is $2 \cdot 2^m - 1$ [i.e., $O(2^m)$], which grows exponentially with the network size.

Second, the complexity of the run-length probing scheme is proportional to the number of executions of Step 2 in the algorithm, since each execution of Step 2 in the RLPA costs roughly a constant time. In the worst case, all the edges fail simultaneously, and the run-length probing scheme employs $\log K$ probes to identify each edge failure. This corresponds to executing Step 2 $\log K$ times for each edge failure. For

a network with m edges, the total number of executions of Step 2 is at most $m \cdot \log K$, which suggests that the complexity of the RLPA is $O(m \cdot \log K)$. It follows that the run-length probing scheme is a polynomial-time algorithm.

V. GREEDY PROBING ALGORITHMS

The information-theoretic perspective in the context of fault diagnosis is in general also very useful in understanding, comparing, and improving the network probing schemes based on existing heuristics. In particular, we will, in this section, study the fault-diagnosis approaches based on dynamic programming.

As previously mentioned, dynamic-programming solutions to the network-diagnosis problem (equivalently, the optimal binary decision-tree design) are in general Co-NP Complete. As a compromise, various suboptimal greedy algorithms [25] are proposed based on local optimization heuristics. The performance of such heuristic algorithms is usually studied only via simulations. With our information-theoretic viewpoint of the problem, it is natural to connect these problems to their counterparts of source-coding problems with dynamic-programming approaches, which have been thoroughly studied for decades.

In this section, we will first review the dynamic-programming formulation of the network-diagnosis problem, with the focus of a particular greedy algorithm that maximizes the local information gain at each stage [25]. We will then compare the performance of this scheme with that of the run-length algorithm to gain more insights.

A. Dynamic-Programming and Greedy Algorithms

We first introduce some useful notations. The design of optimal fault-diagnosis algorithms is equivalent to the design of optimal binary decision trees. We denote \mathcal{I}_π as the set of inner nodes of the decision tree π . Let $\varsigma \in \mathcal{I}_\pi$ denote one of the inner nodes, and P_ς the probability that ς is reached. It follows that P_ς equals the sum of the prior probabilities of the network states that are descendants of the node ς [26]. Let t_ς be the probe employed at inner node ς , $\Pr(0|\varsigma)$ and $\Pr(1|\varsigma)$ be the probabilities that this test returns 0 and 1, corresponding to the probabilities that the network state lies in the left or right subtrees of inner node ς , respectively. Furthermore, let \mathcal{L}_ς be the average number of successive probes required when the inner node ς is reached.

Now, to design the diagnosis algorithm with the minimum average number of probes, we need to, at each inner node ς , choose a probe t_ς to minimize

$$\mathcal{L}_\varsigma = 1 + \Pr(0|\varsigma) \times \mathcal{L}_{\varsigma,0}^* + \Pr(1|\varsigma) \times \mathcal{L}_{\varsigma,1}^* \quad (20)$$

where $\mathcal{L}_{\varsigma,0}^*$ and $\mathcal{L}_{\varsigma,1}^*$ are the minimum average number of additional probes required by the left and right subtrees from the inner node ς , respectively. In particular, taking ς as the root of the entire tree, the solution of this optimization problem gives the optimal fault-diagnosis scheme.

Note that the difficulty of such a problem comes from the fact that the optimization problems at different steps are coupled. In choosing t_ς , one needs to cater for the future optimizations of $\mathcal{L}_{\varsigma,0}^*$ and $\mathcal{L}_{\varsigma,1}^*$. As a result, the computational complexity of this problem grows exponentially with the number of edges m . Some results of using dynamic programming in designing binary decision trees can be found in [18] and [19].

Now, from an information-theoretic perspective, the performance, in terms of the average number of probes, can be computed from the local information efficiencies as follows. For a given probing tree π , the average number of probes required to reach the leaves can be computed as

$$\begin{aligned} \mathcal{L}_\pi &= \sum_{s \in \mathcal{S}} \Pr(s) \times (\text{number of probes to reach state } s) \\ &= \sum_{s \in \mathcal{S}} \Pr(s) \times (\text{number of ancestors of } s) \\ &= \sum_{\varsigma \in \mathcal{I}_\pi} \Pr(\varsigma) \times 1. \end{aligned} \quad (21)$$

On the other hand, one can write $H(\varsigma)$ as the amount of information in bits, obtained by employing the probe t_ς as node ς is reached, that is

$$H(\varsigma) \triangleq H_b(\Pr(0|\varsigma), \Pr(1|\varsigma)). \quad (22)$$

By running this fault-diagnosis algorithm, one can always find out the network state, which contains, on the average, $m \cdot H_b(p)$ bits of information and can be viewed as the sum of the information obtained in each step, i.e.,

$$m \cdot H_b(p) = \sum_{\varsigma \in \mathcal{I}_\pi} \Pr(\varsigma) \times H(\varsigma). \quad (23)$$

Hence, the total inefficiency of the algorithm, in terms of the average number of probes required in excess of the information minimum $m \cdot H_b(p)$, is

$$\mathcal{L}_\pi - m \cdot H_b(p) = \sum_{\varsigma \in \mathcal{I}_\pi} \Pr(\varsigma) \times (1 - H(\varsigma)) \quad (24)$$

where $(1 - H(\varsigma))$ is referred as the local inefficiency of the algorithm π at the inner node ς . Intuitively, one probe is used to return only $H(\varsigma)$ bits of information. Hence, the difference between the two measures is the information inefficiency of employing this probe and the weighted sum of the inefficiency over the tree gives the total number of extra probes required by the given probing scheme.

Following such discussions, we observe that in order to design efficient network-diagnosis algorithms, it is desirable to minimize the local inefficiency at each stage. Intuitively, by always asking the question to which the answer is completely

without bias, one would expect to figure out the network state with fewer questions. This corresponds to making the left and right subtrees as balanced as possible, i.e., to choose a probe t_ζ to minimize

$$t_\zeta^* = \arg \min_{t_\zeta} |\Pr(0|\zeta) - \Pr(1|\zeta)|. \quad (25)$$

Such intuition of balancing the probabilities of the outcomes of a probe is, in general, very useful. We have just used that to design the maximum probing length K for the run-length algorithm. For the first probe over K links, the probabilities of UP and DOWN are, respectively, q^K and $1 - q^K$. It can be shown that the choice of K in inequality (13) indeed minimizes the difference between these two probabilities.

It is important to note that such an approach, by maximizing the local information gain $H(\zeta)$, may not necessarily be the globally optimum choice. As an example, in the example of locating a single failure in Section III-C, it is globally optimal to split the path as in (8), to make sure that the length of one of the subpaths is an integer power of 2. On the other hand, a greedy design based on local optimizations would simply split the path into equal halves.

The greedy algorithm presented above, to maximize the local information gain, is in fact one of many variations [25]. Such algorithms are sometimes preferred due to their conceptual simplicity. However, the run-length algorithm has the same order of computational complexity as these algorithms. In the next section, we will investigate the performance of the greedy algorithm based on maximizing local information gains.

B. Performance Comparison

Using Monte Carlo simulation, we simulate the performance of the greedy probing algorithm and compare it with that of the RLPA in this section.

To compare these two algorithms in a finer scale, we plot the probing inefficiency of the greedy probing scheme over the run-length probing scheme, i.e.,

$$\eta(p) = \frac{\bar{\mathcal{L}}_{\text{greedy}} - \bar{\mathcal{L}}_{\text{run-length}}}{\bar{\mathcal{L}}_{\text{run-length}}} \quad (26)$$

as a function of the edge-failure probability p in Fig. 7. In (26), $\bar{\mathcal{L}}_{\text{run-length}}$ and $\bar{\mathcal{L}}_{\text{greedy}}$ are the average number of probes per edge for the RLPA and the greedy probing algorithm, respectively.

We observe that for some range of link-failure probability, both the run-length probing scheme and the greedy probing scheme have the same average number of probes per edge. It can be verified that this happens for all the edge-failure probabilities such that the maximum probing length K is an integer power of 2. Under this scenario, splitting a path of length K automatically gives subpaths with lengths as powers of 2; hence, the local and global optimums coincide. On the other hand, when the edge-failure probability is in the range such that K is not an integer power of 2, the greedy algorithms are strictly suboptimal. As a result, in Fig. 7, there is a periodic

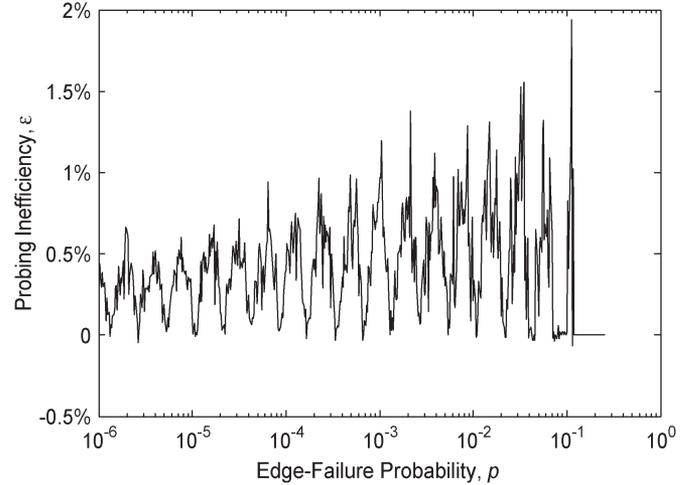


Fig. 7. Performance comparison between the RLPA and the greedy probing algorithm.

pattern in the log plot: When p is such that $K(p)$ equals a power of 2, the probing inefficiency is equal to 0; as p increases or decreases such that $K(p)$ does not equal a power of 2, the probing inefficiency is strictly nonzero.

Therefore, although both probing schemes have the same computational complexity, the run-length probing scheme provides some cost-saving over the greedy probing scheme. However, the difference between the two algorithms is quite limited. Intuitively, this is because the global optimum solution always makes sure that all but one subpath have lengths of powers of 2, in which case the greedy algorithm is also optimum.

VI. APPLICATIONS OF RUN-LENGTH PROBING SCHEMES FOR ALL-OPTICAL WDM NETWORKS

In this section, we first address some practical issues in the implementation of the run-length probing scheme and then apply the run-length probing scheme to different WDM network scenarios.

To employ the run-length probing scheme, we assume, for simplicity, that a centralized fault-management agent communicates with all the network nodes through a reliable out-of-band control channel. In real life, this limit is approached even with in-band signaling for fairly reliable networks and when flood routing for signaling is used, as long as the network remains connected. First, the fault-management agent processes all the previous probe syndromes to decide the attributes of the successive probe, including the source node, the destination node, and the routing path of the probe. Then, the agent signals the source node to send the probe along the specified lightpath to the destination node. After the destination node determines the probe syndrome, it reports the syndrome to the agent through the reliable out-of-band control channel. This process continues until all the edge states have been identified and the entire process is then repeated in a predetermined or adaptive fashion.

In developing the run-length probing scheme, we have made another assumption that the network is abstracted as an

undirected graph, meaning that each optical link is bidirectional. However, in practical all-optical networks, each optical link is seldom bidirectional due to the prohibitively high cost of incorporating two-way traffic in a single fiber. Instead, each bidirectional logical edge includes two parallel unidirectional optical fibers with signals in only one direction. To apply the run-length probing scheme to such networks, we replace each edge in the undirected graph with two parallel directed edges. It follows that the undirected network is converted into a directed graph. Under such a conversion, it can be shown that an Euler trail always exists for a directed graph converted from any undirected graph. To diagnose failures, we simply employ the run-length probing scheme over a directed Euler trail in the directed graph. When the maximal probing length is much less than the number of edges, we can approximate the average number of probes per edge required for a practical network with unidirectional edges with the information entropy of individual directed edge.

Without loss of generality, we assume that there are W wavelengths per fiber in an all-optical WDM network, which is subjected to different kinds of failures. According to the scale of their effect, these failures can be classified into two categories. One category is wavelength-level failures, which affect a particular wavelength channel, e.g., transmitter/receiver failures with one dedicated transmitter/receiver per wavelength and single-bandwidth optical filter or single-channel frequency-selective switch failures. The other category is fiber-level failures, which affect all the wavelength channels within an individual fiber, such as fiber cuts, EDFA breakdowns, and transmitter/receiver failures in the case of only one tunable transmitter/receiver per fiber (which rarely happens).

We observe that, although the wavelength-level and the fiber-level failures are statistically independent, all the wavelength channels passing through an EDFA fail simultaneously when the EDFA fails. This suggests that failures in different wavelength channels on the same fiber are dependent in that knowing one particular wavelength channel fails reveals some information about the failures of other wavelength channels. Therefore, the fault-diagnosis algorithm for practical all-optical WDM networks must consider interdependence among failures in different wavelength channels. In particular, the application of the run-length probing scheme over practical all-optical networks depends on the relative dominance between wavelength-level and fiber-level failures. In other words, the relationship between p_F (i.e., the prior probability of individual fiber-level failure) and p_W (i.e., the prior probability of individual wavelength-level failure) determines how the run-length probing scheme should be implemented over practical all-optical WDM networks.

In one extreme, for an all-optical network where the wavelength-level failures dominate the fiber-level failures (i.e., $p_W \gg p_F$), we can view the network as a graph where each physical link is represented with W parallel edges, and employ the run-length probing scheme over an Euler trail of the resulting graph. We call this the wavelength-level implementation. For a reasonable large number with $m \cdot W \gg K$, the average number of probes required by the run-length probing scheme can be approximated by $W \cdot m \cdot H_b(p_W)$.

In the other extreme, for an all-optical WDM network where the fiber-level failures dominate the wavelength-level failures (i.e., $p_F \gg p_W$), we can view the network as a graph where each physical link is represented with one edge, and employ the run-length probing scheme over an Euler trail of the resulting graph. We call this the fiber-level implementation. For a large network of $m \gg K$, the average number of probes required is approximately equal to $m \cdot H_b(p_F)$.

Finally, for an all-optical WDM network subjected to a comparable (in terms of probability of occurrence) mixture of both fiber-level and wavelength-level failures, we can still use the fault-diagnosis/source-coding equivalence to obtain a useful lower bound for the minimum average number of probes required as the information entropy of network states, i.e.,

$$\mathcal{L}^* \geq m \cdot H(F_1, F_2, \dots, F_W) \quad (27)$$

where F_i 's are dependent random variables indicating states of wavelength channel i 's, and $H(\cdot)$ is the information entropy function. The entropy function $H(F_1, F_2, \dots, F_W)$ can be calculated through the summation of a sequence of conditional entropy functions, i.e.,

$$H(F_1, F_2, \dots, F_W) = \sum_{i=1}^W H(F_i | F_1, \dots, F_{i-1}). \quad (28)$$

However, we are not yet clear whether the entropy lower bound (27) can be achieved, or if achievable, how we can develop probing schemes to achieve this lower bound. Fortunately, in real life, either wavelength or fiber integrity, but not both at the same time, is the dominant reliability factor.

Actually, we can view the network as a graph where each physical link is represented with W parallel edges and employ the run-length probing scheme over an Euler trail of the resulting graph. The performance of this wavelength-level implementation, which is hard to obtain due to the complicated dependence among failures in different wavelength channels, can certainly serve as an upper bound for the minimum average number of probes required by an optimal probing scheme. However, since the wavelength-level implementation does not consider the dependencies among failures in different wavelength channels, the run-length probing scheme is not optimum in general. The same conclusion that the information entropy is a lower bound and the run-length probing scheme might not be near optimum can be extended to a more general failure model that accommodates dependent failures and/or heterogeneous failures.

VII. CONCLUSION

In this paper, we investigated the fault-diagnosis problem for all-optical wavelength-division-multiplexing (WDM) networks under a probabilistic edge-failure model. Our investigation reveals that the complexity of the fault-management system of all-optical networks can be reduced and thus, the operational network cost can be kept low. This research can further the understanding of the relationship between the amount of

network information gathered and the performance of network management.

Using the average number of probes required as the cost metric, we characterized the optimal fault-diagnosis algorithms from an information-theoretical perspective. In particular, the mathematical equivalence between the fault-diagnosis problem and the source-coding problems suggests that the minimum average number of probes is lower bounded by the entropy of the network states. We also showed that the “ 2^m -splitting” probing scheme is optimum for the special case of single failure over a linear topology. Based on these heuristics, we developed a class of efficient network-diagnosis algorithms, i.e., run-length probing schemes. Its performance is uniformly bounded above by $(1 + \varepsilon)H_b(p)$ and converges to the entropy lower bound as the edge-failure probability decreases to 0. We also compared the run-length probing scheme to the “greedy” probing schemes, which indicates that the run-length probing scheme is the algorithm of choice since it outperforms the greedy algorithm of the same computational complexity. Several practical issues in implementing the run-length probing scheme in future all-optical WDM networks were also addressed. The guideline for efficient probing schemes is that each probe should provide approximately 1 bit of state information, and thus, the number of probes required is approximately equal to the information entropy of the network states.

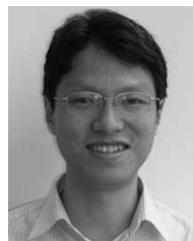
Although this research is done for all-optical networks, full understanding of the network-diagnosis problem for all-optical networks can lead to deeper insight into network management under more general models (e.g., transient failures). For future work, the network-diagnosis problem with simultaneous edge and node failures needs to be addressed. Other possible future work includes extending the centralized run-length probing scheme into a distributed implementation. Finally, using appropriate cost metrics, these diagnostic approaches can probably be applied to the diagnosis of network failures of other networks including wireless and sensor networks.

ACKNOWLEDGMENT

The authors wish to thank Prof. M. Medard, E. Swanson, Dr. P. Saengudomlert, Dr. C. Liu, L. Chen, J. Sun, and G. E. Weichenberg for helpful discussions.

REFERENCES

- [1] V. W. S. Chan *et al.*, “A precompetitive consortium on wide-band all-optical networks,” *J. Lightw. Technol.*, vol. 11, no. 5/6, pp. 714–735, May/June 1993.
- [2] V. W. S. Chan, “All-optical networks,” *Sci. Amer.*, vol. 273, no. 3, pp. 56–59, Sep. 1995.
- [3] Y. G. Wen and V. W. S. Chan, “Ultra-reliable communication over vulnerable all-optical networks via lightpath diversity,” *J. Sel. Areas Commun.*, vol. 23, no. 8, pp. 1572–1587, 2005.
- [4] F. P. Preparata, G. Metze, and R. T. Chien, “On the connection assignment problem of diagnosable systems,” *IEEE Trans. Electron. Comput.*, vol. 16, no. 6, pp. 848–854, Dec. 1967.
- [5] A. K. Somani, V. K. Agarwal, and D. Avis, “A generalized theory for system-level diagnosis,” *IEEE Trans. Comput.*, vol. c-36, no. 5, pp. 538–546, May 1987.
- [6] M. Barborak, M. Malek, and A. Dahbura, “The consensus problem in fault-tolerant computing,” *ACM Comput. Surv.*, vol. 25, no. 2, pp. 171–220, Jun. 1993.
- [7] S. Lee and K. G. Shun, “Probabilistic diagnosis of multiprocessor systems,” *ACM Comput. Surv.*, vol. 26, no. 1, pp. 121–139, Mar. 1994.
- [8] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*, 2nd ed. San Mateo, CA: Morgan Kaufmann, 2002.
- [9] D. C. Kilper, R. Bach, D. J. Blumenthal, D. Einstein, T. Landolsi, L. Ostar, M. Preiss, and A. E. Willner, “Optical performance monitoring,” *J. Lightw. Technol.*, vol. 22, no. 1, pp. 294–304, Jan. 2004.
- [10] W. G. Yang, “Sensitivity issues of optical performance monitoring,” *IEEE Photon. Technol. Lett.*, vol. 14, no. 1, pp. 107–109, Jan. 2002.
- [11] G. Rossi, T. E. Dimmick, and D. J. Blumenthal, “Optical performance monitoring in reconfigurable WDM optical networks using subcarrier multiplexing,” *J. Lightw. Technol.*, vol. 18, no. 12, pp. 1639–1648, Dec. 2000.
- [12] R. Gallager and D. Van Voorhis, “Optimal source codes for geometrically distributed integer alphabets,” *IEEE Trans. Inf. Theory*, vol. IT-21, no. 2, pp. 228–230, Mar. 1975.
- [13] B. Bollobas, *Modern Graph Theory*. New York: Springer-Verlag, 1998.
- [14] T. Ho, M. Medard, and R. Koetter, “An information-theoretic view of network management,” in *Information Communication (INFOCOM)*, San Francisco, CA, May 2003, pp. 1456–1466.
- [15] R. Ahlswede and I. Wegener, *Search Problems*. New York: Wiley, 1987.
- [16] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [17] E. Wong, “A linear search problem,” *SIAM Rev.*, vol. 6, no. 2, pp. 168–174, Apr. 1964.
- [18] H. J. Payne and W. S. Meisel, “An algorithm for constructing optimal binary decision trees,” *IEEE Trans. Comput.*, vol. C-26, no. 9, pp. 905–916, Sep. 1977.
- [19] K. R. Pattipati and M. G. Alexandridis, “Application of heuristic search and information theory to sequential fault diagnosis,” *IEEE Trans. Syst., Man, Cybern.*, vol. 20, no. 4, pp. 872–886, Jul./Aug. 1990.
- [20] V. Raghavan and A. Tripathi, “Sequential diagnosability is Co-NP complete,” *IEEE Trans. Comput.*, vol. 40, no. 5, pp. 584–595, May 1991.
- [21] H. Tanaka and A. Leon-Garcia, “Efficient run-length encodings,” *IEEE Trans. Inf. Theory*, vol. IT-28, no. 6, pp. 880–890, Nov. 1982.
- [22] J. K. Wolf, “Born again group testing: Multi-access communications,” *IEEE Trans. Inf. Theory*, vol. IT-31, no. 2, pp. 185–191, Mar. 1985.
- [23] P. Ungar, “The cut-off point for group testing,” *Commun. Pure Appl. Math.*, vol. 13, pp. 49–54, 1960.
- [24] G. Weichenberg, “High-reliability architectures for networks under stress,” M.Sc. thesis, Dept. Elect. Eng. Comput. Sci., Massachusetts Inst. Technol., Cambridge, MA, 2003.
- [25] S. R. Safavian and D. Landgerebe, “A survey of decision tree classifier methodology,” *IEEE Trans. Syst., Man, Cybern.*, vol. 21, no. 3, pp. 660–674, May/June 1991.
- [26] R. W. Yeung, *A First Course in Information Theory*. New York: Kluwer/Plenum, 2002.



Yonggang Wen (S'00) was born in Nanchang, Jiangxi, China, in 1977. He received the B.Eng. degree from the University of Science and Technology of China (USTC), Hefei, Anhui, China, in 1999 and the M.Phil. degree from the Chinese University of Hong Kong (CUHK), Hong Kong, in 2001. Since 2001, he has been working toward the Ph.D. degree at the Department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology (MIT), Cambridge.

He worked from 1997 to 1999 as a Research Assistant in the Personal Communication Network and Spread Spectrum Laboratory at USTC, focusing on the G3 mobile communication. From 1999 to 2001, he was a Research Assistant in the Lightwave Communication Laboratory at CUHK, working on optical code-division multiplexing systems and high-rate optical routing networks. He is currently a Research Assistant at the Laboratory for Information and Decision Systems (LIDS) of MIT, where his research interests focus on mesh-based network architecture including ultra-reliable all-optical network architecture, efficient network-management algorithms and packet-scheduling over heterogeneous networks.



Vincent W. S. Chan (S'69–M'88–SM'92–F'94) received the B.S., M.S., E.E., and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1971, 1972, and 1974, respectively, all in the area of communications.

From 1974 to 1977, he was an Assistant Professor with the School of Electrical Engineering at Cornell University. He joined Lincoln Laboratory in 1977 as a Staff Member of the Satellite Communication System Engineering Group working on military communications and networking. In January

1981, he became the Assistant Leader of the Communication Technology Group starting a research and development program on optical space communications. In July 1983, he formed and became Leader of the Optical Communication Technology Group and Manager of the Laser Intersatellite Transmission Experiment (LITE) Program. He became the Head of the Communications and Information Technology Division of Lincoln Laboratory until joining Laboratory for Information and Decision Systems (LIDS) in 1999. He is the Joan and Irwin Jacob Professor of Electrical Engineering and Computer Science and Aeronautics and Director of the LIDS at MIT. He also previously served as the Principal Investigator of a Next Generation Internet Consortium (ONRAMP) formed among AT&T, Cabletron, MIT, and JDS Fitel, and a Satellite Networking Research Consortium formed between MIT, Motorola, Teledesic and Globalstar. He is a Member of the Board of Directors of Vitesse Semiconductor Corporation and the Chairman of its Technical Advisory Board and also serves on the Technical Advisory Boards of Agility Communications, Axiowave Networks (TAB Chairman), and as a Member of the Corporation of Draper Laboratory. His research interests are in optical communications, wireless communications, space communications, and networks.

Dr. Chan is a Fellow of the Optical Society of America (OSA).



Lihong Zheng (S'00–M'02) received the B.S. and M.S. degrees from the Department of Electronic Engineering, Tsinghua University, China, in 1994 and 1997, respectively, and the Ph.D. degree from the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley (UC Berkeley), in 2002.

Since 2002, he has been working as an Assistant Professor in the Department of Electrical Engineering and Computer Sciences and the Laboratory of Information and Decision Systems at Massachusetts Institute of Technology, Cambridge. His research interests include information theory, wireless communications, and wireless networks.

Dr. Zheng received the Eli Jury award from UC Berkeley in 2002 and the IEEE Information Theory Society Paper Award in 2003.