Efficient Fault Detection and Localization for All-Optical Networks

Yonggang Wen, *Student Member, IEEE*, Vincent W.S. Chan, *Fellow, IEEE*
and Lizhong Zheng, *Member, IEEE*
Laboratory for Information and Decision Systems
Massachusetts Institute of Technology
Cambridge, MA 02139
{eewyg, chan, lizhong}@mit.edu

***Abstract*—We investigate the fault diagnosis problem for all-optical networks with probabilistic link and node failures in this paper. Our major contribution is the development of diagnosis algorithms that minimize the operating effort to identify failures. We achieve this by employing the fault diagnosis approach based on *proactive probing with perfect feedback*: knowledge of the network state is progressively refined through a sequence of optical probe signals, each of which is determined upon the results of previous probe signals (i.e. probe syndromes). To detect and localize failures in all-optical networks with probabilistic node and link failures, we introduce a network transformation that maps both link and node failures in an undirected graph into arc failures in a directed graph and apply our previously developed run-length probing scheme [1] to the directed graph. Our analytical and numerical investigation verifies our previously established guideline for efficient fault diagnosis algorithms: *each probe should provide approximately 1-bit of state information, and thus the total number of probes required is approximately equal to the entropy of the network state*. Hence the complexity of optical network fault management functionality is fundamentally related to the information entropy of the network state.**

I. INTRODUCTION

Owing to the recent explosion in internet traffic [2], optical fiber, with its vast transmission bandwidth (~35THz) [3], has emerged as the only realistic transmission medium for backbone networks. Moreover, all-optical networks [4], where data traverses lightpaths without any optical-to-electrical conversion, will be increasingly prevalent in future broadband networks as a result of its expected lower cost and full transparency to different signal formats and protocols. However, as in the case of other networks, all-optical networks are vulnerable to physical failures [5] such as fiber cuts, switch node failures, optical amplifiers and transceivers breakdowns. These failures can lead to costly disruptions in communication, and their detection and localization can constitute a significant fraction of reoccurring network operating costs. To ensure specified levels of quality of service at an affordable cost, an efficient network management system – including efficient fault diagnosis capability - should be in place when all-optical networks are fully deployed in future. In this work, we focus on developing efficient fault diagnosis algorithms, which detect and localize failures in the optical layer, for all-optical networks.

Presently, Synchronous Optical Network (SONET) infers the health of each SONET link by verifying the parity bits embedded in the overhead of data frames [3]. This approach is a manifestation of the fault diagnosis paradigm based on passive monitoring, as illustrated in Fig. 1(a). The monitoring module generates the events - alarms, warnings, parameters of network elements - as inputs to the fault diagnosis engine. Using various algorithms, such as neural networks [6] and Finite-state Machines [7], the fault diagnosis engine identifies a set of network elements whose failures may have caused the input events. Because the monitoring module is decoupled from the fault diagnosis engine, the network manager can follow a "divide-and-conquer" approach in designing different modules separately, and thus reducing design complexity. However, the absence of feedback from the diagnosis engine to the monitoring module could entail tremendous inefficiency in the fault diagnosis process. For example, one single failure could trigger a large number of redundant alarms, all of which are fed into the fault diagnosis engine. Combined with the network growth and faster switching speed, the redundancy in the input events could generate a large amount of management information. It will consume a fair amount of network source to transfer and store this large amount of management information, and thus the operating cost will be increased significantly. To make matters worse, because all these measurements are piggybacked onto real traffic, 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.

Motivated by these shortcomings of the passive monitoring approach to fault diagnosis, we proposed the fault diagnosis scheme in [1], based on proactive probing with perfect feedback. In an all-optical network, optical signals traverse a lightpath - a succession of interconnected links and nodes - without being detected by the intermediate nodes. This property of optical networks permits lightpath probes to test the health of several links/nodes simultaneously. Our approach exploits this fact via a dynamic algorithm which chooses subsequent lightpath probes based upon the results of previous probe signals with the objective of minimizing the total number of probes (i.e., the overhead cost of fault diagnosis). This approach corresponds to the proactive fault diagnosis paradigm shown in Fig 1(b), where feedback from the fault diagnosis engine to the probing module (i.e., the event generator) provides the flexibility to minimize the number of probes required.

Physically, these probing measurements can be implemented via two methods, depending on whether the chosen probing lightpath presently carries traffic or not. If

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.

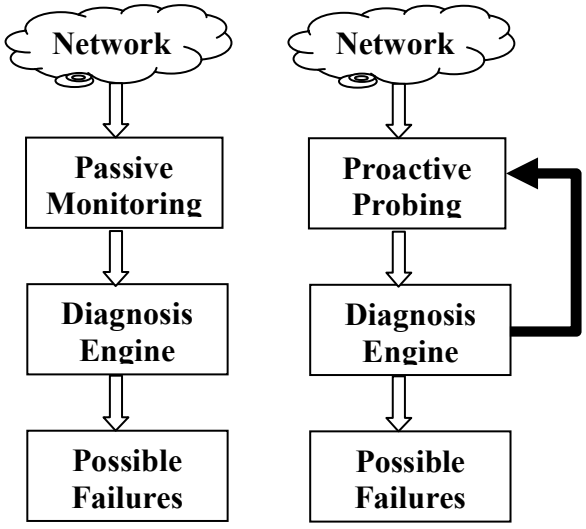

(a) (b)

Fig. 1. Two alternative paradigms for fault diagnosis systems: (a) Fault diagnosis based on passive monitoring of events; (b) Fault diagnosis based on proactive probing of network elements.

traffic exists along the chosen lightpath, the probing signal can be piggybacked onto the existing traffic, similar to the SONET parity check technique. On the other hand, if there is no traffic along the chosen lightpath, the probing signal can be sent out by an unused transmitter at the source node and received by an unused receiver at the destination node. On rare occasions when no transmitter/receiver can be used for probing, an alternative lightpath may be chosen to probe or the length of the original lightpath may be adjusted. For large networks, the impact of a small set of deviations from the optimal probe sequence on the overall efficiency of the fault diagnosis process can be negligible.

Compared with the passive monitoring solution, our proposed adaptive lightpath probing approach has two significant advantages. First, because all of the probes can be implemented without additional hardware provisions, no additional capital expenditure is required. Second, because each successive probe is dynamically chosen according to the results of previous probes (i.e, probe syndromes), we can reduce the operating cost of diagnosing failures to its minimum by solving an optimization problem, as shown in [1].

Our present work complements and extends our previous research in [1] from a more practical perspective. In [1], we have investigated the fault diagnosis problem for all-optical networks with Eulerian topologies [9] under a probabilistic link failure model. In particular, we have established the mathematical equivalence between the fault diagnosis problem and the source-coding problem, which implies an entropy lower bound on the minimum average number of probes required and an information theoretic approach to translating efficient source coding algorithms into efficient fault diagnosis algorithms under the physical probing constraints (e.g., the run-length probing scheme based on the run-length code [10]). Here, we focus on fault diagnosis for all-optical networks with both node and link failures. To diagnose both node and link failures, we introduce a network transformation that converts

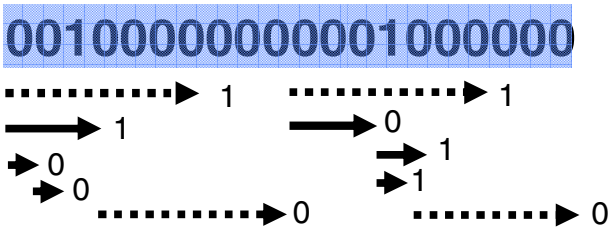

Fig. 2. Demonstration of run-length probing scheme over a network. It contains a sequence of concatenations of two phases: the failure detection phase (dotted lines) and the failure localization phase (solid lines).

the original undirected graph into a directed graph: each link in the original graph is replaced by two parallel directed arcs in opposite directions, and each node of degree d is replaced by a $d \times d$ directed complete bi-partite graph [9], where any node in the left column is connected to any node in the right column via a directed arc. Under this transformation, both link and node failures in the undirected graph are mapped into arc failures in the directed graph. Moreover, the directed topology can always be made Eulerian, rendering the run-length probing scheme applicable. Depending on the relative dominance between link failure probability and node failure probability, different probing strategies are obtained through analytical and numerical investigations.

This paper is organized as follows. In Section II, we review the near-optimum run-length probing scheme and present its performance in closed-form. In Section III, we introduce the network transformation that converts both link and node failures in the undirected graph into arc failures in the directed graph. In Section IV, we apply the run-length probing scheme to all-optical networks with probabilistic link/node failures, and characterize its corresponding performance compared to the entropy lower bound.

II. Run-Length Probing Algorithms For Fault Diagnosis In All-Optical Networks

In [1], we have developed and characterized the run-length probing scheme to diagnosis link failures for Eulerian networks, which contains an Euler trail (i.e., a path containing all the links in the network topology without repetition). The run-length probing scheme has two attractive features. First, the computational complexity of the scheme is on the polynomial order of network size (i.e., the number of edges in the network m). Second, the run-length probing scheme is asymptotically optimal because it achieves the minimum average number of probes per link for large networks.

Essentially, the run-length probing scheme performs a two-phase procedure to identify each individual failure along an Euler trail of the network: the *failure detection* phase and the *failure localization* phase. In the failure detection phase, a detection probe is sent over a set of $K = \left| -\log_{1-p}(2-p) \right|$ consecutive links along the Euler trail, where p is the edge failure probability. If all the links are fault-free, we move onto the next set of K consecutive links along the Euler trail. On the other hand, if the detection probe reveals that there is at least one faulty link in this set of links, the algorithm enters the failure localization phase. In this phase, given that there is

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.

some failure in the detection probe, the "2^m-splitting" binary searching algorithm [1] is employed to locate the closest faulty link to the source. After the fault is localized, the algorithm resumes the failure detection phase by sending another probe spanning K links along the trail immediately following the failure. Fig. 2 illustrates, through a simple example, the two-phase probing scheme for efficient fault diagnosis with $K = 8$, where the dashed lines denote the fault detection phases and the solid lines denote the fault localization phases.

Using the information theoretical interpretation of the run-length probing scheme, we have approximated in [1] the average number of probes per link required by the run-length probing scheme by

$$\bar{L}_\infty(p) \approx \left[1+\varepsilon(p)\right] H_b(p) \quad 0 < p \le \frac{1}{2}, \tag{1}$$

where $H_b(p) = -p\log_2 p - (1-p)\log_2(1-p)$ is the entropy function, and $\varepsilon(p) < 5\%$ for $0 < p \le 0.5$. Asymptotically, it can be shown that

$$\lim_{p\to 0} \varepsilon(p) \approx \frac{c}{-\log_2 p}, \tag{2}$$

where $c = 2 - \left(\log_2 e + \log_2(\log_2 e)\right)$ is a constant. (2) indicates that $\varepsilon(p)$ approaches zero asymptotically as the failure probably decreases, and thus the run-length probing scheme is asymptotically optimal when the network is reliable.

III. Network Transformation From Undirected Topologies to Directed Topologies

A. Optical Links and Link Failure Model

In optical networks, bidirectional communication between adjacent nodes is typically achieved by means of two contra-directional optical fiber links. Thus, an optical link may be abstracted as an undirected graph edge in an undirected graph, or equivalently as a pair of contra-directional arcs in a directed graph. In the following, we adopt the latter abstraction. We further assume that, in addition to being a representation of the physical optical fiber, each directed arc represents the optical amplifier to compensate for the power loss of optical signals. Any of these components can fail due to physical defects or mechanical fatigues. We assume that the state of transmitter and receiver associated with a directed arc can be locally monitored, so that the network management system can poll this information when it signals the probe.

We assume that each directed optical link fails independently with probability of p $(0 \le p \le 0.5)$ over a fixed interval of time, which represents the time duration between fault diagnoses. This assumption of statistical independence among failures is reasonable when "normal" operation of the network is considered, because the equipments (mostly, optical amplifiers) abstracted into each arc operates independently from the equipment abstracted into other arcs. In the event of a catastrophic failure, however, this model is inapplicable and other approaches to ensure network reliability, such as lightpath diversity [5], can be resorted.

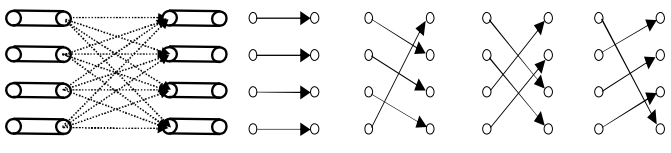

(a) (b)

Fig. 3. Optical network node model: (a) an illustration of 4x4 optical switch fabric; (b) an illustration of some non-blocking directed configurations of a 4x4 optical switch fabric.

B. Optical Nodes and Node Failure Model

In an undirected graph representation of an optical network, a node is an abstraction of an optical switch that is responsible for optically routing signals from input fibers to output fibers. We assume that each network node of degree d is equipped with a $d \times d$ optical switch fabric, switching the optical beam from each input port to any desired output port, as shown in Fig 3(a). We further assume that each input/output port of the optical switch is equipped with a low-cost transponder (economically viable due to the VCSEL technology [11]), whose state of health is locally monitored and reported to the network management system upon polling. We focus on the active components (e.g., the mirrors in MEMS optical switches) in the switch fabric, which could fail from manufacture defects and/or fatigue from normal use.

Under these assumptions, each node i of degree d with a $d \times d$ optical switch fabric can be modeled by a directed bipartite graph, defined as follows:

1. d virtual input nodes correspond to all the input ports of the switch, denoted as $i_k^I, k = 1,2,\cdots,d$;
2. d virtual output nodes correspond to all the output ports of the switch, denoted as $i_k^O, k = 1,2,\cdots,d$;
3. Each virtual input node is connected to all the virtual output nodes via d directed arcs, as shown in Fig. 3(a).

For each node of degree d, there exist $d!$ different non-blocking directed configurations, each comprised of a set of d directed arcs from input nodes to output nodes where no two arcs share the same head/tail node. For example, Fig. 3(b) shows some of the possible configurations for a node of degree 4. At any instance, the switch can take one and only one non-blocking configuration. Therefore, we can use one sample non-blocking directed configuration to model the corresponding network node for the purpose of fault diagnosis.

In an analogy to the link failure model, we assume an independent failure model for each configuration of the optical switch: each input-output connection in the switch fabric fails independently with probability $q\,(0 \le q \le 0.5)$. We assume that, once an input-output connection fails, all switching functions related to the responsible physical component fail simultaneously, that is, the optical signal in the input port cannot be switched to any output port. For any possible configuration of the network node, the independent component assumption suggests that each arc fails independently with probability of q. This simplified node failure model captures the essence of practical switching node failures, and more practical node failure models can be addressed by appropriate extension of this simple model.

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.

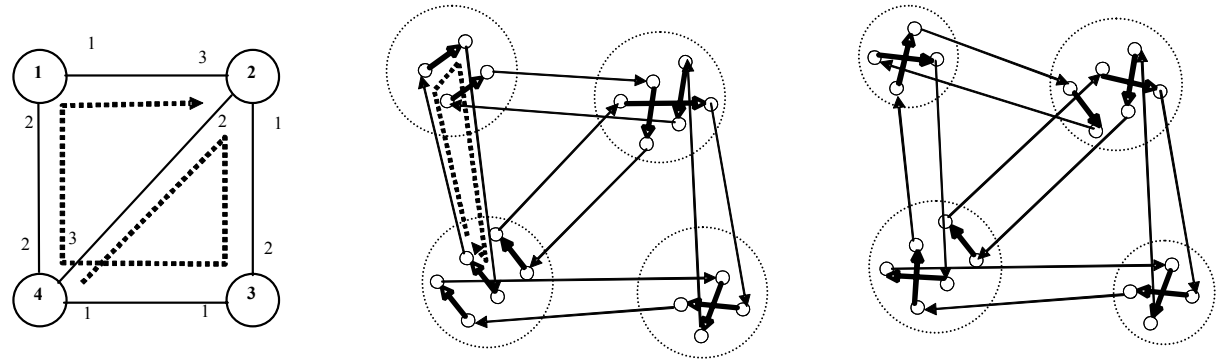

(a) (b) (c)

Fig. 4 An illustration of the network transformation from undirected topologies to directed topologies. (a) the original undirected graph, (b) one transformed directed graph without any Euler trail, (c) one transformed directed graph with an Euler trail.

C. Network Transformation from Undirected Topologies to Directed Topologies

The run-length probing algorithm developed in [1] can only diagnose link failures, whereas, nodes are also vulnerable to failures in practical all-optical networks. To diagnose failures in all-optical networks with directed optical links and possible node failures, we propose a network transformation as follows:

1. We index all the links of each node in the undirected topology, as shown in Fig 4(a).
2. We replace each link (i, j) in the undirected graph with two directed arcs, $i \to j$ and $j \to i$, in opposite directions.
3. We replace each network node of degree d with an empty bipartite graph comprised of $2d$ nodes (i.e., two columns of indexed nodes without any arc connecting them).
4. For each link (i, j) in the original graph, if its index in node i is k and its index in node j is l, we connect the output node i_k^O to the input node j_l^I with the directed arc $i \to j$, and also connect the output node j_l^O to the input node i_k^I with the directed arc $j \to i$.
5. For each node, we choose an appropriate directed configuration such that the transformed graph is Eulerian (details will be elaborated later in this section.)

Fig. 4 demonstrates how the network transformation described above converts an undirected graph into a directed graph. In particular, Fig. 4 (a) depicts the original undirected graph, and Figures 4(b) and 4(c) depict two different directed graphs resulted from choosing different network node configurations.

The transformed directed graph can always be made Eulerian for some appropriate set of node configurations. To see this, first note that the transformation replaces each undirected link with two directed arcs in opposite directions, and the in-degree of each network node (upon step 2 of the transformation) is therefore equal to its out-degree. It follows from the Euler Theorem [9] that the directed graph is Eulerian.

The existence of an Euler trail in the directed graph depends on how the configurations are chosen for all the network nodes. As shown in Fig. 4(a), the original graph has an Euler trail of 4-2-3-4-1-2. However, Fig. 4(b) shows that the directed graph is decomposed into two disjointed cycles and is thus non-Eulerian. On the other hand, in Fig. 4(c), the Eulerian property of the graph is maintained by appropriately choosing the configurations of all the nodes. The set of appropriate node configurations can be identified as follows. After step 2 of replacing each link in the undirected graph with two parallel directed arcs, the resulting directed graph has an Euler trail since the in-degree of each node is equal to its out-degree. To maintain the Eulerian property, each node should be configured to match the way in which some chosen Euler trail passes through the node. For example, in Fig 4(a), the Euler trail passes node 1 from its link 2 to its link 1. It follows that node 1 should be configured as the across state as in Fig. 4(c), instead of the through state as in Fig 4(b). As a result, Fig 4(c) is Eulerian, while Fig. 4(b) is non-Eulerian.

Through the proposed network transformation, both links and nodes in the undirected graph are mapped into directed arcs in the directed graph. Any directed arc connecting two virtual nodes in different switches corresponds to a directed optical fiber link in the all-optical network, which we shall call a link arc; and any directed arc connecting two virtual nodes in the same switch corresponds to a switching component (e.g., a MEMS mirror) in the all-optical network, which we shall call a switch arc. For an original undirected graph of m links and n nodes, the directed graph has $2m$ links arcs, and $2m$ switch arcs. In this way, the transformation maps both link and node failure in the original graph into arc failures in the transformed graph, which can be identified by the run-length probing scheme as shown in next section.

IV. Applications of Run-Length Probing Schemes for All-Optical Networks with Probabilistic Link/Node Failures

In this section, we address how to employ the run-length probing scheme for fault diagnosis to all-optical networks with probabilistic link/node failures through our proposed network transformation.

To detect and localize node/link failures in an all-optical network, we first employ the network transformation to obtain the directed network topology, and then identify a directed Euler trail in the transformed graph. We observe that link arcs and switch arcs appear alternatively along any Euler trail. Without loss of generality, we assume that the directed Euler

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.

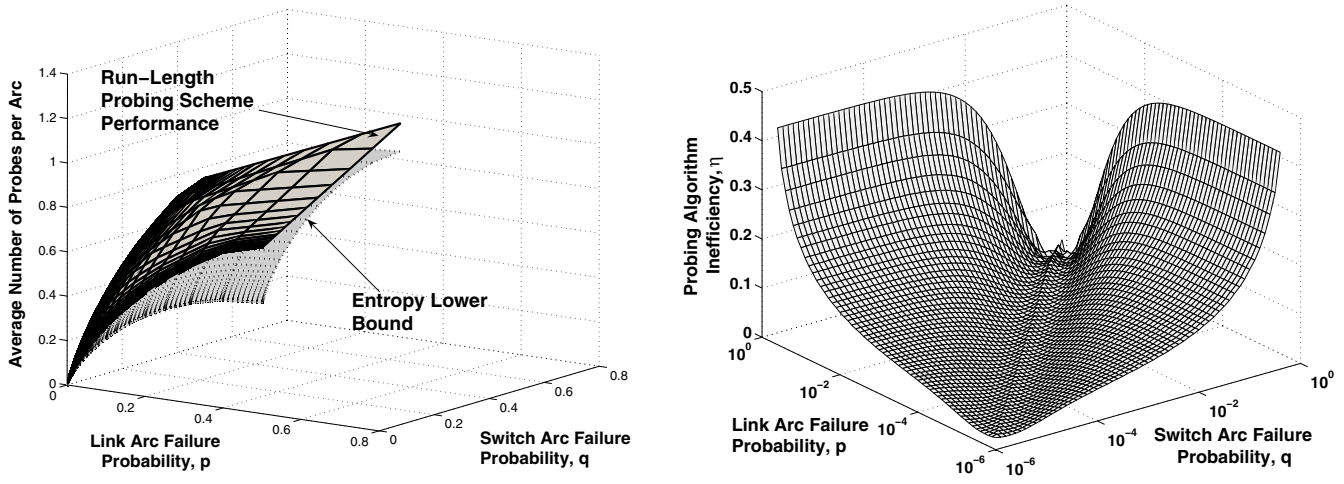

(a) (b)

Fig. 5 The performance of the run-length probing scheme for all-optical networks with probabilistic link/node failures. (a)The average number of probes per component is compared to the entropy lower bound. (2) The probing algorithm inefficiency is plot for different link arc failure probability and switch arc failure probability pairs.

trail starts from a link arc and ends with a switch arc. However, the difficulty here is that the failure probability of each arc along the Euler trail is heterogeneous, while the run-length probing scheme requires that all the arcs fail independently with the same probability. Our proposed solution is to combine an adjacent switch arc and link arc into a virtual arc with failure probability of:

$$r = p + q - pq\,. \tag{3}$$

This combination results in a directed Euler trail of length $2m$, in which the failure probability of each virtual arc in the Euler trail is homogenous. Hence, the run-length probing scheme is now applicable.

After the virtual arc combination step, the probing scheme follows a two-stage procedure. In the first stage, we employ the run-length probing scheme along the Euler trail to identify all the faulty virtual arcs. For a reasonably large network, the average number of probes per virtual arc is approximately equal to:

$$\bar{L}^{arc} = \begin{cases} \bar{L}_\infty\left(p+q-pq\right) & \text{if } 0 < p+q-pq < \dfrac{3-\sqrt{5}}{2} \\ 1 & \text{if } p+q-pq \geq \dfrac{3-\sqrt{5}}{2} \end{cases}. \tag{4}$$

where $\bar{L}_\infty\left(p+q-pq\right)$ is defined in (1). Note that when the failure probability is higher than $\left(3-\sqrt{5}\right)/2$ (the golden ratio), the run-length probing scheme always probes each virtual arc individually. Hence, in this case, the average number of probes per arc is always equal to 1 as indicated in (4).

After this first stage, among all the $2m$ virtual arcs, the average number of failures is $2m(p + q - pq)$. Conditioning on the fact that one virtual arc fails, there are three possible failure scenarios: (1) a single switch arc failure with probability of $(1-p)q/r$, (2) a single link arc failure with probability of $p(1-q)/r$, or (3) a combined switch/link arc failure with probability of pq/r.

During the second stage of the scheme, we deploy additional probes, using the built-in lasers in the optical switch, to determine which of the above three scenarios has occurred for each faulty virtual arc. At this point, we have two alternatives to proceed. If we probe the link arc first, with probability of $(1-p)q/r$ the link arc is fault-free and we conclude, with only one probe, that the switch arc has failed; or with probability of $1-(1-p)q/r$ that the link arc has failed and we have to continue to probe the switch arc with a second probe. Under this alternative, the number of additional probes for each faulty virtual arc is given by

$$\bar{L}_1^c = 1\cdot\frac{(1-p)q}{r} + 2\cdot\left[1-\frac{(1-p)q}{r}\right] = 2+\frac{pq}{r}-\frac{q}{r}\,. \tag{5}$$

Similarly, the second alternative is to probe the switch arc first and probe the link arc subsequently, if necessary. Under this alternative, the number of additional probes for each faulty virtual arc is given by:

$$\bar{L}_2^c = 1\cdot\frac{p(1-q)}{r} + 2\cdot\left[1-\frac{p(1-q)}{r}\right] = 2+\frac{pq}{r}-\frac{p}{r}\,. \tag{6}$$

Comparing (5) with (6), we conclude that the optimal strategy depends on the relationship between the link arc failure probability p and the switch arc failure probability q. If $p > q$, we have $\bar{L}_1^c > \bar{L}_2^c$, which suggests that the switch arc should be probed first. On the other hand, if $p < q$, we have $\bar{L}_1^c < \bar{L}_2^c$, which suggests that the link arc should be probed first. It follows that the average number of additional probes for each faulty virtual arc is given by:

$$\bar{L}^c = 2+\frac{pq}{r}-\max\left\{\frac{p}{r},\frac{q}{r}\right\}. \tag{7}$$

Now, by combining our efforts in first identifying all the faulty virtual arcs, and then determining the sources of failure for each faulty virtual arc, we obtain the average number of probes per directed arc (or vulnerable component) as

1-4244-0357-X/06/$20.00 ©2006 IEEE
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE GLOBECOM 2006 proceedings.

$$\bar{L}_{mixed-failure} \approx \frac{1}{4m}\left[2m\bar{L}^{arc} + 2m(p+q-pq)\bar{L}^{c}\right] = \frac{1}{2}\bar{L}^{arc} + \frac{1}{2}(p+q-pq)\bar{L}^{c} \quad (8)$$

We compare (8) with the entropy lower bound of $(H_b(p) + H_b(q))/2$ in Fig. 5. An immediate observation from Fig. 5(a) is that the average number of probes per arc is close to the entropy lower bound, as expected from our previous results on the near-optimality of the run-length probing scheme. This observation also lends support to our approach to fault diagnosis involving network transformations. A second observation, from Fig. 5(b), is that the probing algorithm inefficiency, which is defined in [1, equation (16)] as the ratio between the number of additional probes compared to the entropy lower bound and the entropy lower bound, increases as the difference between the link arc failure probability and the switch arc failure probability increases. This can be understood as follows. When the difference between p and q increases, one kind of failure occurs more probably than the other kind of failure. The general approach treats both the link arc failure and the switch arc failure equivalently. As a result, we pay the penalty for not exploiting in our algorithm the fact that one type of failure dominates the other. Our third observation, from Fig. 5(b), is that when p is equal to q and both approach zero, the probing algorithm inefficiency does not converge to zero as in the link-failure case [1]. In fact, if $p = q$, the link/node failure diagnosis problem is equivalent to the link failure diagnosis problem with twice as many links. It would be better to treat switch arcs and link arcs on equal basis, and thus employ the run-length probing scheme along an Euler trail of $4m$ links. On the other hand, the combination of switch arcs and link arcs definitely sacrifices performance when the failure probability is fairly low, because the two-stage probing procedure is different from the optimal run-length probing scheme.

In summary, the numerical analysis suggests the following rules of thumb for applying the run-length probing scheme to all-optical networks with probabilistic link/node failures. First, when the link failure probability is equal to the switch failure probability, it is better to treat them equivalently and employ the run-length probing scheme over an Euler trail of $4m$ links. Second, when one type of failures dominates, we should focus on the dominant failure. Finally, for all other cases between the aforementioned two extremes, we should choose the proposed virtual arc approach.

V. Conclusion

In this paper, we investigated the fault diagnosis problem for all-optical networks with vulnerable links and nodes. To apply our previous run-length probing scheme to all-optical networks with probabilistic link/node failures, we introduced a network transformation which converts the original undirected graph to the transformed directed graph. Both the link and node failures in the original graph are mapped into arc failures in the transformed graph, thus rendering the run-length probing scheme applicable.

Complementing our previous work in [1], our present investigation reveals that the complexity of the fault diagnosis problem for all-optical networks is related to the information entropy of the network state. In particular, the operational network diagnosis cost can be kept as low as the information entropy of the network state by cleverly designed fault diagnosis algorithms, such as our run-length probing scheme. Our performance analysis verifies that the information entropy is the correct metric in the design of efficient network fault management system, and suggests different approaches to employ the run-length probing scheme efficiently. More broadly, the connection between the distinct research areas of information theory and network management suggests that additional insights into network management issues can be gained from the more mature field of information theory.

As a further step toward the practical fault diagnosis applications, we would like to investigate how often the run-length probing scheme should be employed as a trade-off between the total operating cost and the timeliness of network state.

Acknowledgement

The research in this paper was supported in part by the Defense Advanced Research Projects Agency (DARPA) and the National Science Foundation (NSF). The authors also would like to thank Mr. Guy E. Weichenberg of MIT for his tremendous help in preparation of this manuscript.

References

[1] 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.

[2] http://www.rhk.com/

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

[4] 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.

[5] Y. G. Wen and V. W. S. Chan, "Ultra-reliable communication over vulnerable all-optical networks via lightpath diversity," IEEE Journal of Selected Areas of Communication (JSAC) - Optical Communications and Networking Series, Vol. 23, No. 8, August 2004, pp.1572-1587

[6] C. Rodriguez, S. Rementeria, J. I. Martin, A. Lafuente, J. Muguerza and J. Perez, "A modular neural network approach to fault diagnosis," IEEE Transactions on Neural Networks, Vol. 7, No. 2, March 1996, pp. 326-340.

[7] E. Athanasopoulou and C. N. Hadjicostis, "Realistic approaches to fault detection in networked discrete event systems," IEEE Transactions on Neural Networks, Vol. 16, No. 5, September 2005, pp. 1042-1052.

[8] F. Khunjush, M. W. El-Kharashi, K. F. Li and N. J. Dimopoulos, "Network processor design: issues and challenges," 2003 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing, Vol.1, August 2003, pp. 164-168.

[9] B. Bollobas, Modern *Graph Theory*, Springer-Verlag New York, Inc. 1998.

[10] R. Gallager and D. Van Voorhis, "Optimal source codes for geometrically distributed integer alphabets," *IEEE Transactions on Information Theory*, Vol. IT-21, March 1975, pp. 228-230.

[11] C. J. Chang-Hasnian, "Tunable VSCEL," IEEE Journal on Selected Topics in Quantum Electronics, Vol. 6, No. 6, November/December 2000, pp. 978-987.