Efficient Fault Diagnosis for All-Optical Networks: An Information Theoretic Approach

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

***Abstract*—Network management and control contribute to at least half of the operating cost of current optical networks. All-optical networks with end-to-end transparent lightpaths promise significant cost savings using optical switching at network nodes. However, this cost saving cannot be realized unless the cost of network management is also reduced. In this paper we explore a promising technique towards that goal. The fault diagnosis problem for all-optical networks is investigated via an information theoretic approach, with the objective to minimize the operating 'cost' of failure detection and localization in the optical layer. Under a probabilistic link failure model, we first interpret the *run-length probing scheme* previously developed for Eulerian graphs as a constrained source-coding algorithm, and characterize its performance via the code rate of its corresponding run-length code. We then extend the run-length probing scheme to non-Eulerian graphs via two alternative approaches: the *disjoint-trail decomposition approach* and the *path-augmentation approach*, and obtain their performance analytically. The analytical and numerical results indicate that the run-length probing scheme is asymptotically optimum for both Eulerian and non-Eulerian graphs of large size. The property of the run-length probing scheme also suggests that each probe in an efficient probing scheme should provide approximately one bit of network state information and thus the total number of probes (or equivalently, the operating cost of failure identification) is lower-bounded and approximated by the entropy of the network states. We believe that our approach using Information Theory in an inter-disciplinary effort can provide new insights on network management, and substantial cost-reduction for all-optical networks can be realized.**

I. Introduction

With the emerging deployment of all-optical networks, broadband network services have the potential to become available to the mass population at much lower cost than what can be achievable today. Present day network management constitutes a significant fraction of the cost of operating a network, (~50%) [1]. Future all-optical networks promise significant cost savings via optical switching of high data rate lightpaths at network nodes, reducing electronic processing costs. However, the side benefit of electronic switching namely parity checks at the end of each SONET line is now absent and the network management system must develop a new mechanism to diagnose link and node failures. Otherwise the promised cost savings will not be realized. Therefore, it is desirable to architect a cost-efficient network management system for all-optical networks of the future.

Network management consists of five functions: fault, configuration, performance, security and account management. For optical networks, a lot of emphasis has been put on fault management, whose cost is dominated by detection and isolation of problems that cause failures. Current standard Synchronous Optical Networks (SONET) infer the health of each SONET link by verifying the parity bits embedded in the overhead of the data frames. This approach is a manifestation of the "*single-hop*" test model [2]: signals are transmitted between adjacent nodes in a network to determine the state of the link connecting them. For future all-optical networks, due to the unique property that optical signals can be carried over a lightpath of many interconnected links without necessarily being detected by optically-switched intermediate nodes, we proposed a "*multi-hop*" test model [3] to diagnose several links simultaneously. Specifically, probing signals are sequentially sent along a set of lightpaths over an all-optical network to probe their state of health; and the network state (i.e., failure pattern) is then inferred from this set of end-to-end measurements (i.e., probe syndromes). Each successive probe is dynamically chosen among the set of permissible probes according to the previous probe syndromes to minimize the number of probes.

Theoretically, fault diagnosis can be understood from an information theoretic perspective. The network state can be viewed as a collection of binary-valued random variables; where each variable is associated with a network element, indicating failure/no-failure of that element. A fault diagnosis algorithm uses a number of tests, whose results are called the 'syndromes', to uniquely identify the network state. The objective of the fault diagnosis process is to encode the set of network states with the set of probe syndromes such that the average syndrome length (thus the operating cost of diagnosis) is minimized. The source-coding problem in Information Theory shares a similar objective, which suggests that the single-hop and the multi-hop test models can be compared using an information theoretic framework. As an example, consider a linear network with two links, each of which fails independently with probability of 0.2. With the single-hop tests, the result of each test is '0' (for no failure) with a probability of 0.8 and '1' (for failure) with a probability of 0.2. The information contained in this syndrome is the entropy $H_b(0.8) = 0.72$ bit where $H_b(x) = -x\log_2 x - (1-x)\log_2(1-x)$ is the Shannon entropy function. In comparison, the two-hop test contains $H_b(0.64) = 0.94$ bit of state information. This observation indicates that multi-hop tests can be more informative than single-hop tests, and should be chosen when the failure probabilities are in the 'right' ranges.

Using the information theoretic approach [3], we have first established a mathematical equivalence between the fault diagnosis problem and the source-coding problem, and have developed the asymptotically-optimal run-length probing scheme for Eulerian networks. In this paper, we extend the application of the run-length probing scheme to general non-Eulerian networks via two alternative mechanisms: the *disjoint-trail decomposition approach* and the *path-augmentation approach*. This investigation verifies the guideline for efficient probing schemes [3]: *each probe should provide approximately 1-bit of network state information and the number of probes required is approximately lower-bounded by and equal to the information entropy of the network states.*

II. Fault Diagnosis Problem for All-Optical Networks with Probabilistic Link Failures

A. Formulation of Fault Diagnosis Problem

In our research, all-optical networks are abstracted as undirected graphs. An *undirected graph* G is a pair of sets (V,E), where V is the set of network nodes of size n, and E is the set of optical links of size m. To illustrate the technique succinctly, the nodes are assumed to be invulnerable in this work (the vulnerable node case is being treated in [4]); links are assumed to fail independently with probability p $(0 \le p \le 0.5)$. Moreover, the states of links are assumed not to change over the duration of the fault diagnosis process. It follows that each link state can be modeled by a Bernoulli random variable, taking the value 1 with probability p for link failure and the value 0 with probability $1-p$ for no failure. A network state $s \in S$ is referred to as a realization of all link states, where S denotes the set of all possible network states.

To detect and localize all failures, optical probing signals are sequentially sent along a set of permissible lightpaths in the network. The result of each probe is called the *probe syndrome*, denoted as $r_t = 0$ if all the links along the probe are UP (no failure) and the probing signal arrives successfully; and $r_t = 1$ if any of the links along the probe is DOWN (at least one failure) and the probing signal never reaches the destination. In addition, to reduce the diagnosis effort, each successive probe is determined sequentially according to the previous syndromes. A sequential employment of permissible probes to identify any network state is called a *probing scheme* $\pi \in \Pi(G)$, where $\Pi(G)$ is the set of all probing schemes for the network G.

We can associate a probing scheme with the network operating cost by assigning each probe with a cost value according to a pre-determined cost function. In this research, each probe, if employed, is assumed to cost one unit of diagnosis effort. Consequently, the probing cost of state s, denoted by l_s^π, is equal to the number of probes to identify network state s with the probing scheme π. Normalizing over the network size (i.e., the number of links m), we associate the probing scheme π with a figure of merit (cost) of the average number of probes per link as

$$\overline{\mathcal{L}}_\pi = \frac{1}{m}\sum_{s\in S}\Pr(s)\, l_s^\pi , \tag{1}$$

where $\Pr(s)$ is the prior probability of occurrence of the network state s.

For a network G, the objective is to find the probing scheme that minimizes the average number of probes per link, and thus minimizing the operating cost of fault diagnosis. This can be achieved by solving the following optimization problem,

$$\begin{aligned} \min_{\pi} \quad & \overline{\mathcal{L}}_\pi = \frac{1}{m}\sum_{s\in S}\Pr(s)\, l_s^\pi \\ s.t. \quad & \pi \in \Pi(G) \end{aligned} . \tag{2}$$

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

$$\overline{\mathcal{L}}^* = \min_{\pi\in\Pi(G)}\left\{\frac{1}{m}\sum_{s\in S}\Pr(s)\, l_s^\pi\right\} = \frac{1}{m}\sum_{s\in S}\Pr(s)\, l_s^{\pi^*} , \tag{3}$$

where π^* is the optimum probing scheme.

B. A Source-Coding View of Fault-Diagnosis Problem

The fault-diagnosis problem can be understood as the source-coding problem in Information Theory under some physical constraints. This association has important implications in designing efficient network diagnosis schemes.

For the probing scheme π, we denote the probe syndrome of the network state s as $r(s) = r(t_1^s)r(t_2^s)\cdots r(t_{l_s^\pi}^s)$, where $\{t_1^s, t_2^s, \ldots t_{l_s^\pi}^s\}$ is the sequence of probes applied to identify the network state s. The set of all probe syndromes is denoted as $R = \{r(s), s \in S\}$. The diagnosability of any valid probing scheme requires that a one-to-one mapping exists between the set of network states and the set of probe syndromes. In fact, it has been shown [3] that the set of probe syndromes R of a probing scheme π forms a uniquely-decodable code. It follows that the fault diagnosis problem is mathematically equivalent to the well-established source-coding problem in Information Theory. Specifically, the objectives of both problems are equivalent, i.e. to design a probing/coding scheme mapping the set of network-states/source-alphabets into a set of probe-syndromes/codewords such that the average syndrome/codeword length is minimized.

This mathematical equivalence between the fault-diagnosis problem and the source-coding problem suggests that we can exploit the rich set of results from the source-coding literature to construct efficient fault-diagnosis algorithms. First, it follows from the lossless source coding theory that the minimum average number of probes per link is lower bounded by the information entropy of individual link, i.e.,

$$\overline{\mathcal{L}}^* \ge H_b(p), \tag{4}$$

where $H_b(p)$ is the Shannon information entropy function. Second, this alternative interpretation also suggests that we can translate optimal and/or sub-optimal source coding algorithms into efficient fault diagnosis schemes. However, not all source coding algorithms, e.g. the optimal Huffman coding algorithm,

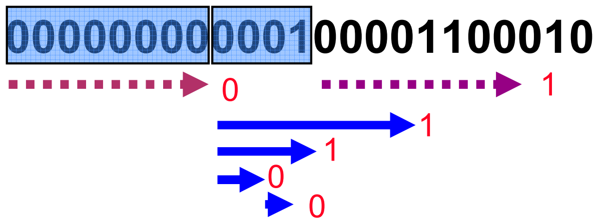

Fig.1. Demonstration of the run-length probing scheme over an Eulerian network. It contains a sequence of concatenations of two alternating phases: the failure detection phase (dotted lines) and the failure localization phase (solid lines).

can be transformed into fault diagnosis algorithms. The Huffman algorithm can be understood as a sequence of YES/NO questions in the form of "Is the source realization in some set A?" Translated into the context of fault diagnosis, they correspond to questions such as "Is link 1 UP?" or "Is Link 1 UP and link 3 DOWN and link 5 UP?" Not all of such questions are physically realizable probes, which can only probe consecutive links, corresponding to one particular class of questions such as "Are links 2, 3, 4 all UP?" Thus, the nature of permissible probes posts an extra restriction that only a special class of questions can be asked. In our research, we refer to this restriction as the "*consecutive probing constraint*", and study the fault diagnosis problem, or the equivalent source coding problems, under this constraint.

III. RUN-LENGTH PROBING SCHEMES FOR EULERIAN ALL-OPTICAL NETWORKS

In this section we interpret, from an information theoretic perspective, the run-length probing scheme [3] previously developed for any Eulerian network, which contains at least one Euler trail (i.e., a sequence of interconnected links containing all the links in the topology without repetition).

For an Eulerian network, we can introduce a natural order to any network state by indexing all the link states along an Euler trail in the network. Specifically, any network state must have the format of $0^{i_1}10^{i_2}1\cdots0^{i_{L-1}}10^{i_L}$ where $i_1, i_1, \cdots i_L$ are non-negative integers and 0^i means a run of i '0', and each of the segments, 0^i1, is called a sub-state. Considering that any probe can locate at most one faulty link at a time, each of such sub-states should be encoded separately. This idea suggests that we should, instead of coding for binary input streams, code on the symbol set of $Z_0 = \{0^i1\}_{i=0}^{\infty}$ with a geometrical probability distribution. In the context of source coding, the optimal code for the set $Z_0 = \{0^i1\}_{i=0}^{\infty}$ with geometrical distributions has been shown as the run-length code [5]. A natural question is whether the run-length coding algorithm can be translated to some corresponding fault diagnosis algorithm under the consecutive probing constraint. The run-length codeword of alphabet 0^i1 is a concatenation of two prefix codes: the unary code for the integer $\lfloor i/K \rfloor$ followed by the Huffman code for the alphabet $0^{i \bmod K}1$, where $K = \lceil -\log_{1-p}(2-p) \rceil$. The unary code for an integer j is the codeword with j zeros followed by a single one, i.e., $u(j) = 0^j1$. In the fault diagnosis context, such a

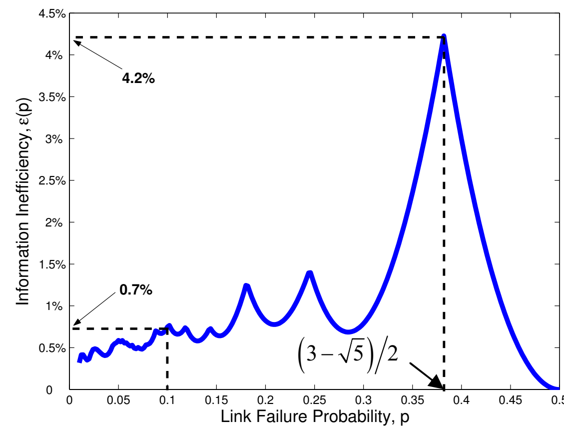

Fig. 2. The information inefficiency of run-length probing schemes for different link failure probability (adapted from [3]) is uniformly upper bounded by $\varepsilon(p) \le 4.2\%$, where the equality is achieved when the link failure probability equals to the golden ratio of $p = \left(3-\sqrt{5}\right)/2$.

unary code can be implemented by sending $j+1$ back-to-back probes of length K along the Euler trail. At the same time, the Huffman codeword for the alphabet 0^k1 ($0 \le k \le K-1$) can be implemented by the 2^{m}-splitting binary searching algorithm developed in [3], which balances the Huffman code tree and maximizes the information gain of each probe. Therefore, the concatenation structure of the run-length code guarantees its transferability to a corresponding fault diagnosis algorithm, called the *'run-length probing scheme'*.

Algorithmically, as a mirror of its concatenation structure in the run-length code, the run-length probing scheme alternates two diagnosis phases to identify each faulty link: the *failure detection* phase and the *failure localization* phase. In the failure detection phase, a detection probe is sent over a set of K (called the maximum probing length) consecutive links along the Euler trail. If all the links are fault-free, we move onto the next set of K consecutive links along the Euler trail. If on the other hand the detection probe returns the syndrome '1', the algorithm enters the failure localization phase. In this phase, given that there is some failure in the detection probe, the "2^m-splitting" binary searching algorithm [3] is employed to locate the leftmost faulty link. After the fault is localized, the algorithm resumes the failure detection phase by sending another probe spanning K links along the trail right after the failure. As an illustration, Fig. 1 demonstrates how to employ the two-phase probing scheme for efficient fault diagnosis.

Since the probe syndrome of any network state under the run-length probing scheme is a concatenation of a series of run-length codewords for the set of sub-states $\{0^i1 : i \ge 0\}$, we can approximate the average number of probes per link required for the run-length probing scheme by the code rate of its corresponding run-length code [5], i.e.,

$$\overline{\mathcal{L}}_{\infty}(p) = p \cdot \left(\lfloor \log_2 K \rfloor + 1 + \frac{(1-p)^L}{1-(1-p)^K} \right) \qquad 0 < p \le \frac{1}{2}, \qquad (5)$$

where $K = \lceil -\log_{(1-p)}(2-p) \rceil$ and $L = 2^{\lfloor \log_2 K \rfloor + 1} - K$. Comparing

(5) to the entropy bound (4), we have also proved that the run-length probing scheme is ε-optimal, i.e.,

$$H_b(p) \le \overline{\mathcal{L}}_\infty \le \left[1+\varepsilon(p)\right] H_b(p), \qquad (6)$$

where the information inefficiency, $\varepsilon(p)$, tends to decrease with smaller link failure probability and is uniformly upper bounded by $\varepsilon(p) \le 4.2\%$ where the equality is achieved at $p = \left(3-\sqrt{5}\right)/2$ (the golden ratio) as illustrated in Fig. 2. This indicates that the performance of the run-length probing scheme is always less than 5% larger than the entropy lower bound. In practical networks with fairly reliable components, both the upper and the lower bounds in (6) are reduced to the entropy of individual link, suggesting that the run-length probing scheme is asymptotically optimum for large Eulerian all-optical networks, even if much of topological information is suppressed by probing over an Euler trail.

IV. Run-Length Probing Schemes for Non-Eulerian All-Optical Networks

To employ the run-length probing scheme, we assume in Section III that the network is Eulerian. This requires that all (or except two) the nodes in the network have even degrees [6]. However, practical all-optical networks may not satisfy this condition and thus the run-length probing schemes cannot be applied directly. In this section, we propose two alternative approaches to apply the run-length probing scheme to non-Eulerian topologies and characterize their corresponding cost performance analytically.

A. The Disjoint-Trail Decomposition Approach

Any non-Eulerian graph can be decomposed into a set of link-disjointed trails, among which no two trails share the same link. The set of link-disjointed trails can be identified via a sequential deletion procedure. We start from any node and walk along the graph until we have to pass some link twice. The set of passed links forms a trail (a sequence of interconnected links without repetition), and are deleted from the graph. The same trail deletion process is resumed from any other node of non-zero degree until the graph is empty. For example, in Fig. 3(a), the sequential deletion procedure results in two link-disjointed trails in the non-Eulerian network, i.e., trail A-B-C-D-E-F-G-H-I-J-B and trail C-M-L-K-J.

After the decomposition step, the run-length probing scheme can be applied to each link-disjointed trail sequentially. The network state is uniquely identified after all the trails have been probed.

Unfortunately, the decomposition could potentially break one sub-state 0^i1 into two sub-states of $0^{i'}$ and $0^{i-i'}1$ on two link-disjointed trails. The number of probes to identify sub-state 0^i1 is at least less than the number of probes to identify two sub-states of $0^{i'}$ and $0^{i-i'}1$, where the additional number of probes is upper bounded by 1. If the number of individual link-disjointed trails is T, the average number of probes per link is given by

$$\overline{\mathcal{L}}_\infty(p) \le \overline{\mathcal{L}}^{non-Euler}_{run-length} \le \overline{\mathcal{L}}_\infty(p) + T/m\ . \qquad (7)$$

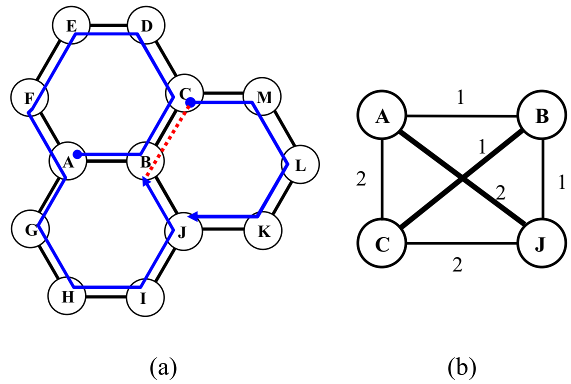

Fig. 3. (a) Each non-Eulerian graph can be decomposed into a set of non-overlapping trails. (b)The complete graph M to identify the minimum set of replicated links.

Since each link-disjointed trail reduces the number of odd-degree nodes by two, we conclude that $T = n_o/2$, where n_o is the number of odd-degree nodes in the network, and thus the upper bound becomes $\overline{\mathcal{L}}_\infty(p) + n_o/2m$.

In particular, we are interested in the class of non-Eulerian regular topologies, considered for cost-optimized all-optical network architectures in [7]. A graph is said to be *regular* of degree d if the degrees of all the nodes are equal to d. For example, the d-nearest neighbors Graph, the symmetric Hamilton Graph and the Moore Graph (with the fully-connected graph as a special case) are the most popular regular graphs considered for all-optical network architectures. The non-Eulerian property suggests that degree d is odd and thus $n_o = n$. Notice that for a regular graph of degree d, the handshake property suggests $n/2m = 1/d$. It follows that, for a non-Eulerian regular graph of degree d, the average number of probes per link is given by

$$\overline{\mathcal{L}}_\infty(p) \le \overline{\mathcal{L}}^{non-Euler}_{run-length} \le \overline{\mathcal{L}}_\infty(p) + 1/d\ . \qquad (8)$$

For cost-optimized architectures (to the first order) of all-optical networks with optical cross-connect (OXC) switches, Guan and Chan [7] have recently shown that, under the assumption of all-to-all uniform traffic, the optimal node degree d asymptotically approaches infinity as the network size (in particular, the number of nodes) approaches infinity while their ratio approaches zero. It follows that, for a cost-optimally architected all-optical network, the upper bound in (8) converges to the lower bound, indicating that the run-length probing scheme is asymptotically optimum for large non-Eulerian regular networks with cost-optimized architectures.

B. The Path-Augmentation Approach

In any network, we can replicate each link once along the shortest path between any two nodes of odd degree to make their degrees even. We call the shortest path between two odd-degree nodes an augmenting path and the above replication operation a path augmentation. Notice that the path augmentation does not change the degree parity of any other nodes along the augmenting path, but reduces the number of odd-degree nodes in the network by two. Since the number of odd-degree nodes in a finite network is always even due to the handshake property (i.e., the sum of node degrees is even) [6], we can convert any non-Eulerian graph into an Eulerian graph via a finite number of path augmentations.

After the path-augmentation step, the run-length probing scheme can be applied along the nominal Euler trail in the resulting Eulerian graph. Upon termination, all the link states have been identified except that a set of redundant links have been probed more than once. If possible, to reduce the diagnosis effort, we can skip those redundant links whose states have been identified previously.

Moreover, we would like to minimize the number of replicated links resulted from the path-augmentation step, via a minimum-weight perfect matching approach. As illustrated in Fig. 3, this approach includes the following four steps:

(1) an all-pair shortest-path algorithm (for example, the Floyd-Warshall algorithm [8]) is run to identify the set of all-pair shortest paths among the set of odd-degree nodes in the original graph (e.g., six distinct shortest paths for the set of odd-degree nodes {A, B, C, J} in Fig. 3(a));

(2) a complete graph M (i.e., Fig 3(b)) is created with the set of odd-degree nodes (i.e., {A, B, C, J}) and the weight of each link as the length of the shortest path connecting the two nodes in the original graph (e.g., the weight of link AJ is 2 because the shortest path connecting node A and node J in Fig. 3(a) is A-B-J);

(3) a minimum-weight perfect matching algorithm (a perfect matching of a graph is a subset of links in the graph that touch all the nodes exactly once [8], which can be identified by the *Edmonds' blossom algorithm* [8]) is run over graph M to obtain a perfect matching (e.g., {AJ, BC} is the minimum weight perfect matching in Fig. 3(b));

(4) the original network G is augmented along the set of paths chosen by the resulted minimum perfect matching except for the augmenting path with the maximum weight, because a graph with two odd-degree nodes is Eulerian. As a result, path B-C is augmented in Fig 3(a) via the dotted link.

In the augmented graph, we can identify a nominal Euler trail, i.e., trail A-*B*-*C*-D-E-F-G-H-I-J-*B*-*C*-M-L-K-J, which passes link B-C twice. Notice that the number of replicated links is 1, which is significantly less than the number of links (14 in this case) in the graph. Moreover, this observation is in general true, as shown for the class of non-Eulerian regular network topologies.

For non-Eulerian regular topologies (i.e., the d-nearest neighbors Graph, the symmetric Hamilton Graph and the Moore Graph) considered for all-optical networks [7], each contains a Hamilton path (a path containing each node exactly once) of size n. It follows that all the augmenting paths can reside along the path and the number of replicated links is $n/2$. Therefore, for a regular graph of degree d, the average number of probes per link under the path-augmentation approach is given by

$$\overline{\mathcal{L}}_\infty(p) \le \overline{\mathcal{L}}_{run-length}^{non-Euler} \le \overline{\mathcal{L}}_\infty(p) + \overline{\mathcal{L}}_\infty(p)\big/d . \qquad (10)$$

For cost-optimized architecture whose optimal node degree d asymptotically approaches infinity as the network size (in particular, the number of nodes) tends to infinity [7], the upper bound in (10) converges to the lower bound, verifying that the run-length probing scheme is asymptotically optimum for large non-Eulerian regular networks with cost-optimized architectures.

V. Conclusion

In this paper, we address a very important cost driver for future networks by proposing a new network diagnosis technique that can substantially reduce network operating costs. We investigated the fault diagnosis problem for all-optical networks with probabilistic link failures via an information theoretic approach. Our research reveals that the complexity of the fault management system of all-optical networks can be related to the problem of how to represent the network states efficiently. In particular, we have shown that the failure identification cost can be kept as low as the information entropy of the network state by our proposed run-length probing scheme, which exploits the unique property of all-optical networks. We believe that this research suggests fruitful connections between the two distinct research areas, i.e., Information Theory and network management, and will provide substantial insights and cost savings over current practices.

For future research, it would be interesting to address the problem of how to diagnose failures when the network management system has no or limited prior knowledge of the link failure probability via an information theoretic approach.

Acknowledgement

The research in this paper was supported in part by the Defense Advanced Research Projects Agency (DARPA) and the National Science Foundations (NSF). The authors also would like to thank Prof. Muriel Medard of MIT for her encouragement to present this work to the Information Theory community.

References

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

[2] F. P. Preparata, G. Metze and R. T. Chien, "On the connection assignment problem of diagnosable systems," *IEEE Transaction on Electronic Computers*, Vol. 16, No. 6, December 1967, pp. 848-854.

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

[4] Y.G. Wen, V. W. S. Chan and L. Z. Zheng, "Efficient Fault Diagnosis for All-Optical Networks with Probabilistic Link/Node Failures," *Submitted to GlobeCom'06*.

[5] 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.

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

[7] C. K. Guan and V. W. S. Chan, "Topology design of OXC-switched WDM networks," *IEEE Journal of Selected Areas in Communications*, Vol. 23, No. 8, August 2005, pp. 1670-1686.

[8] R. K. Ahuja, T. L. Magnanti and J. B. Orlin, *Network Flows*, Prentice Hall, 1993.