Active Monitoring and Alarm Management for Fault Localization in Transparent All-Optical Networks

Sava Stanic, *Member, IEEE*, Suresh Subramaniam, *Senior Member, IEEE*, Gokhan Sahin, *Member, IEEE*, Hongsik Choi, and Hyeong-Ah Choi

***Abstract*—Achieving accurate and efficient fault localization in large transparent all-optical networks (TONs) is an important and challenging problem due to unique fault-propagation, time constraints, and scalability requirements. In this paper, we introduce a novel technique for optimizing the speed of fault-localization through the selection of an active set of monitors for centralized and hierarchically-distributed management. The proposed technique is capable of providing multiple levels of fault-localization-granularity, from individual discrete optical components to the entire monitoring domains. We formulate and prove the NP-completeness of the optimal monitor activation problem and present its Integer Linear Program (ILP) formulation. Furthermore, we propose a novel heuristic whose solution quality is verified by comparing it with an ILP. Extensive simulation results provide supporting analysis and comparisons of achievable alarm-vector reduction, localization coverage, and time complexity, for flat and hierarchically distributed monitoring approaches. The impact of network connectivity on fault localization complexity in randomly generated topologies is also studied. Results demonstrate the effectiveness of the proposed technique in efficient and scalable monitoring of transparent optical networks.**

***Index Terms*—Transparent optical networks, fault detection, fault localization, monitoring, alarm processing.**

I. Introduction

EMERGING transparent optical networks (TONs) introduce many advantages, including the ability to efficiently leverage large bandwidth potential and provide transparent support for diverse transmission protocols. However, optical transparency also introduces a requirement for a new monitoring and fault-localization approach at the optical layer. Due to the lack of optical/electro/optical (O/E/O) regeneration in TONs, a single fault may propagate throughout the network, thus generating a flood of redundant alarms, increasing the processing overhead and localization time, and ultimately delaying service restoration. With inherently high data rates, even a short interruption may have a catastrophic effect. Frequent disruptions in TONs include: bending or cutting of fiber, equipment failure, human error, and sophisticated attacks [1].

Ideally, physical failures should be detected, localized, and resolved at the optical layer before they are noticed and handled by the higher layer protocols. Legacy monitoring techniques in SONET/SDH provide 50ms optical layer restoration, but cannot be directly applied in TONs due to required O/E/O conversion at each node.

Most of the recently proposed TON monitoring techniques are based on different schemes for establishment of dedicated supervisory cycles or paths which are provisioned on dedicated wavelengths through the monitored network elements [3], [5], [14], [17]–[19]. However, these schemes introduce additional overhead in terms of the required bandwidth, dedicated transponders and monitors, supervisory-path computation and provisioning time, added channel interference, and necessary maintenance. Furthermore, only monitoring of the edge failures is commonly considered, and the proposed techniques may become inefficient for detailed localization of individual fiber-spans and discrete optical components in large topologies.

In order to address the above mentioned deficiencies, we propose a new approach to fault-localization in TONs that efficiently utilizes the existing traffic lightpaths and commonly integrated monitors [9], [13], [29]. The redundant alarms due to optical transparency result in suboptimal alarm codes that increase the complexity of the fault-localization problem. Minimizing the number of alarms reported to the network manager without compromising on fault coverage is crucial for rapid and accurate localization of faults as well as the stability of management systems against ever-increasing amounts of alarm data [7].

We formulate the problem for optimization of alarm-vectors and present the optimal and heuristic solutions. To address the localization accuracy in large network topologies while maintaining low complexity, we also expand this problem into a hierarchically-distributed monitoring model.

The proposed hierarchical monitoring model enables parallel optimization of fault-localization within independent monitoring domains, scalability, and multiple fault-localization granularity levels. Since the approach is not dependent on dedicated supervisory lightpaths, it is applicable for efficient integration within the existing TON monitoring and management systems.

Manuscript received September 14, 2009; revised February 26, 2010 and March 7, 2010. The associate editor coordinating the review of this paper and approving it for publication was A. Sethi.

Parts of the work presented in this paper were presented in [9], [10], [24].

S. Stanic and S. Subramaniam are with the Dept. of ECE, and H.-A. Choi is with the Dept. of Computer Science, The George Washington University, Washington DC, USA (e-mail: {stanic, suresh, hchoi}@gwu.edu).

G. Sahin is with the Dept. of ECE, Miami University, Oxford OH, USA (e-mail: sahing@muohio.edu).

H. Choi is with the Dept. of Computer Science, Virginia Commonwealth University, Richmond VA, USA (e-mail: hchoi@vcu.edu).

Digital Object Identifier 10.1109/TNSM.2010.06.I9P0343

A. Related Work

Previous work on fault diagnosis and localization includes [1]–[28]. The problem of identifying faults assuming that a fault propagates downstream on all lightpaths from the point of fault origin is considered in [12], [13]. A single-link failure detection scheme in [4] performs failure detection by assigning monitors to each optical multiplexing and transmission section. In [11], a central manager is assumed, and an algorithm for single fault identification is presented. The central manager periodically tests all source and destination powers using the routing table information. If some node's power is out of expected bounds, the possible source of the fault is identified. In [12], fault identification through filtering alarms is discussed. Using a fault identification tree of depth equal to the number of alarming components, the set of potential fault sources is narrowed down. A survey of fault detection capability at each layer is presented in [13]. Approaches to the fault location problem are also classified.

A problem on monitor placement was considered in [20]. The authors show that a wavelength to be used for probing other nodes from a monitor node is highly likely to be available in dynamic traffic conditions. Based on this, it is claimed that fault detection and identification can be done successfully with high probability with a small number of monitor nodes. A heuristic algorithm for monitor placement that is based on clustering, as well as the idea of setting up additional test connections to collect diagnostic information, is proposed. In [14], a fault detection scheme for optical mesh networks based on decomposition of network topology into monitoring-cycles (test-lightpath loops) is presented. The authors present heuristics for the construction of a monitoring cycle cover that minimizes cycle overlap for a given network topology.

A monitoring scheme in [15] partitions the network into islands and determines a node or link failure using an island-by-island restoration protocol. Islands are recomputed whenever the topology changes. A fault-localization protocol in [16] determines a limited-perimeter around the shortest affected path through flooding, and then locates the failure within limited-perimeter through the exchange of multicast fault-vectors. Pointurier et al. [3] consider estimating the Q factor for all lightpaths given that only some nodes have monitors by establishing additional test-lightpaths and using Q factor correlation between established paths to monitor paths that terminate at nodes without monitors.

In [17], authors present a scheme of utilizing supervisory cycles and paths to localize a single link fault from a single monitoring location. Wu et al. [18] propose the use of m-trails as a more flexible monitoring scheme for minimizing the total monitoring cost. It is shown that the use of m-trails outperforms the m-cycle counterpart. A work in [19] investigates the use of m-trails for fault-localization such that the sum of monitoring cost and bandwidth cost is minimized. Introduced is a simple heuristic based on random code assignment and swapping for m-trail design problem.

B. Contributions

This work investigates the optimization of fault localization in transparent optical networks through the selective activation of available monitors in flat and hierarchically-distributed monitoring architectures. The paper makes the following contributions:

- Introduces a novel approach for efficient TON monitoring that utilizes the available traffic lightpaths and reduces the complexity of fault localization through the fault-vector optimization for centralized and hierarchically-distributed fault-management. Resulting rapid fault localization enables the use of fast and efficient link-based restoration techniques.
- Proves that the proposed hierarchically-distributed monitoring has identical fault-detection capabilities as centralized fault monitoring.
- Formulates the monitor optimization problem, presents complexity analysis, and proves the NP-completeness.
- Formulates an Integer Linear Program (ILP) for the optimal monitor set activation problem.
- Introduces an effective heuristic algorithm for centralized and distributed monitoring and presents its complexity.
- Provides extensive numerical results and performance comparisons for the centralized and hierarchically-distributed monitoring in common and random network topologies.

The paper expands our work in [9], [10], [24] by contributing novel hierarchical and flat localization and performance comparisons for several network topologies, studies the time complexity for both monitoring models, introduces and analyzes a new problem of achieving full fault-localization monitoring through the use of randomly added shortest paths and compares it to results obtained by selecting the minimum number of all shortest paths. Furthermore, it analyzes the impact of network connectivity on fault localization complexity in randomly generated network topologies.

C. Paper Organization

The rest of the paper is organized as follows. Section II covers common optical monitoring techniques and fault propagation models, defines monitoring models used throughout the rest of the paper, and proves equivalent fault-detection between hierarchical and centralized monitoring. Section III provides a simple illustrative example of the optimized monitor activation. In Section IV we formulate the problem, address its complexity, and prove its NP-completeness. Section V provides an ILP formulation for the fault-vector optimization problem. Section VI presents an efficient heuristic algorithm for the problem. In Section VII we provide extensive numerical results and performance comparisons for different monitoring scenarios and topologies. Finally, Section VIII concludes the paper with a brief summary of contributions and presents interesting ideas for possible extensions.

II. Network Model

A. TON Monitoring

Fig. 1 shows the model of a proposed monitoring approach. For a given topology and established set of lightpaths, the proposed optimization algorithm computes the optimal activation set of optical monitors and provides that information to

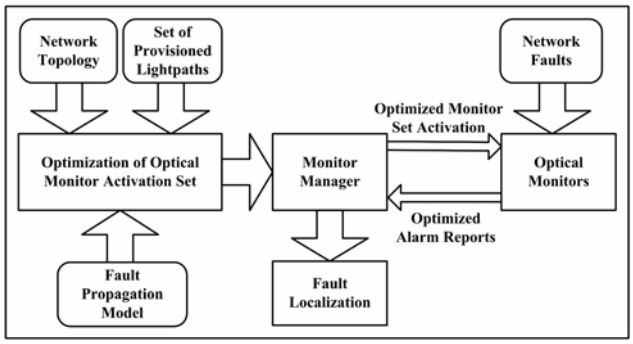

Fig. 1. TON monitoring model.

TABLE I
OPTICAL MONITORING TECHNIQUES

Monitoring Technique	Monitoring Domain	Monitoring Information
Wide-band optical power detection	Optical	Measures the aggregate optical power on fiber.
Wavelength optical power detection	Optical	Measures the optical signal power on a specific optical wavelength.
Optical Spectrum Analyzer	Optical	Measures the channel power, power stability, channel center wavelength and spacing, wavelength stability, OSNR, and total optical power.
Oscilloscope (Eye-Diagram)	Optical	Used for monitoring Eye-Diagram for determining transmission system's performance and Inter-Symbol-Interference (ISI).
Optical Pilot Tones (Out of Band or Sub-Carrier Multiplexed)	Optical	Specially designed test signals that are generated on selected lightpaths to determine channel degradations.
Q-factor	Optical	Q-factor provides quality-of-signal information and can be used to estimate BER.
Optical Time Domain Reflectometry (OTDR)	Optical	Injects a series of optical pulses into the fiber under test and monitors light that is reflected back from points where the index of refraction changes.
BERT	Electrical	Measures the Bit Error Rate in received electrical signal.

the monitor manager, which configures monitors accordingly. When a fault occurs the optimized set of activated optical monitors reports alarms, which are then processed by the assigned fault manager to perform fast fault localization. If some monitors are not configurable, the manager could also filter the required set of received monitors.

B. Optical Layer Monitoring

In TONs the optical layer includes the optical channel (OCh) layer, the optical multiplex section (OMS) layer (line layer), and the optical transmission section (OTS) layer [24]. Our approach considers all available monitors within these three optical sub-layers. Table I shows some of the commonly used monitoring techniques. They include: Wide-band and wavelength optical power monitoring, optical spectrum analysis, eye-diagram, optical pilot tones, Q-factor, Optical Time Domain Reflectometry (OTDR), and Bit-Error-Rate (BER).

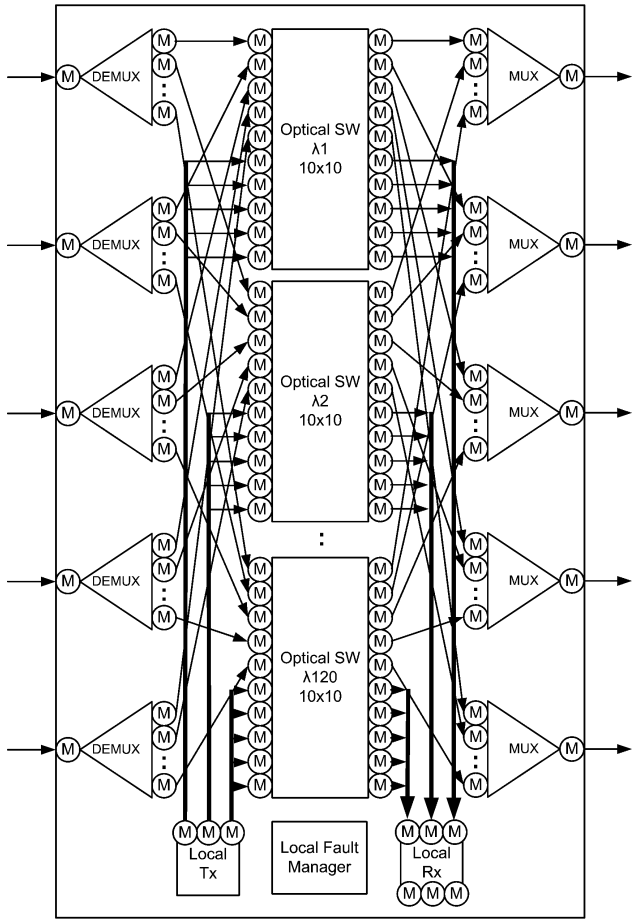

Fig. 2. Monitoring of a small OXC.

Based on their monitoring capabilities, common optical network components can be categorized as passive or active. Some passive optical components (i.e. optical fiber) do not have any monitoring capabilities, and depend on other components' monitors for fault localization and restoration. Other optical components (i.e. optical amplifiers) have monitors and are able to report alarms, but cannot handle restoration. Finally, active optical components such as Optical Cross Connect (OXC) have built-in monitoring and service restoration capabilities and are able to report and manage alarms when an abnormal condition occurs. Throughout this paper we assume that the optical layer is monitored with optical power monitors due to their common availability in network components [9]. Each monitor is capable of reporting a binary alarm to its assigned network manager when the input optical power level is outside of the configured threshold level. Furthermore, each monitor's alarm reporting can be activated or deactivated by its assigned fault-manager.

The following figures depict monitor placement (with letter M) for various optical components that form OXC (Fig. 2) and bi-directed edge (Fig. 3) [9], [13], [29]. These models were used in simulations for obtaining discrete optical-component localization results in Section VII.

C. Fault Propagation Model

Due to the lack of O/E/O regeneration at each optical node, faults may propagate throughout various parts of the

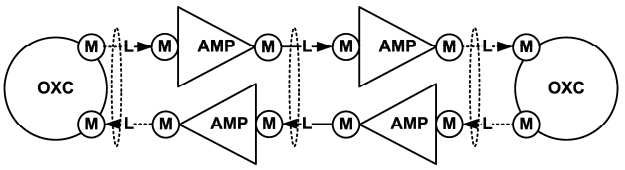

Fig. 3. Monitor placement on a bi-directional fiber edge.

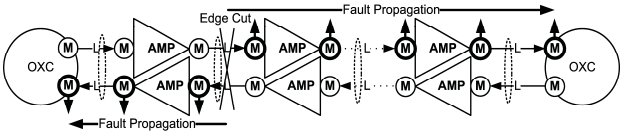

Fig. 4. Fault propagation due to bi-directional edge cut.

network resulting in a large number of redundant alarms. Similarly, imperfections in optical devices may lead to complex fault/attack propagation models due to optical crosstalk, amplifier gain competition, and fault masking. Failure types can be categorized by the effect they have on the individual wavelengths into wavelength-specific and shared risk link groups [17]. Some failures, such as switch-fabric failure, affect individual wavelengths. Other failures, such as optical-amplifier-failures, affect all wavelengths that traverse that component. We consider both types of failures. It is assumed that faults propagate in the downstream direction along all the lightpaths that pass through the fault origin. Such a model is consistent with the ones used in [10], [11], [13]. Fig. 4 shows a simple example of the fault propagation model due to a bidirectional edge-cut for an expanded view of the network edge. All downstream monitors through which affected lightpaths pass will report an alarm to their assigned fault-manager.

Similarly, Fig. 5 shows the fault propagation due to a failure of the switch fabric component inside an OXC node for the wavelength $\lambda 3$. Assuming that four lightpaths $\{LP1(\lambda 3), LP2(\lambda 2), LP3(\lambda 3), LP4(\lambda 2)\}$ are established as shown, a failure of the switch $\lambda 3$ will only propagate downstream along the lightpaths that use this specific wavelength. Lightpaths on different wavelengths that pass through this node will not be affected and will not propagate this failure downstream.

D. Monitor and Alarm Management

The complexity of the monitor optimization and fault localization problems grows with increasing network size and granularity of fault-localization. Management systems can be divided into: centralized management, distributed management, and hierarchical management. The centralized management is primarily used in small networks; distributed management is scalable and can be applied to networks of moderate size, and hierarchical management is usually used in large scale networks. In this paper we consider fault localization within centralized-flat and hierarchically-distributed management models.

1) Centralized Monitoring: A centralized monitoring model consists of a single central fault-manager which receives alarms from all monitors in the network and processes

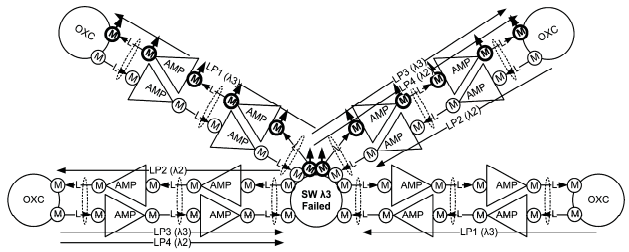

Fig. 5. Fault propagation due to failure of a wavelength switch.

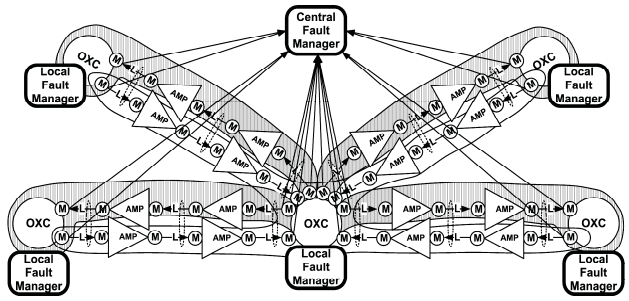

Fig. 6. Hierarchically-distributed fault monitoring.

them as a fault-vector to localize a fault. As shown in Table II, we consider two centralized monitoring models: Flat Centralized Top Level (OXC and edge level monitoring) and Flat Centralized Detail Level (discrete optical component level monitoring), Table II defines the monitor placement as well as the fault detection and localization capability for both models. With centralized monitoring, fault propagation can flood the central manager with a large number of redundant alarms, delaying fault localization and service restoration. Optimizing the set of activated monitors can effectively reduce the number of redundant alarms reported to the central manager and simplify centralized localization problem.

2) Hierarchically Distributed Monitoring: In hierarchically-distributed monitoring, the network topology is partitioned into monitoring domains, each of which is assigned a hierarchy level. We consider two hierarchical levels: Hierarchically Distributed Top Level (HDTL) and Hierarchically Distributed Detail Level (HDDL). Table II defines the monitor placement as well as the fault detection and localization capability for both models. Fig. 6 shows a simple two-level monitoring hierarchy. In accordance with our assumed fault-propagation model, we define a HDTL monitoring domain to consist of all nodes' input and output ports. Each HDDL domain consists of all discrete optical components within that OXC and all incoming edges to that OXC. This is shown in Fig. 6 by shaded regions. Every domain is assigned a local fault-manager responsible for computing the optimal set of activated monitors and performing the fault localization for all of its components. The HDTL manager receives only alarms from each node's port numbers, and is able to localize failures to node and edge granularity level. When it localizes failure to a particular monitoring domain, it can query the corresponding HDDL manager for detailed fault localization, or the HDDL manager could report it when it detects a fault. This approach can also be extended to multiple hierarchy levels.

TABLE II
MONITORING SCHEMES AND FAULT DETECTION AND LOCALIZATION CAPABILITY

Monitoring Scheme	Monitor Assignment	Fault Detection Capability	Fault Localization Capability
Flat Centralized Top Level (OXC-Edge) Monitoring	Active monitors on every OXC's input and output ports and all lightpath-terminating receivers report to central network manager	Within network: ▪ Any OXC ▪ Any directed-edge, directed-fiber-span, optical amplifier, or other component along any directed-edge ▪ Any bi-directed-edge, bi-directed-fiber-span, or other component along any bi-directed edge	Within network: ▪ Any OXC ▪ Any directed-edge between two OXCs ▪ Any bi-directed-edge between two OXCs
Flat Centralized Detail Level (Optical Component) Monitoring	All active monitors in network report to central network manager	Within network: Any transmitter, receiver, directed optical fiber, bi-directed optical fiber span, optical wavelength switch fabric, multiplexer, de-multiplexer, optical amplifier	Within network: Any transmitter, receiver, directed optical fiber, bi-directed optical fiber span, optical wavelength switch fabric, multiplexer, de-multiplexer, optical amplifier
Hierarchically Distributed Top Level (OXC-Edge) Monitoring	Active monitors on domain's OXC's input and output ports, and domain's lightpath-terminating receivers report to local domain manager	Within monitoring domain: ▪ Domain's OXC ▪ Any directed-fiber-span, optical amplifier, or other component along any incident edge ▪ Any bi-directed-fiber-span, or other component along a bi-directed edge whose one edge is domain's incident edge	Within monitoring domain: ▪ Domain's OXC ▪ Any incident directed-edge between two OXCs ▪ Any bi-directed-edge between two OXCs*
Hierarchically Distributed Detail Level (Optical Component) Monitoring	All active monitors within domain's OXC and along all incident edges report to a local manager	Within monitoring domain: Any transmitter, receiver, directed optical fiber, bi-directed optical fiber span, optical wavelength switch fabric, multiplexer, de-multiplexer, optical amplifier	Within monitoring domain: Any transmitter, receiver, directed optical fiber, bi-directed optical fiber span*, optical wavelength switch fabric, multiplexer, de-multiplexer, optical amplifier
* Hierarchical bi-directed edge localization requires two domain managers, each localizing failure on its corresponding incident-edge of that bi-directed edge.			

3) Fault-Coverage Completeness: In order to be able to compare performance results of the two localization schemes, we next show that the localization coverage of the two approaches is identical; i.e. a fault that is localizable in one approach is also so in the other, and vice versa. The hierarchical monitoring scheme creates a full cover over the network topology using a set of disjoint monitoring domains (except for monitors located on the domain's borders, which report to the local and central fault-managers) as shown in Fig. 6. Furthermore, each domain manager monitors failures within the node and on all incoming edges. All other components are located within the monitoring domain and their failures are handled by the corresponding manager.

Proof: First, let us assume the existence of a single fault $f(i)$ that can be detected by the centralized-flat monitoring scheme, but not by the 2-level hierarchically-distributed monitoring scheme. If a fault $f(i)$ is detected by the central-flat monitoring then at least one lightpath is affected by such fault. All monitors downstream from the location of a fault $f(i)$ along all affected lightpaths will detect this fault. Since every lightpath passes through at least one network node (monitoring-domain), any single failure within such a domain will also be detected by the monitors that are logically assigned to the second-level-domain (along the domain-border). Therefore, the 2nd-level manager is able to detect any fault detectable by the central-flat manager, implying that it is impossible for fault $f(i)$ to exist. This is a contradiction.

Next, let us assume the existence of a single fault $f(j)$ that can be detected by the 2-level-hierarchically distributed monitoring, but not by the centralized-flat monitoring. If a fault $f(j)$ is detected by the 2-level-hierarchical monitoring, then at least one lightpath is affected. In this case, the 2nd-level manager localizes fault $f(j)$ to its domain granularity, and the local manager localizes the fault to a specific component within this domain. Since in central-flat monitoring the central manager receives alarms from all monitors in the network it will also detect fault $f(j)$. Therefore, the central-flat manager is able to detect any fault detectable by the 2nd-level manager, implying that it is impossible for fault $f(j)$ to exist. This is again a contradiction, and we have proved the following:

Theorem 1: **In TONs with downstream fault propagation and monitors at every port of each component, two-level hierarchically-distributed monitoring achieves identical fault detection as centralized-flat monitoring.**

III. A SIMPLE EXAMPLE OF OPTIMAL MONITOR ACTIVATION

For this example we consider a simple network topology with 8 optical nodes (OXCs) and 14 bidirectional edges as shown in Fig. 7. The number of bidirectional ports on each node is determined by the degree of each node. Furthermore, a common assumption is made that every node has monitoring capabilities at each input/output port and that each node has two local transponder monitored ports (each with receiver and transmitter monitor) for sourcing or sinking a wavelength. With this configuration the total alarm vector length is 88 bits where each bit represents a unique monitor in the network. The problem is to minimize alarm vector length while maintaining

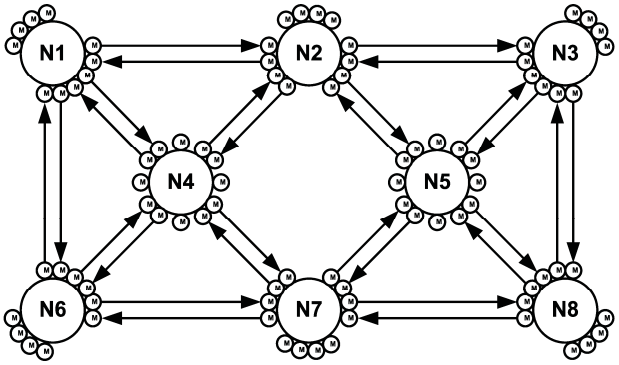

Fig. 7. A simple example network topology with port monitors.

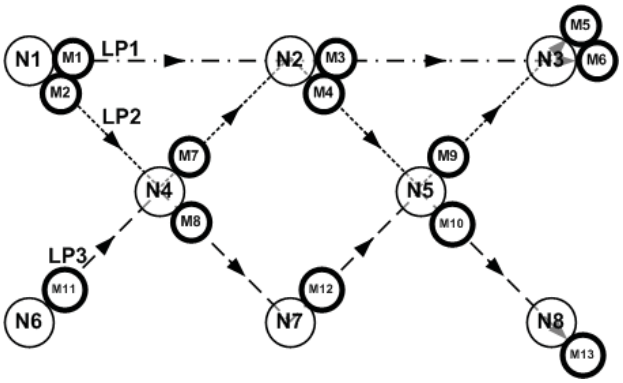

Fig. 8. Reduced topology graph with naively activated monitors.

full fault-localization coverage. We initially introduced this approach in [9].

In this example we consider only the localization of OXC faults for simplicity, whereas throughout the rest of this paper failures of all OXCs, fiber-spans, and optical amplifiers (at the top level) and all discrete optical components (at the detail level) are considered as uniquely localizable faults. For the topology considered in Fig. 7, let us assume that the following three lightpaths are provisioned:

LP1: Node1 $\rightarrow$ Node2 $\rightarrow$ Node3,

LP2: Node1 $\rightarrow$ Node4 $\rightarrow$ Node2 $\rightarrow$ Node5 $\rightarrow$ Node3, and

LP3: Node6 $\rightarrow$ Node4 $\rightarrow$ Node7 $\rightarrow$ Node5 $\rightarrow$ Node8.

Initially, we start solving the monitor activation problem by determining the naïve monitor activation scenario. This approach activates all downstream monitors on considered components (only OXCs in this example) that are traversed by at least one lightpath. By applying such monitor activation to this scenario, the set of all available monitors $\{M_1, \ldots, M_{88}\}$ is reduced to a subset containing 13 monitors. We label this set of activated monitors as $\{M_1, \cdots, M_{13}\}$, with graph placement as shown in Fig. 8.

Using the information on established lightpaths, network topology, and the active monitor locations, we create the alarm matrix as shown in Table III. For example, when node 5 fails (FN5) in Fig. 8, all activated downstream monitors on all lightpths that traverse node 5 $\{M_5, M_9, M_{10}, M_{13}\}$ report the alarm as binary 1; as shown in Table III.

Next, the preprocessing phase can be used to remove redundant information from the above alarm matrix by removing all zero-vector rows and grouping all identical rows into a single alarm vector. In this example, none of the fault vectors is zero, and all node faults can be uniquely identified; thus, no such reduction is necessary. Finally, the optimal set of active monitors can be obtained by removing all redundant monitors under the following constraints: (1) None of the matrix rows can be zero vectors. (2) All matrix rows must remain distinct. By applying these constraints while minimizing the number of monitors, the above matrix is reduced to the optimal alarm matrix shown in Table IV.

TABLE III
ALARM MATRIX

Faults	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	M13
FN1	1	1	1	1	1	1	1	0	1	0	0	0	0
FN2	0	0	1	1	1	1	0	0	1	0	0	0	0
FN3	0	0	0	0	1	1	0	0	0	0	0	0	0
FN4	0	0	0	1	1	0	1	1	1	1	0	1	1
FN5	0	0	0	0	1	0	0	0	1	1	0	0	1
FN6	0	0	0	0	0	0	0	1	0	1	1	1	1
FN7	0	0	0	0	0	0	0	0	0	1	0	1	1
FN8	0	0	0	0	0	0	0	0	0	0	0	0	1

TABLE IV
OPTIMAL ALARM MATRIX

Faults	M3	M5	M7	M8	M10	M13
FN1	1	1	1	0	0	0
FN2	1	1	0	0	0	0
FN3	0	1	0	0	0	0
FN4	0	1	1	1	1	1
FN5	0	1	0	0	1	1
FN6	0	0	0	1	1	1
FN7	0	0	0	0	1	1
FN8	0	0	0	0	0	1

This procedure has deactivated a redundant set of monitors $\{M_1, M_2, M_4, M_6, M_9, M_{11}, M_{12}\}$ and thus reduced alarm-vector length from 88 original monitors to only 6 monitors while maintaining full localization capability. It has achieved over 93% reduction in the total number of monitors (54% reduction if only monitors on used-output-ports are considered), thus greatly reducing the complexity of the fault-localization problem. The resulting network placement of the optimal set of monitors $\{M_3, M_5, M_7, M_8, M_{10}, M_{13}\}$ is shown in Fig. 9.

Once the optimal alarm matrix is computed as shown, optical component failure is located by employing an efficient fault-vector search within the optimized set of fault-vectors. A fast link-restoration protocol [28] can then be used to restore the traffic on all affected lightpaths.

IV. COMPLEXITY ANALYSIS

Once an alarm matrix is introduced (e.g., Table III), finding an optimal alarm matrix (e.g., Table IV) is equivalent to finding a set of redundant monitors of maximum cardinality whose deactivation satisfies the following properties:

1. The maximum fault detection coverage is still achieved, and
2. Each fault is reported by a unique set of alarms.

We call this problem the Redundant Monitor Deactivation Problem, which is formally stated below.

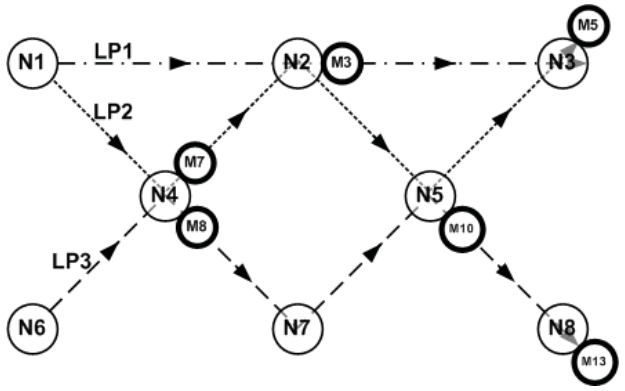

Fig. 9. Optimal monitor activation for the example.

Redundant Monitor Deactivation Problem (RMDP):

Given an integer m, a set of k monitors $M = \{M_1, ..., M_k\}$ and a set of n faults $F = \{f_1, \cdots, f_n\}$ such that $A_M(f_i) \neq A_M(f_j)$ for any $1 \leq i \neq j \leq n$, where $A_M(f_i)$ denotes the set of monitors in M that raise alarms when fault f_i occurs, the problem is to find a subset $D \subseteq M$, where D is the set of redundant monitors to be deactivated, such that:

(A1) $A_{M\text{-}D}(f_i) \neq A_{M\text{-}D}(f_j)$ for any $1 \leq i \neq j \leq n$

(A2) $A_{M\text{-}D}(f_i) \neq \mathbf{0}$ for any i, $1 \leq i \leq n$ and

(A3) $|D|$ is maximum.

In what follows, we show the NP-completeness of this problem by sketching a polynomial time transformation to the decision version of the RMDP (i.e., decide whether there exists a subset $D \subseteq M$ such that $|D| = m$ for a given integer m) from the Exact Three Cover (X3C) problem which is known to be NP-complete [30]. The X3C problem is given below.

X3C Problem:

Given a set of n elements $X = \{x_1, \cdots, x_n\}$ (where n is a multiple of three), and a set of m elements $C = \{c_1, \cdots, c_m\}$ where $c_j \subset X$ and $|c_j| = 3$ for all j, $1 \leq j \leq m$, the objective is to find a subset $C_0 \subset C$ such that $\cup_{c_j \in C_0} c_j = X$ and $|C_0| = n/3$.

A. Polynomial Time Transformation

Let X and C be an arbitrary input to the X3C problem. Define $C(x_i) = \{c_j | x_i \in c_j\}$. From X and C, we construct a monitor set M, a fault set F, and a set of alarms $A_M(f_i)$ corresponding to each $f_i \in F$ as follows.

The monitor set is constructed as $M = \{c_1, \cdots, c_m\} \cup \{d_1, \cdots, d_n\}$. The fault set is constructed as $F = \{x_1, \cdots, x_n\} \cup \{g_1, \cdots, g_n\}$. The alarm set is constructed as $A_M(x_i) = C(x_i) \cup \{d_i\}$ for $1 \leq i \leq n$ and $A_M(g_i) = \{d_i\}$ for $1 \leq i \leq n$.

Consider the following example as an input to the X3C problem: $X = \{x_1, \cdots, x_6\}$ and $C = \{c_1, c_2, c_3\}$, where $c_1 = \{x_1, x_2, x_3\}$, $c_2 = \{x_2, x_3, x_5\}$, $c_3 = \{x_4, x_5, x_6\}$. The input to the RMDP is then constructed as: $M = \{c_1, \cdots, c_3, d_1, \cdots, d_6\}$, $F = \{x_1, \cdots, x_6, g_1, \cdots, g_6\}$ and $A_M(x_1) = \{c_1, d_1\}$, $A_M(x_2) = \{c_1, c_2, d_2\}$, $A_M(x_3) = \{c_1, c_2, d_3\}$, $A_M(x_4) = \{c_3, d_4\}$, $A_M(x_5) = \{c_2, c_3, d_5\}$, $A_M(x_6) = \{c_3, d_6\}$, $A_M(g_1) = \{d_1\}$, $A_M(g_2) = \{d_2\}$, $A_M(g_3) = \{d_3\}$, $A_M(g_4) = \{d_4\}$, $A_M(g_5) = \{d_5\}$, and $A_M(g_6) = \{d_6\}$.

B. Correctness of Transformation

Lemma 1: Let X and C be an input to the X3C problem, and let $I_{X,C}$ instance of the input (i.e., M, F, and $A_M(F)$) to the RMDP constructed by following the above procedure. Then, there exists a solution to the X3C problem if and only if there exists a solution $D \subseteq M$ to $I_{X,C}$ such that $|D| = m - n/3$ satisfying the constraints (A1) and (A2) defined in the RMDP.

Proof: Suppose there exists a solution C_0 to the X3C problem. (Note that if C_0 is a exact cover for X, then $|C_0| = n/3$.) We then define $D = C - \mathbf{C}_0$, i.e., monitors corresponding to elements in $C - C_0$ will be deactivated. It is then easy to see that each fault x_i $(1 \leq i \leq n)$ will have a unique alarm set $A_{M-D}(x_i) = \{d_i, c_j\}$, where $x_i \in c_j$ for some $c_j \in C_0$. Each fault g_i $(1 \leq i \leq n)$ will also have a unique alarm set $\{d_i\}$ of size one. (In the example above, $D = \{c_2\}$ and $C_0 = \{c_1, c_3\}$.)

Now, assume that there exists a solution D for $I_{X,C}$ where $|D| = m - n/3$. We note that $d_i \notin D$ as otherwise fault g_i cannot be detected for $1 \leq i \leq n$. We also note that for each x_i, at least one monitor in $C - D$ should remain activated to monitor x_i since otherwise x_i and g_i both will only be monitored by the unique alarm d_i. This implies that $C - D$ must be a solution to the X3C problem.

We thus have the following result:

Theorem 2: **The RMDP is NP-complete.**

V. Integer Linear Program (ILP) Formulation

It was shown in Section IV that the problem of optimally activating monitors is NP-complete. In this section, we present an ILP formulation that minimizes the number of monitors that are activated, while maintaining the full fault-localization capability. Given the requirements on the speed of the monitor activation in dynamic traffic settings and the computational intensity of solving ILPs of even moderate sizes, we do not suggest solving the ILP as a monitor activation algorithm. Our goal in introducing the ILP here is to make optimization performance comparisons with a simple and fast heuristic that *can* be used. The input to the problem is an $n \times k$ alarm matrix A which indicates the set of monitors that generates alarms in the event of each failure. Let $\{a_{f,m}\}$ represent the elements of the alarm matrix A, where $a_{f,m}$ takes the value 1 if monitor m generates an alarm when fault f occurs, and 0 otherwise. Since $M = \{M_1, ..., M_k\}$, $|M| = k$. For simplicity, we assume that A is the reduced alarm matrix obtained at the end of the preprocessing phase of our solution approach described in Section III, and thus all the rows are non-zero, and identical rows are eliminated. The decision variables are $\{u_m\}$, where u_m takes the value 1 if monitor m is activated, and the value 0 otherwise. The ILP attempts to minimize the number of monitors that are activated such that all constraints are satisfied. The optimization problem can be described by the following ILP:

The objective of the ILP is to minimize the number of activated monitors:

$$\text{Minimize} \quad Z = \sum_{m=1}^{k} u_m$$

Subject to the following three constraints:

Input: An *nxk* matrix *A*
// Where element of *A*, $a_{f,m} = 1$ if fault *f* causes alarm *m* //

Output: An *nxk'* matrix *A'* such that $\lceil \log n \rceil \le k' \le k$

1. **while:** There is unmarked column **Do**
2. find unmarked column with smallest sum;
3. **if** column deletion is allowed
4. delete the column;
5. **else** mark this column;
6. **endwhile**

Algorithm 1. Greedy Min Heuristic algorithm.

1. Uniqueness constraints: Each fault-vector, consisting of the elements corresponding to the set of activated monitors, should be different from all other fault-vectors. Thus, for every pair of rows corresponding to failures f_i and f_j , where $i \neq j$ we have,

$$\sum_{m=1}^{k} \left(a_{f_i,m} - a_{f_j,m}\right)^2 u_m > 0. \quad (1)$$

Note that $a_{f,m}$ are parameters and not decision variables, and the problem is not non-linear as (1) might appear to suggest.

2. Non-zero constraints: For each failure scenario, there should be at least one activated monitor that generates an alarm in the event of that failure to ensure that the failure can be detected. Thus, each fault-vector should contain at least one non-zero element in the columns corresponding to the activated monitors. For each f_i,

$$\sum_{m=1}^{k} a_{f_i,m} u_m > 0. \quad (2)$$

3. Binary Decision Variable Constraints: All decision variables $\{u_m\}$ should be binary, since each monitor's reports are either turned on or off.

$$u_m \in \{0, 1\} \quad \forall m.$$

VI. Efficient Heuristic for Monitor Activation

In Section IV we proved that the optimal monitor activation problem is NP-complete. Thus, in practice, computing the optimal set of activated monitors for a dynamically provisioned set of lightpaths might not be feasible within the allotted time limits. Accordingly, here we present a heuristic algorithm that is far more applicable as it provides near-optimal solutions with much smaller time complexity.

The heuristic algorithm is based on a simple greedy method. It starts by selecting a monitor column in the alarm-matrix with a minimum sum and deletes it if the following two fault-vector constraints are not violated: 1. All rows remain non-zero, 2. All rows remain distinct. The reason for selecting the minimum sum column first is because such a column is least effective in splitting a fault set into two equal fault sets (assuming a binary tree splitting of a fault set) and yields the best results when compared to other orders of column removal for this problem. The algorithm proceeds to test and delete if possible all untested columns in order of increasing column sum until all columns are tried. Pseudo code for this heuristic is presented in Algorithm 1. Applying this heuristic to the simplified alarm matrix in Table III yields the optimal solution of 6 monitors for this specific case as shown in Table V.

TABLE V
The Heuristic Alarm Matrix

Faults	M5	M7	M8	M9	M10	M13
FN1	1	1	0	1	0	0
FN2	1	0	0	1	0	0
FN3	1	0	0	0	0	0
FN4	1	1	1	1	1	1
FN5	1	0	0	1	1	1
FN6	0	0	1	0	1	1
FN7	0	0	0	0	1	1
FN8	0	0	0	0	0	1

A. Time Complexity

The time complexity consists of two components: (1) time required to optimize the alarm matrix using the proposed heuristic, which is computed only when the set of lightpaths changes, and (2) time required to locate received fault vector through matching with the corresponding row within the reduced alarm matrix.

For the first component, it can be easily verified that the worst-case time complexity of the heuristic (given the alarm matrix) is $O\left(n^2 k\right)$ where n is the number of unique fault-vectors and k is the number of monitors in the preprocessed (reduced) alarm-matrix.

In the hierarchically-distributed model, each domain manager solves a monitor optimization sub-problem for the subset of monitors located within its domain, thus solving only a fraction of the overall problem. All domain-managers compute this heuristic independently and in parallel, thus complexity is bounded by a domain with the largest number of possible faults. Assuming uniform distribution of possible faults across d domains, per-domain complexity can be approximated as $O\left(n^2 k/d^3\right)$. Furthermore, only domains affected by lightpath reconfiguration need to re-compute their alarm vectors.

The complexity of the fault localization is given by the time required to match the received fault-vector with one of the vectors in alarm matrix. This can be easily done using one of the efficient search algorithms such as a binary-tree search or a lookup table. Assuming that the number of monitors in optimized alarm-matrix is given by m, this can be done in worst-case $O\left(m\right)$ time. Note that the size of m per domain will be a fraction of the m in centralized monitoring.

VII. Numerical Results

Experimental results were obtained through simulation consisting of flat and hierarchically-distributed network monitoring models, a fault propagation model and alarm matrix computation code, heuristic algorithm, and ILP implementation. ILOG CPLEX [31] was used for the ILP computation and everything else was implemented in C code. All simulation results were computed on a server with two quad-core Intel Xeon 2.5 GHz E5420 CPUs.

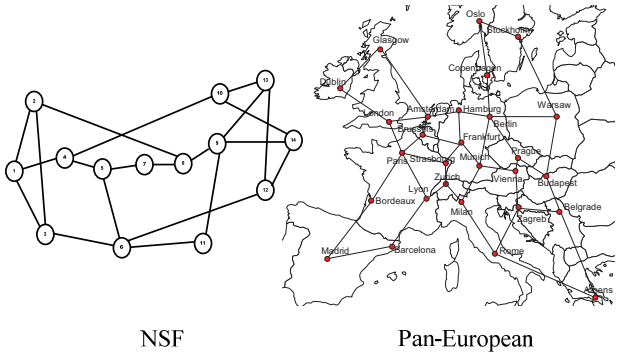

Fig. 10. NSF [32] and Pan-European [33] network topologies.

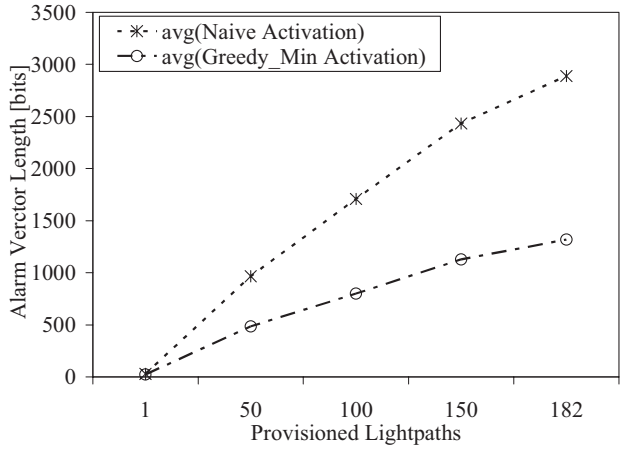

Fig. 11. NSF network centralized-flat discrete-component monitoring alarm-vector length.

Results were obtained for the two common network topologies. The NSF network (Fig. 10) [32] with 14 nodes, 42 directed edges, and 182 ordered node-pairs, and the Pan-European network (Fig. 10) [33] with 28 nodes, 82 directed edges and 756 ordered node-pairs. All provisioned lightpaths were computed using the shortest-path routing between uniformly distributed node pairs, using the actual physical distances as weights. Wavelength assignment for every requested lightpath was performed using the first-fit algorithm without wavelength conversions. The top-level centralized monitoring model for the NSF network consists of 14 OXCs, 42 directed edges, and 21 bi-directed edges, resulting in the total of 77 possible element faults (EF). Similarly, Pan-European top-level centralized monitoring has a maximum of 151 possible EF.

In the detail-component monitoring model, all discrete optical components within an OXC and along each directed edge were considered as possible EF. Each OXC model has a maximum of 5 input and output ports (maximum node-degree for both topologies is 5), and can support up to 120 wavelengths without blocking. The maximum possible EF that can occur within each OXC are as follows: 5 multiplexers, 5 de-multiplexers, (120^*5^*2) interconnecting fiber spans, 120 wavelength switching fabrics, (5^*120) transmitters, and (5^*120) receivers; resulting in the total maximum of 2530 EF per OXC. Every discrete optical element is monitored by power monitors that are available on all input and output ports of that element.

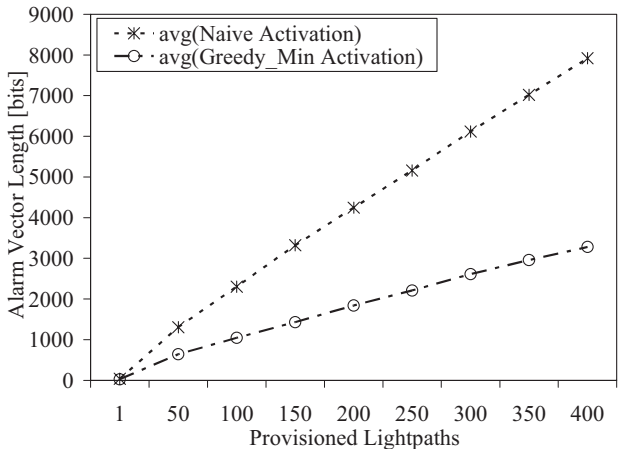

Fig. 12. Pan-European network centralized-flat discrete-component monitoring alarm-vector length.

Each directed edge consists of multiple fiber spans and discrete optical amplifiers. Optical amplifiers are placed along each directed edge every 75km. Actual physical link distances were scaled down by a factor of 10 due to the large number of required optical amplifications and redundant monitors along the full length of the links. The total number of optical amplifiers, fiber spans, and required monitors along each directed edge are computed according to the scaled physical link distances.

A. Centralized-Flat Monitoring Results

Initially, the simulation was run with a single central network manager monitoring all discrete optical components in the network to estimate the complexity of the centralized component-monitoring problem. Results for the NSF topology are shown in Fig. 11, and the results for the Pan-European topology are shown in Fig. 12. Results compare alarm vector lengths of naïve monitor activation and the effectiveness of heuristic alarm-vector reduction. Solving the ILP with CPLEX for optimal activation in a large centralized monitoring scheme was too complex even for a very small number of lightpaths. However, optimal results for all other monitoring models in the next section demonstrate that the heuristic provides a nearly-optimal solution in terms of the alarm-vector length.

For both topologies, heuristic reduction increases with the increasing number of provisioned lightpaths due to the effect that the fault propagation model has on the larger number of lightpaths.

It is interesting to note that for the NSF network with 182 lightpaths, the heuristic deactivates 1,800 redundant monitors, while for the Pan-European network with 400 lightpaths the heuristic deactivates 5,000 redundant monitors while maintaining full fault-localization capability in both scenarios. Such a large reduction of fault-vector size greatly simplifies the fault localization problem.

The discrete optical-component monitoring coverage for each topology is shown in Table VI. The results represent the average number of uniquely localizable network element faults averaged over all simulation runs, each of which considers a

TABLE VI
CENTRALIZED-FLAT DISCRETE-COMPONENT MONITORING ULF COVERAGE

LPs	NSF		Pan Euro	
	EF	ULF	EF	ULF
1	22.9	22.9	27.9	27.9
50	732.2	732.2	1000.7	1000.7
100	1256.6	1256.6	1711.6	1711.6
150	1760.4	1760.4	2390.2	2390.2
182	2074.7	2074.7	2824.6	2824.6
200			3069.3	3069.3
250			3710.1	3710.1
300			4362.0	4362.0
350			4992.8	4992.8
400			5690.0	5690.0

TABLE VII
CENTRALIZED-FLAT DISCRETE-COMPONENT MONITORING ALARM VECTOR OPTIMIZATION TIME COMPLEXITY

LPs	NSF Time [sec]		Pan Euro Time [sec]	
	Avg Greedy	Max Greedy	Avg Greedy	Max Greedy
1	0.0001	0.001	0.0015	0.015
50	40.3	47.2	103.3	112.5
100	268.2	295.8	702.5	760.0
150	887.6	924.9	1900.1	2160.1
182	1625.1	1656.8	4391.2	4782.6
200			5799.6	6268.2
250			13100.1	14210.5
300			21865.3	24337.9
350			31600.7	34780.1
400			41000.3	45590.8

specified number of randomly provisioned shortest-path lightpaths. It is clear that the number of uniquely localizable faults (ULF) is always identical to the corresponding number of EFs implying the 100% localization capability of the proposed scheme. It is also clear that the resulting alarm-vector lengths can be very large (over 5,600 monitors) resulting in a large computation requirement for centralized fault-localization.

Table VII presents heuristic reduction time complexity for both networks. The large size of the original alarm matrix will slow down centralized alarm-vector optimization, while the large alarm-vector length will complicate fault localization every time a fault occurs.

It is possible to compute the fault-vector reduction problem, using the proposed heuristic, within a reasonable time offline for static or manually provisioned networks. However, for emerging dynamically provisioned networks, this problem must be solved as new lightpaths are provisioned, and accurate fault localization must be performed quickly whenever fault occurs.

Although proposed heuristic achieved a large reduction of fault vector length and greatly improved fault-localization speed, centralized component-level monitoring in large networks is a very complex problem for a single manager monitoring scheme, and is usually handled by multiple management entities in existing optical networks. Thus the need for hierarchically-distributed monitoring approach whose results we present next.

B. Hierarchically-Distributed Monitoring Results

The following results were obtained by imposing a two-level monitoring hierarchy on the discrete component-level

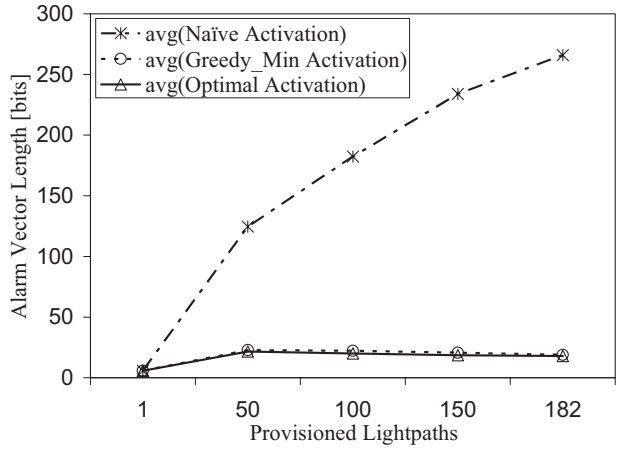

Fig. 13. NSF network hierarchical centralized top-level-manager node-edge monitoring alarm-vector length.

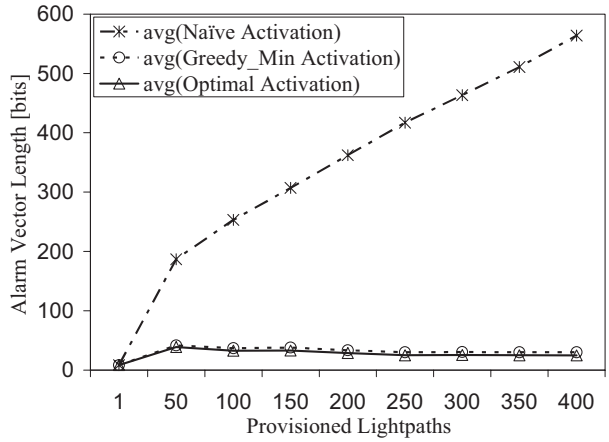

Fig. 14. Pan-European network hierarchical centralized top-level-manager node-edge monitoring alarm-vector length.

NSF and Pan-European topologies as described in Section II. The top level centralized monitoring alarm matrices were computed for node-edge localization granularity. The domain level alarm matrices were computed for each monitoring domain separately at discrete optical component localization granularity for all iterations. All resulting alarm matrices were then reduced using the heuristic algorithm and optimized using the ILP. Results for domain-level monitoring were then averaged over all iterations and all domains in the network for a given number of lightpaths.

The top-level (node-edge) centralized monitoring results for the NSF and Pan-European topology are shown in Fig. 13. and Fig. 14. respectively. It is apparent that for the NSF network with 182 lightpaths, the heuristic deactivates over 250 redundant monitors, while for the Pan-European network with 400 lightpaths the heuristic deactivates over 550 redundant monitors while maintaining full fault-localization capability in both scenarios. It is also clear that in both cases the heuristic provides results that are very close to the optimal solution.

The centralized top-level (node-edge) monitoring localization coverage for both topologies is provided in Table VIII. For both networks the full localization coverage at node-

TABLE VIII
HIERARCHICAL CENTRALIZED TOP-LEVEL-MANAGER NODE-EDGE MONITORING ULF COVERAGE

LPs	NSF		Pan Euro	
	EF	ULF	EF	ULF
1	5.8	5.8	8.4	8.4
50	68.0	68.0	126.4	126.4
100	75.5	75.5	141.0	141.0
150	77.0	77.0	145.7	145.7
182	77.0	77.0	147.8	147.8
200			149.4	149.4
250			150.2	150.2
300			150.4	150.4
350			150.6	150.6
400			151.0	151.0

TABLE IX
HIERARCHICAL CENTRALIZED TOP-LEVEL-MANAGER NODE-EDGE MONITORING ALARM VECTOR OPTIMIZATION TIME COMPLEXITY

LPs	NSF Time [sec]		Pan Euro Time [sec]	
	Avg Greedy	Max Greedy	Avg Greedy	Max Greedy
1	0.00001	0.00002	0.00001	0.00002
50	0.0132	0.017	0.0608	0.094
100	0.0211	0.026	0.1073	0.125
150	0.0242	0.027	0.1326	0.1496
182	0.0250	0.025	0.1425	0.1632
200			0.1504	0.172
250			0.1621	0.1883
300			0.1732	0.203
350			0.1834	0.2112
400			0.1934	0.218

edge granularity is maintained throughout all results while the number of EFs is significantly smaller implying much faster fault localization.

Table IX presents heuristic reduction time complexity for both networks. It is clear that computation time never exceeds 0.027sec for the NSF and 0.218sec. for the Pan-European topology, making it efficient enough for use in dynamically provisioned networks, since the lightpath provisioning time is currently much longer than this.

Figs. 15 and 16 show the results for distributed domain–level (discrete optical component level) monitoring. Similarly to top-level monitoring, the fault-vector reduction of 100 monitors for the NSF topology and 200 monitors for the Pan-European topology is achieved. In both cases, the heuristic also provides nearly optimal results. Also the small alarm-vector size implies fast localization of any fault within the monitoring domain.

The hierarchically-distributed localization coverage for both topologies is provided in Table X and shows that a full localization capability is maintained throughout and a significantly smaller number of faults implies a much smaller alarm matrix and much faster fault detection, localization, and restoration than in centralized-component monitoring. Finally, Table XI shows time complexity for domain level fault-vector reduction. It is clear that the average time is upper-bounded by 0.077sec and 0.388sec for NSF and Pan-European networks respectively, making this hierarchically-distributed monitoring scheme suitable for dynamically-provisioned networks, given that it only needs to be computed when network topology or set of lightpaths change, which takes much longer in current optical networks. Furthermore, when a fault occurs within

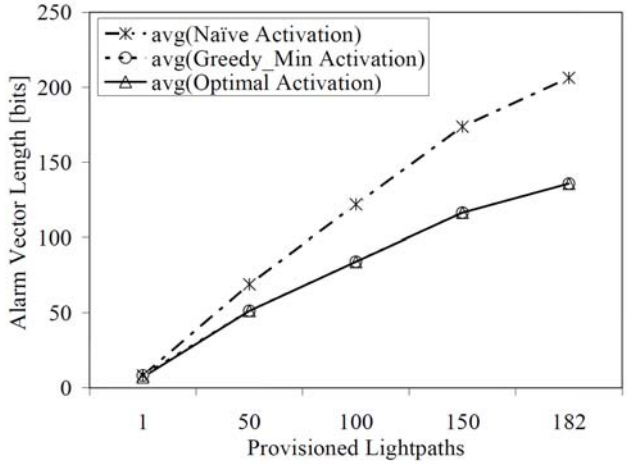

Fig. 15. NSF network hierarchically distributed domain-manager discrete component monitoring alarm-vector length.

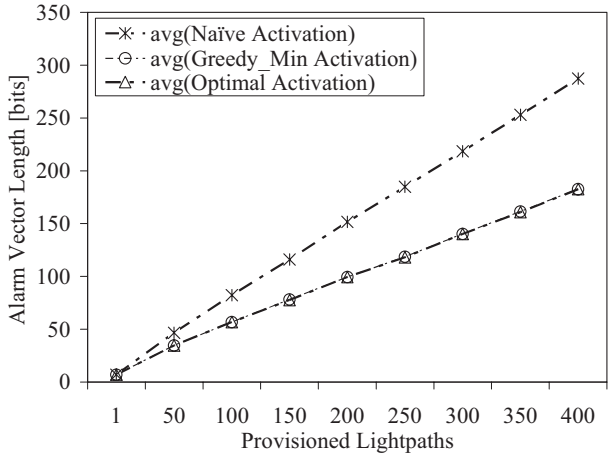

Fig. 16. Pan-European network hierarchically distributed domain-manager discrete component monitoring alarm-vector length.

a specific domain, fault-localization (vector search) is much faster given the small subset of possible faults and monitors within each domain. Another advantage of hierarchically-distributed monitoring is the limited impact of lightpath reconfiguration on re-computation of fault vectors. When a lightpath is added or removed, only those domains affected by such change need to re-compute their fault vectors, thus greatly reducing the overall impact on computation and the average frequency of domain updates.

Next, we apply the proposed approach to obtain the required number of monitors for localization of all possible bi-directed edge failures in common network topologies shown in Fig. 17. Only traffic lightpaths are used, eliminating the overhead of supervisory lightpaths. Randomly generated directed-shortest-lightpaths are added until all bi-directed-edge failures are uniquely localizable. Corresponding alarm matrices are computed and optimized with the proposed heuristic and optimal ILP solution while maintaining full localization capability.

Results in Table XII, suggest that even in the case of full fault-localization coverage using only already existing traffic lightpaths, the number of required monitors (provided by the proposed heuristic or the optimal solution) is small enough

TABLE X
HIERARCHICALLY DISTRIBUTED DOMAIN-MANAGER DISCRETE-COMPONENT MONITORING ULF COVERAGE

LPs	NSF		Pan Euro	
	EF	ULF	EF	ULF
1	8.3	8.3	7.0	7.0
50	57.1	57.1	38.9	38.9
100	94.7	94.7	64.4	64.4
150	130.7	130.7	88.4	88.4
182	153.1	153.1	103.8	103.8
200			113.0	113.0
250			135.8	135.8
300			159.1	159.1
350			182.9	182.9
400			206.9	206.9

TABLE XI
HIERARCHICALLY DISTRIBUTED DOMAIN-MANAGER DISCRETE-COMPONENT MONITORING ALARM VECTOR OPTIMIZATION TIME COMPLEXITY

LPs	NSF Time [sec]		Pan Euro Time [sec]	
	Avg Greedy	Max Greedy	Avg Greedy	Max Greedy
1	0.000008	0.001	0.000001	0.000002
50	0.002	0.015	0.0007	0.016
100	0.014	0.059	0.0057	0.047
150	0.043	0.192	0.020	0.108
182	0.077	0.289	0.032	0.217
200			0.040	0.281
250			0.044	0.721
300			0.147	1.388
350			0.262	2.483
400			0.388	3.884

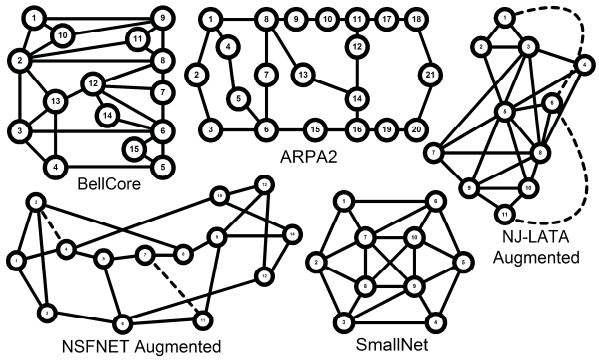

Fig. 17. Network topologies used in the simulations.

TABLE XII
REQUIRED MONITORS FOR LOCALIZATION OF BI-DIRECTED EDGE FAULTS IN COMMON TOPOLOGIES WITHOUT SUPERVISORY LIGHTPATHS AND ONLY RANDOM TRAFFIC LIGHTPATHS

Monitors	RND MIN	RND AVG	RND MAX	All-Pairs SPs	% Monitor Reduction
ARPA2					
Naïve	66	79.3	88	100	
GreedyMin	11	13.1	15	8	92.0%
ILP	10	12.3	14	7	93.0%
BellCore					
Naïve	69	82	88	92	
GreedyMin	14	15.6	18	14	84.8%
ILP	13	14.5	17	12	87.0%
NJ LATA Augmented					
Naïve	64	67	71	73	
GreedyMin	18	18.5	20	18	75.3%
ILP	17	17.4	18	17	76.7%
NSF					
Naïve	49	62.6	75	81	
GreedyMin	10	12.8	15	9	88.9%
ILP	10	12	15	8	90.1%
NSF Augmented					
Naïve	53	69.4	79	87	
GreedyMin	12	14	16	10	88.5%
ILP	11	13	16	9	89.7%
SmallNet					
Naïve	50	60.3	70	76	
GreedyMin	15	16.6	18	15	80.3%
ILP	15	16.6	18	15	80.3%
AVG % Monitor Reduction Heuristic					85.0%
AVG % Monitor Reduction Optimal (ILP)					86.1%

to provide fast and accurate fault localization of all failures. The second, third, and fourth columns in Table XII show the minimum, average, and maximum results respectively for the simulation scenario in which shortest-path (SP) traffic lightpaths are randomly (RND) added until full fault localization is achieved. The last two columns in the table show the results for the scenario when all-pairs shortest paths are added to the network and the resulting alarm matrices are computed and optimized. The large average (AVG) reduction of fault-vector size (85% avg. reduction with heuristic and 86.1% avg. reduction with optimal ILP) with the proposed heuristic provides almost optimal results in all scenarios.

Finally, we examine the impact of the network's connectivity on the fault localization complexity of the proposed scheme for a set of random network topologies. In this simulation, 10 random networks are generated for each specified node-degree. Node degrees range from 3 to $N-1$ where N is the number of nodes in the network. For each such random network, 10 different sets of random-shortest-path traffic lightpaths are generated by adding random-shortest-paths between uniformly distributed node-pairs until all faults in the network can be uniquely localized. Thus, 100 alarm matrices are generated for each node-degree of a given network size. Results were obtained for 10-node and 16-node random networks.

Results in Fig. 18 show the relation between the node degree and the alarm vector length for 10-node and 16-node networks. For both network sizes, alarm vector length increases almost linearly with the increasing network connectivity. Furthermore, there is a notable reduction of the alarm vector length that grows with increasing network size. This suggests the scheme's scalability in large networks. Also, the quality of the proposed heuristic results remains very close to the optimal ILP solutions (solutions directly overlap in figure) for all random network cases, suggesting its applicability for a wide range of network topologies. It is apparent that the gap between Naïve and Heuristic/Optimal solutions increases with increasing network size, implying the scalability of our approach with increasing network size.

Note that these results are based purely on utilizing already existing traffic lightpaths to perform fault-localization. Given the nature of existing DWDM core-networks, there is a high probability of having at least one traffic wavelength used on every fiber-span, thus yielding full fault-localization of all fiber-span failures. There is no additional overhead in terms of the supervisory-path computation time, provisioning-time, maintenance, and bandwidth-utilization that are associated with dedicated supervisory-path based schemes.

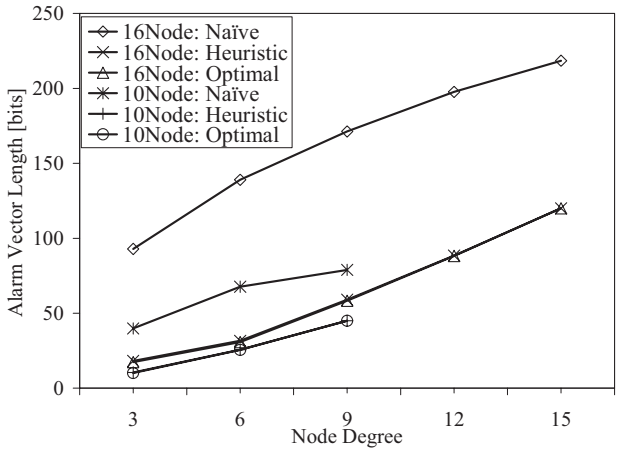

Fig. 18. Impact of network connectivity on fault localization complexity in random network topologies.

VIII. Conclusions and Future Work

In this work we addressed the important problem of optimizing the fault-localization in transparent all-optical networks through minimization of activated optical monitoring equipment while maintaining full localization coverage. The presented approach enables scalable and rapid monitor optimization and multiple granularity levels of fault localization through the creation of a monitoring hierarchy. It enables distributed optimization of alarm-vector length and fault-localization. We defined the construction of monitoring domains and proved that such hierarchically distributed fault localization achieves identical fault-localization as the centralized monitoring. We also proved that the minimum monitor activation problem is NP-hard, and introduced a novel heuristic algorithm with nearly optimal performance. The ILP formulation with CPLEX implementation provided an optimal reference for evaluating performance of the proposed heuristic. Extensive simulation was conducted to evaluate the performance of the proposed approach. We have shown that the proposed algorithm can provide high quality results in much shorter computation time than ILP. Combined with the proposed hierarchical monitoring, we have shown that it is applicable for use in emerging dynamically-provisioned TONs.

This work can be extended in several interesting ways. For example, including other types of monitors would allow detailed analysis of the optical layer. Different fault-propagation models could be considered to examine their impact on fault localization. Furthermore, although the proposed heuristic is fast enough to be recomputed when lightpaths change, a more adaptive heuristic could be developed to allow progressive optimization of the alarm-vectors as the lightpaths are reconfigured. Finally, the proposed approach can be easily extended to other areas of monitoring and sensor data reduction (assuming that appropriate propagation models are considered) including: sensor array networks, signal cross-talk monitoring, process monitoring, circuit design and testing, perimeter surveillance, and other related applications.

References

[1] M. Médard, D. Marquis, R. A. Barry, and S. G. Finn, "Security issues in all-optical networks," *IEEE Network Mag.*, vol. 11, no. 3, pp. 42–48, May/June 1997.

[2] C. Assi, Y. Ye, A. Shami, S. Dixit, and M. Ali, "A hybrid distributed fault-management protocol for combating single-fiber failures in mesh-based DWDM optical networks," in *Proc. IEEE GLOBECOM*, 2002, vol. 3, pp. 2676–2680.

[3] Y. Pointurier, M. Coates, and M. Rabbat, "Active monitoring of all-optical networks," in *Proc. IEEE ICTON*, 2008, pp. 169–172.

[4] Y. Hamazumi, M. Koga, K. Kawai, H. Ichino, and K. Sato, "Optical path fault management in layered networks," in *Proc. IEEE GLOBECOM*, 1998, pp. 2309–2314.

[5] T. Y. Chow, F. Chudak, and A. M. French, "Fast optical layer mesh protection using pre-cross-connected trails," *IEEE/ACM Trans. Networking*, vol. 12, no. 3, pp. 539–548, 2004.

[6] Y. Kobayashi, Y. Tada, S. Matsuoka, N. Hirayama, and K. Hagimoto, "Supervisory systems for all-optical network transmission systems," in *Proc. IEEE GLOBECOM*, 1996, pp. 933–937.

[7] R. D. Gardner and D. A. Harle, "Pattern discovery and specification techniques for alarm correlation," in *Proc. Network Operations Management Symposium*, 1998, pp. 713–722.

[8] S. Ramamurthy, L. Sahasrabuddhe, and B. Mukherjee, "Survivable WDM mesh networks," *IEEE J. Lightwave Technol.*, vol. 21, no. 4, pp. 870–883, 2003.

[9] S. Stanic, S. Subramaniam, H. Choi, G. Sahin, and H.-A. Choi, "On monitoring transparent optical networks," in *Proc. Int'l Conf. Parallel Processing Workshop*, pp. 217–223, Aug. 2002.

[10] S. Stanic, G. Sahin, H. Choi, S. Subramaniam, and H.-A. Choi, "Monitoring and alarm management in transparent optical networks," in *Proc. IEEE BROADNETS*, 2007.

[11] I. Katzela, G. Ellinas, W. S. Yoon, and T. E. Stern, "Fault diagnosis in optical networks," *J. High Speed Networks*, vol. 10, no. 4, pp. 269-91, 2001.

[12] C. Mas, I. Tomkos, and O. Tonguz, "Failure location algorithm for transparent optical networks," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 8, pp. 1508–1519, Aug. 2005.

[13] C. Mas and P. Thiran, "An efficient algorithm for locating soft and hard failures in WDM networks," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 10, pp. 1900–1911, 2000.

[14] H. Zeng, C. Huang, A. Vukovic, and M. Savoie, "Fault detection and path performance monitoring in meshed all-optical networks," in *Proc. IEEE GLOBECOM*, 2004.

[15] A. G. Hailemariam, G. Ellinas, and T. Stern, "Localized failure restoration in mesh optical networks," in *Proc. IEEE Conf. OFC*, 2004, pp. 23–27.

[16] A. V. Sichani and H. T. Mouftah, "Limited-perimeter vector matching fault-localization protocol for transparent all-optical communication networks," *IEEE IET Commun.*, pp. 472–478, 2007.

[17] S. S. Ahuja, S. Ramasubramanian, and M. Krunz, "SRLG failure localization in all-optical networks using monitoring cycles and paths," in *Proc. IEEE INFOCOM*, 2008.

[18] B. Wu, P.-H. Ho, and K. L. Yeung, "Monitoring trail: on fast link failure localization in WDM mesh networks," *IEEE/OSA J. Lightwave Technol.*, vol. 27, no. 23, Dec. 2009.

[19] J. Tapolcai, B. Wu, and P.-H. Ho, "On monitoring and failure localization in mesh all-optical networks," in *Proc. IEEE INFOCOM*, 2009.

[20] W. Tao and A. K. Somani, "Cross-talk attack monitoring and localization in all-optical networks," in *IEEE/ACM Trans. Networking*, Dec. 2005.

[21] N. Harvey, M. Patrascu, Y. Wen, S. Yekhanin, and V. W. S. Chan, "Non-adaptive fault diagnosis for all-optical networks via combinatorial group testing on graphs," in *Proc. IEEE INFOCOM*, 2007, pp. 697–705.

[22] C. Li, R. Ramaswami, I. Center, and Y. Heights, "Automatic fault detection, isolation, and recovery in transparent all-optical networks," *IEEE J. Lightwave Technol.*, vol. 15, no. 10, pp. 1784–1793, 1997.

[23] Y. Wen, V. Chan, and L. Zheng, "Efficient fault-diagnosis algorithms for all-optical WDM networks with probabilistic link failures," *IEEE J. Lightwave Technol.*, vol. 23, pp. 3358–3371, 2005.

[24] S. Stanic and S. Subramaniam, "Distributed hierarchical monitoring and alarm management in transparent optical networks," in *Proc. IEEE ICC*, 2008.

[25] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*. Morgan Kaufman, 1998.

[26] S. Stanic, S. Subramaniam, H. Choi, G. Sahin, and H.-A. Choi, "Efficient alarm management in optical networks," *DISCEX-III*, vol. I, pp. 252–260, Apr. 2003.

[27] A. Todimala and B. Ramamurthy, "Survivable virtual topology routing under shared risk link groups in WDM networks," in *Proc. IEEE BROADNETS*, 2004, pp. 130–139.
[28] M. Medard, S. G. Finn, R. A. Barry, W. He, and S. S. Lumetta, "Generalized loop-back recovery in optical mesh networks," *IEEE Trans. Networking*, pp. 153–164, 2002.
[29] H. T. Mouftah and P.-H. Ho, *Optical Networks: Architecture and Survivability*. Norwell, MA: Kluwer Academic Publisher Group, 2003.
[30] M. R. Garey and D. J. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: Freeman, 1979.
[31] ILOG optimization products, ILOG CPLEX. [Online]. Available: http://www.ilog.com/products/cplex/
[32] A. Betker, C. Gerlach, R. Hülsermann, M. Jäger, M. Barry, S. Bodamer, J. Späth, C. Gauger, and M. Köhn, "Reference transport network scenarios," technical report, BMBF Multi Tera Net, 2003.
[33] S. D. Maesschalck, D. Colle, I. Lievens, M. Pickavet, P. Demeester, C. Mauz, M. Jaeger, R. Inkret, B. Mikac, and J. Derkacz, "Pan-European optical transport networks: an availability-based comparison," *Photonic Network Commun.*, vol. 5, no. 3, 203–225, 2003.

Sava Stanic (S'98-M'01) received his B.Sc. in electrical engineering from the University of Maryland, College Park, MD, in 1998 and his M.Sc. in electrical and computer engineering from The George Washington University, Washington, DC, in 2000. He is currently a Ph.D. candidate in electrical and computer engineering at The George Washington University, and is a recipient of the ARCS (General Dynamics Corp. Scholar) and Phi Delta Gamma graduate scholarships for his Ph.D. research. His professional experience includes research, development, implementation, and deployment of mission-critical advanced broadband wireless, satellite, and optical communications systems for Lockheed Martin, DISA, JTRS, COMSAT Laboratories, Nextel Communications, Astrolink, and Booz|Allen| Hamilton. His research interests are in the areas of performance analysis, emerging architectures and algorithms for wireless ad-hoc, sensor, and optical networks, network monitoring and resilience, and distributed systems/algorithms.

Suresh Subramaniam (S'95-M'97-SM'07) received the Ph.D. degree in electrical engineering from the University of Washington, Seattle, in 1997. He is a Professor in the Department of Electrical and Computer Engineering at The George Washington University, Washington, DC. His research interests are in the architectural, algorithmic, and performance aspects of communication networks, with particular emphasis on optical and wireless ad hoc networks. Dr. Subramaniam is a co-editor of the books *Optical WDM Networks - Principles and Practice* and *Emerging Optical Network Technologies: Architectures, Protocols, and Performance*. He has been on the program committees of several conferences including Infocom, ICC, Globecom, OFC, and Broadnets, and served as TPC Co-Chair for the optical networks symposia at Globecom 2006 and ICC 2007. He currently serves on the editorial boards of the IEEE/ACM TRANSACTIONS ON NETWORKING, *Optical Switching and Networking*, and the *KICS Journal of Communications and Networks*. He is a co-recipient of Best Paper Awards at the ICC 2006 Symposium on Optical Systems and Networks, and at the 1997 SPIE Conference on All-Optical Communication Systems.

Gokhan Sahin (S'96-M'02) is an associate professor in the Department of Electrical and Computer Engineering at Miami University. He received his master's degree in 1997 and his Ph.D. in 2001, from the University of Washington, Seattle, and his bachelor's degree from Bilkent University, Turkey in 1996, all in electrical engineering. Dr. Sahin joined Miami in 2004 as one of the inaugural faculty members of the ECE department. Prior to that, he held research positions at AT&T Labs, Corvis Corporation, the The George Washington University, and the University of Nebraska. His research interests are in the design and analysis of computer and communication networks, with an emphasis on IP/optical and wireless networks. Dr. Sahin is a co-inventor of AT&T's ROLEX (Robust Optical Layer End-to-End X-connection) restoration method, currently utilized in DARPA CORONET, a major national initiative aiming to design the next-generation Internet. He has co-authored over 35 peer-reviewed articles, and serves on the editorial board of the IETE TECHNICAL REVIEW.

Hongsik Choi (M'97-SM'05) received the D.Sc. degree in computer science from The George Washington University, Washington, DC, in 1997. He is an Assistant Professor in the Department of Computer Science, Virginia Commonwealth University, Richmond, VA. Before joining VCU, he was an Assistant Professor in the Department of Computer Engineering, Hallym University, Korea from 1997–2001 and in the Department of Computer Science, The George Washington University, Washington DC from 2001–2003. His research interests include network design and analysis, combinatorial design, and application of algorithms.

Hyeong-Ah Choi has been a Professor in the Department of Computer Science at The George Washington University since 1988. Prior to her employment at GW, she was an Assistant Professor in the Department of Computer Science at Michigan State University from 1986–1988.Dr. Choi has engaged in optical and wireless networks research and published over 150 research articles in journals and conference proceedings. She has served as a Guest Editor of special issues of journals and as a member of technical program committees for numerous conferences and workshops. Her research has been supported by NSF, NSA, NIST, DARPA, DISA, and SKTelecom.