Lightpath Restoration in WDM Optical Networks

Gurusamy Mohan, National University of Singapore
C. Siva Ram Murthy, Indian Institute of Technology, Madras

Abstract

Optical networks employing wavelength-division multiplexing and wavelength routing are potential candidates for future wide area networks. Because these networks are prone to component failures and carry a large volume of traffic, maintaining a high level of service availability is an important issue. This article discusses providing fault tolerance capability to the optical layer in WDM-based transport networks. It presents a survey on restoration schemes available in the literature, explains the operation of these schemes, and discusses their performance.

The ever-increasing demand for bandwidth is posing new challenges for transport network providers. A viable solution to meet this challenge is to use optical networks based on wavelength-division multiplexing (WDM) technology. WDM divides the vast transmission bandwidth available on a fiber into several nonoverlapping wavelength channels and enables data transmission over these channels simultaneously.

The transport network architecture consists of three layers: circuit, path, and physical media [1]. Existing networks utilize WDM technology as a physical media layer for point-to-point transmission; these networks are called *point-to-point WDM networks*. In these networks, an optical signal is converted to an electrical signal, buffered, and transmitted again as an optical signal at every intermediate node before reaching the destination node. While these networks enhance transmission capacity, they do not possess sufficient node processing capability. Due to electro-optical conversion at intermediate nodes, the message delay increases; also, large buffers and an increased number of optical receivers and transmitters are required at the nodes.

A viable alternative to overcome the shortcomings of point-to-point WDM networks is to apply WDM technology to the path layer. This layer is also known as the *optical layer*. Here, a message is transmitted between two nodes using a *lightpath* without requiring any electro-optical conversion and buffering at the intermediate nodes. This is known as *wavelength routing*. A *lightpath* is an optical communication path between two nodes, established by allocating the same wavelength throughout the route of the transmitted data. A lightpath is uniquely identified by a wavelength and a physical path. At the optical layer, lightpaths are established between a subset of node pairs, forming a *virtual topology*. Networks with the optical layer have several advantages, such as enhanced node processing capability and protocol transparency. They are becoming a reality due to advances in optical technology. The application of WDM to the circuit layer is thought to still be premature since the number of wavelengths on a fiber is restricted to a small value [1].

The optical layer is formed between the lower physical media layer and the higher circuit layer in the transport network architecture. This layer is protocol-transparent, and can support different kinds of services and protocols at the higher layer such as synchronous optical network (SONET)/synchronous digital hierarchy (SDH), asynchronous transfer mode (ATM), and Internet protocol (IP). The higher circuit layer is also known as the *client layer* [2] because it is served by the optical layer.

A WDM network consists of *wavelength cross-connects* (WXCs) interconnected by fiber links. A *wavelength selective cross-connect* (WSXC) routes a message arriving at an incoming fiber on some wavelength to an outgoing fiber on the same wavelength. It is realized by wavelength demultiplexers, wavelength multiplexers, and optical switches. This architecture may not achieve the best network performance due to the *wavelength continuity constraint* which requires that the same wavelength be used on all the fiber links throughout the route of data transmission. An alternative is to use a *wavelength interchanging cross-connect* (WIXC) which employs wavelength conversion. *Wavelength interchange* (*convertible*) *networks* yield better performance than *wavelength selective networks*. However, wavelength converters are very expensive. Due to technological constraints, the number of wavelengths that can be supported on a fiber is limited (up to 160 wavelengths are used in commercially available systems). In order to meet the growing bandwidth demands of users, multiple fibers can be used in a link. Networks that employ multiple fibers in a link are known as *multifiber networks*.

WDM networks are prone to failures of components such as links, nodes, and wavelength multiplexers/demultiplexers. Since these networks carry high volumes of traffic, failures may have severe consequences. Therefore, it is imperative that these networks be fault-tolerant. Fault tolerance refers to the ability of the network to reconfigure and reestablish communication upon failure, and is widely known as *restoration* [3]. A network with restoration capability requires redundant capacity or spare resources. Restoration may be provided at the optical layer or the higher client layer, each of which has its own merits [4]. Restoration at the optical layer has several advantages, such as shorter restoration time, efficient resource utilization, and protocol transparency, over that at the client layer.

According to International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) Recommendation G.872, the optical layer can be further decomposed into three sublayers: the *optical channel layer*, *optical multiplex section layer*, and *optical transmission section layer* [2]. The optical channel layer provides end-to-end networking of optical channels (lightpaths) for transparently conveying client information such as SDH and ATM. The optical multiplex section layer provides functionality for networking of an aggregate optical signal with multiple wavelengths. The optical transmission section layer provides functionality for transmitting aggregate optical signals on different kinds of optical media. This layer is served by the lower physical media layer. Restoration at the optical channel layer deals with service recovery at the individual lightpath level. On the other hand, restoration at the optical multiplex section layer deals with recovering all the lightpaths (which are multiplexed together) simultaneously as a single group. It does not distinguish between the individual lightpaths multiplexed together.

This article is primarily concerned with restoration at the optical channel layer. It discusses several restoration algorithms, performance issues, and research results available in the literature. Most of the discussion pertains to the single-link failure model. This model assumes that at any instant of time at most one link has failed. The key ideas and approaches used for link failures can be extended to handle node failures and multiple component failures.

The rest of the article is organized as follows. We begin by presenting classification of restoration methods. Various multiplexing techniques that can be used to improve wavelength channel utilization are discussed. Several methods of designing restorable networks for a static traffic demand are presented. Restoration algorithms for routing dynamic traffic are dealt with. Distributed control protocols for restoration are described. Survivability mechanisms in ring networks are explained, and finally, some concluding remarks are made.

Classification of Restoration Methods

WDM networks are prone to component failures. A fiber cut causes a link failure. When a link fails, all its constituent fibers will fail. A node failure may be caused due to failure of the WXC. A fiber may fail due to failure of its end components. The lightpath that carries traffic during normal operation is known as the *primary lightpath*. When a primary lightpath fails, the traffic is rerouted over a new lightpath known as the *backup lightpath*. Failure detection, correlation, and root cause analysis are a difficult problem in WDM optical networks [5]. The nodes adjacent to the failed link can detect the failure by monitoring the power levels of signals on the links [5].

The restoration schemes differ in their assumption about the functionality of cross-connects, traffic demand, performance metric, and network control. Networks with WIXCs do not impose any wavelength continuity constraint. As a result, the wavelength channel utilization is higher in wavelength convertible networks than in wavelength selective networks.

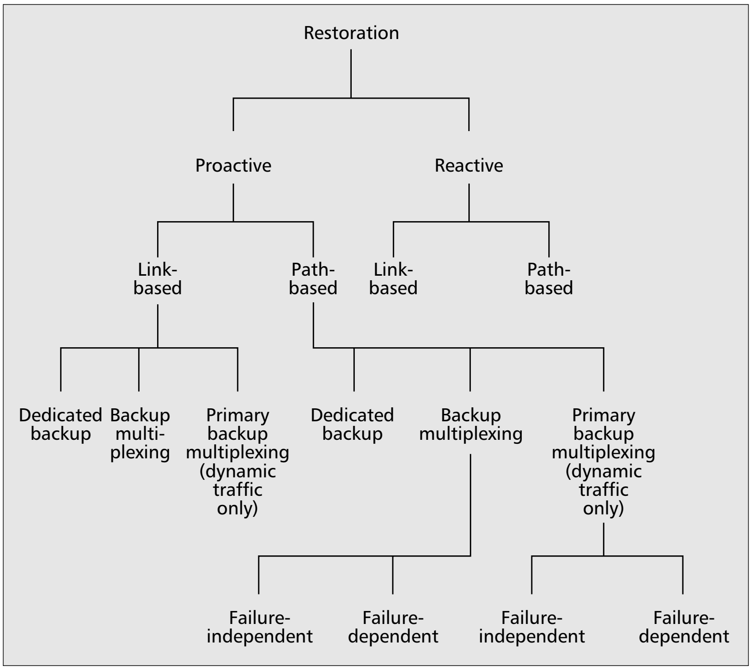

Figure 1. *Classification of lightpath restoration methods.*

The traffic demand can be either static or dynamic. In static traffic demand, the set of demands (or connection requests) is given a priori. The objective is to assign lightpaths with restoration capability to all the demands to minimize the spare resources (wavelengths or fibers) required. This problem is relevant for the capacity planning phase to determine the capacity needed in the near future based on current and expected demands [4]. Alternatively, the objective could be to satisfy as many demands as possible for a fixed amount of network resources. This problem is valid in a situation where there are new demands, and the objective is to route as many demands as possible using the available capacity on a network [4]. In a dynamic traffic environment, the demands arrive at a network one by one in a random manner. Once a demand is honored, it is held for a random finite time before being terminated. Here, the objective is to increase the acceptance ratio (or equivalently, to decrease the *blocking probability*) of demands. Dynamic traffic demand results in several situations in transport networks. First, it may become necessary to reconfigure the network in response to changing traffic demand patterns. Second, with the rise in broadband traffic it is expected that the leased line rates for private virtual networks and Internet service provider links will grow higher and higher. The demand for such services will change with time.

A restoration scheme may assume either centralized or distributed control. For large networks, distributed control is preferred over centralized control. A distributed control protocol requires several control messages to be exchanged between nodes. Also, there is a possibility of reservation conflicts between two simultaneous attempts to find paths.

The restoration methods can be classified as illustrated in Fig. 1. They are broadly classified as reactive and proactive methods. The reactive method is the simplest way of recovering from failures. In this method, when an existing lightpath fails, a search is initiated to find a new lightpath which does not use the failed components. This has an advantage of low overhead in the absence of failures. However, this does not guarantee successful recovery, since the attempt to establish a new lightpath may fail due to resource shortage at the time of failure recovery. Also, in case of distributed implementation,

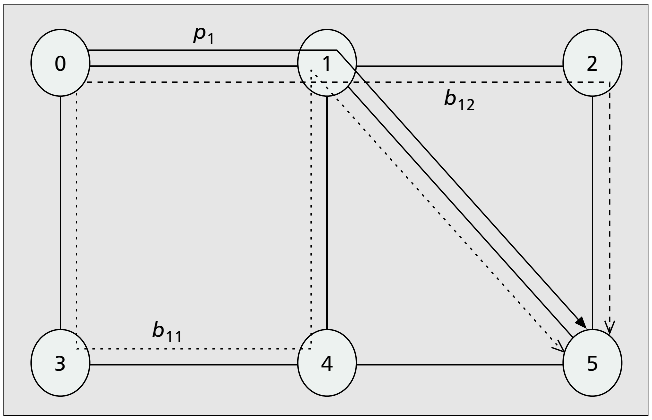

■ Figure 2. *An illustration of a link-based restoration method.*

contention among simultaneous recovery attempts for different failed lightpaths may require several retries to succeed, resulting in increased network traffic and restoration time. To overcome the shortcomings of reactive methods, proactive methods can be employed. In a proactive method, backup lightpaths are identified, and resources are reserved along the backup lightpaths at the time of establishing the primary lightpath itself. In doing so, this method yields a 100 percent *restoration guarantee*. This metric refers to the guarantee that a failed path finds its backup path readily available upon failure [6]. The backup lightpath takes over the role of the primary lightpath when it fails. Since the backup lightpath is established before a failure actually occurs, one can use it immediately upon occurrence of a failure on the primary, without invoking the time-consuming connection reestablishment process. Hence, the restoration time of a proactive method is much lower, leading to fast recovery.

The technique that uses preassigned capacity to ensure survivability is referred to as *protection*, and the technique that reroutes the affected traffic after failure occurrence by using available capacity is referred to as *restoration* in ITU-T Recommendation G.872 [2]. In [1, 3, 4, 7], terms such as *restoration network*, *restoration method*, and *restoration guarantee* have been used to refer to protection with a mention that the resources are preassigned. In [6, 8], the term *proactive* is used to refer to protection and the term *reactive* to refer to restoration. Throughout this article, we use the terms proactive and reactive to mean protection and restoration, respectively.

A proactive or reactive method is either *link-based* or *path-based* [4, 9]. The link-based method employs *local detouring*, while the path-based method employs *end-to-end detouring*. The link-based method reroutes traffic around the failed component. When a link fails, a new path is selected between the end nodes of the failed link. This path, along with the working segment of the primary path, will be used as the backup path. This method is unattractive for several reasons [4]. The choice of backup paths is limited; also, the backup paths are usually longer. Also, in wavelength selective networks, the backup path must necessarily use the same wavelength as the primary path since its working segment is retained. Furthermore, handling node failures this way is very difficult. The link-based restoration method is illustrated in Fig. 2. It shows a primary lightpath, p_1, and two backup lightpaths, b_{11} and b_{12}, on a wavelength. When link 0→1 fails, backup path b_{11} is used. It can be observed that b_{11} is routed around link 0→1 while retaining the working segment of p_1. When link 1→5 fails backup path b_{12} is used. It can be observed that b_{12} is routed around link 1→5 while retaining the working segment of p_1. Note that the working segment of the primary lightpath is retained in the backup path.

In the path-based restoration method, a backup lightpath is selected between the end nodes of the failed primary lightpath. Unlike in the link-based method, in the path-based method a backup lightpath need not retain the working segment of the primary lightpath. This method shows better resource utilization than the link-based restoration method. The backup path can use any wavelength independent of the one used by the corresponding primary lightpath. The path-based restoration method is illustrated in Fig. 3. It shows a primary lightpath, p_1, and its backup lightpath, b_1, on a wavelength. Note that b_1 is established between the end nodes of p_1, and the working segment of p_1 is not retained by b_1.

A proactive restoration method may use a dedicated backup lightpath for a primary lightpath. This is known as the *dedicated backup reservation* method. This method is illustrated in Fig. 4. The figure shows two primary lightpaths, p_1 and p_2, and their respective backup lightpaths, b_1 and b_2, on a wavelength. It can be observed that b_1 and b_2 do not share any wavelength channel. The dedicated backup reservation method has an advantage of shorter restoration time, since the WXCs are configured for the backup path when establishing the primary path itself. However, this method reserves excessive resources. Therefore, this method is not attractive and is not emphasized in this article.

For better resource utilization, multiplexing techniques can be employed. If two primary lightpaths do not fail simultaneously, their backup lightpaths can share a wavelength channel. This technique is known as *backup multiplexing*. In a dynamic traffic scenario, a proactive method can employ *primary backup multiplexing* [6] to further improve resource utilization. This technique allows a wavelength channel to be shared by a primary and one or more backup paths. By doing so, the blocking probability of demands decreases at the expense of reduced restoration guarantee.

A path-based restoration method is either *failure-dependent* or *failure-independent*. In a failure-dependent method, associated with the failure of every link used by a primary lightpath, there is a backup lightpath. When a primary lightpath fails, the backup lightpath that corresponds to the failed link will be used. A backup lightpath can use any link, including those used by the failed primary lightpath, except the failed link. Different backup lightpaths of a primary lightpath can share channels since they do not fail simultaneously in a single-link failure model. In a failure-independent method, a backup lightpath link-disjoint with the primary lightpath is chosen. This backup path is used upon link failure, irrespective of which of its links has failed. When this method is employed, a source node of a failed primary lightpath need not know the identity of the failed component. However, this method does not allow a backup path to use the channels used by the failed primary lightpaths. This will result in poorer resource utilization.

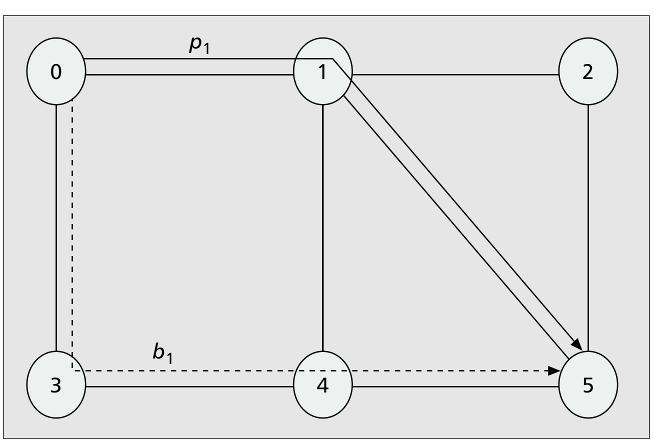

■ Figure 3. *An illustration of a path-based restoration method.*

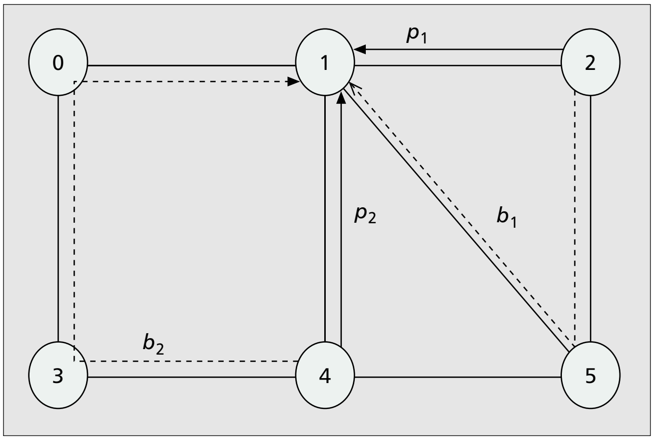

■ Figure 4. *An illustration of dedicated backup path reservation.*

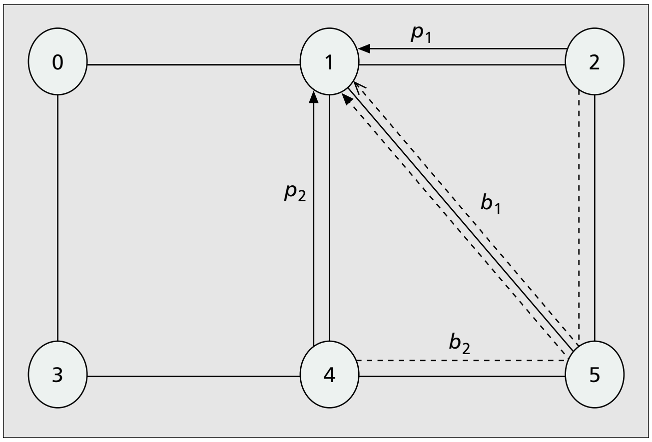

■ Figure 5. *An illustration of backup multiplexing.*

Multiplexing Techniques

In proactive methods, wavelength channels may be shared to improve their utilization. In this section we describe two multiplexing techniques: backup multiplexing and primary-backup multiplexing. We assume the single-link failure model.

Backup Multiplexing

The backup multiplexing technique allows two or more backup lightpaths to share a channel if the corresponding primary lightpaths do not fail simultaneously. When a failure occurs, control signals may be sent to nodes on the backup lightpath to configure the WXCs. Then message transmission starts on the backup lightpath.

This technique is illustrated in Fig. 5. It shows two primary lightpaths, p_1 and p_2, and their respective backup lightpaths, b_1 and b_2, on a wavelength. Since p_1 and p_2 are link-disjoint, they may not fail simultaneously. Therefore, b_1 and b_2 can share the wavelength on link 5→1. This shared channel will be used by b_1 when link 2→1 fails and by b_2 when link 4→1 fails.

Primary-Backup Multiplexing

The primary-backup multiplexing technique allows a primary lightpath and one or more backup lightpaths to share a channel [6]. By using this technique, an increased number of lightpaths can be established at the expense of reduced restoration guarantee. This technique is useful for dynamic traffic where the lightpaths are short-lived. This technique is motivated by the following factors. First, failures do not occur frequently enough in practice to warrant full reservation. Second, every lightpath does not necessarily need fault tolerance to ensure network survivability. Third, at any instant of time, only a few primary lightpaths critically require fault tolerance. For such critical paths, backup lightpaths may be exclusively reserved; for others the restoration guarantee could be less than 100 percent.

A lightpath loses its recoverability when a channel on its backup lightpath is used by some other primary lightpath. It regains its recoverability when the other primary lightpath terminates. A lightpath loses its recoverability only when the following three events occur simultaneously:
- A link fails.
- The failed link is used by the lightpath.
- A channel on its backup lightpath is used by some other primary lightpath.

However, such a situation is less probable. When a primary lightpath fails, the source node sends control messages along the route of the backup path to configure WXCs. If any channel in the route is used by a primary lightpath, the establishment process fails; otherwise, the backup path is established.

This technique is illustrated in Fig. 6. It shows three primary lightpaths, p_1, p_2, and p_3, and their respective backup lightpaths, b_1, b_2, and b_3, on a wavelength. Backup lightpaths b_1 and b_2 share the channel on link 4→1, since p_1 and p_2 are link-disjoint. The channel on link 1→2 is shared by p_3 and b_1, and the channel on link 5→2 is shared by p_1 and b_3. Therefore, both p_1 and p_3 are nonrecoverable. However, if one of them terminates, the other regains recoverability immediately.

Lightpath Restoration for Static Traffic

In this section we discuss methods of designing restorable WDM networks with static traffic demand. The process of assigning network resources to static traffic demand is called *provisioning a network*. Given a set of demands and their primary lightpaths, the restoration network design problem is to find backup lightpaths for all the demands to minimize the capacity required. The capacity is measured in terms of number of wavelengths for a single-fiber network and number of fibers for a multifiber network.

Provisioning Restorable Multifiber Networks

In this section we briefly describe the design methods available in the literature for provisioning restorable multifiber networks.

Restoration Network Design — The restoration network design for multifiber networks was first considered in [1]. Two schemes, *virtual wavelength path* (VWP) and *wavelength path* (WP), were proposed. They assume wavelength interchange and wavelength selective cross-connects, respectively. Both

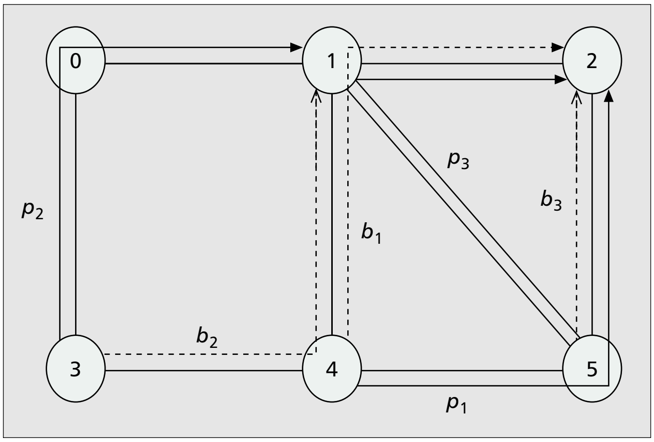

■ Figure 6. *An illustration of primary-backup multiplexing.*

schemes are proactive, path-based, and failure-dependent, employing backup multiplexing. A weighing function w_i is associated with every link i. The *path load*, denoted p_i, defines the number of paths traversing link i. The *residual wavelengths*, denoted c_i, defines the number of used wavelengths on the last fiber in link i and is given by p_i mod W. Here, W is the number of wavelengths per fiber. The value of w_i is set to p_i if c_i is 0; otherwise, it is set to $1/c_i$. The design algorithms choose primary and backup paths so that the path loads are minimized, and the residual wavelengths on the links are as close to W as possible. This is expected to improve wavelength utilization and reduce fiber requirements.

The VWP scheme uses an iterative algorithm. It is required to select a backup path for each link used by a primary path among several precomputed candidate paths. The algorithm starts with the primary network. It proceeds in two phases, *initial backup path selection* and *backup path optimization*. Link failures are considered one by one. A link fails, and a backup path is chosen for each failed path. Other links fail in a similar way, and initial backup paths are chosen. While choosing these paths, the weighing functions as described above are used. The initial backup paths are then optimized to reduce the fiber requirements. Again, link failures are considered one by one. A link fails, and the failed lightpaths are replaced with the initial backup paths. A new backup path is chosen in place of the initial backup path only if it results in reduction of fiber requirements.

Like the VWP scheme, the WP scheme also uses an iterative algorithm that works in two phases. Initial backup lightpaths are set up using a similar procedure to that of the VWP scheme. Initially, the backup paths use the same wavelength as their corresponding primary paths. Then the backup paths are optimized to further reduce fiber requirements. The set S of links whose fiber requirements need to be reduced is determined. The links in set S fail one by one. When a link fails, the failed lightpaths are replaced with the initial backup lightpaths. The initial backup lightpath is then replaced with a new backup lightpath if it reduces fiber requirements. Two methods are considered while choosing the new backup lightpath. In *method 1*, a backup lightpath uses the same wavelength as the primary lightpath. In *method 2*, a backup lightpath can use any wavelength. While method 2 is flexible in choosing a wavelength for a backup path, it requires coordination from the client layer.

The performance of the proposed design methods was evaluated in terms of fiber requirements and accommodation efficiency on a polygrid network. The accommodation efficiency is defined as the ratio between the number of laid out paths on a link and the maximum number of paths that can be accommodated. The VWP scheme performs better than WP since it does not impose a wavelength continuity constraint. However, the hardware cost of the cross-connects in the VWP scheme is high. The results also show that method 2 performs better than method 1, and the difference in performance becomes smaller when the number of wavelengths is larger.

Independent and Coordinated Design — The restoration network design problem for multifiber wavelength selective networks was addressed in [3]. The objective is to minimize the weighted number of fibers required. A fiber between a pair of nodes is weighted by its length. The work considers a proactive failure-dependent path-based restoration scheme. Backup multiplexing is used to improve wavelength channel utilization. Two iterative design methods, *independent* and *coordinated design*, were proposed, and their performance was evaluated through simulation.

A network may have different failure scenarios, each corresponding to a different set of component failures. A set may refer to one or more link/node failures. It is necessary to design a restoration network for a given primary network. For each primary lightpath in a failure scenario, there is a set of candidate backup paths. Associated with each failure scenario is an *impact set*. The impact set consists of a set of primary lightpaths whose backup lightpaths need to be determined for the failure scenario. This set must include the failed primary paths. The independent design method finds the set of backup paths for each failure scenario independent of other scenarios. An *integer linear programming* (ILP) formulation is solved for a failure scenario to minimize the resource requirement on links. The wavelengths/fibers required on a link will be the maximum value required by any failure scenario found by the ILP. Since the failure scenarios are considered independently, the order in which they are considered does not matter; therefore, computational complexity is significantly reduced. It is to be noted that computational complexity also depends on the size of each impact set and the number of failure scenarios. Since the spare capacity information calculated for a failure scenario is not exploited by other scenarios, this method requires more spare resources. To overcome this shortcoming, coordinated design was proposed.

The coordinated design method has a number of iterations. The failure scenarios are processed in some order. For each failure scenario, a backup path for each node pair is determined. A metric $M^*(p, w)$ is associated with every lightpath whose path is p and wavelength w. The metric $M^*(p, w)$ is calculated as the summation of another metric $m^*(e, w)$ which is associated with each edge e on w used by the lightpath. These metrics are a representative of load on the links. The pair $<e, w>$ refers to a (wavelength) channel. A channel contributes to the lightpath metric only if there is a failure scenario under which the channel becomes fully loaded. In such a case, its contribution is inversely proportional to the number of similar channels on the link. The algorithm starts with the primary network and a failure scenario. For this scenario, node pairs are considered in some order. For a node pair, those lightpaths that are members of the corresponding impact set are replaced by suitable backup lightpaths. A lightpath that maximizes M^* is chosen. Among several possible backup candidate paths, the one that minimizes M^* in the current state is chosen. This is repeated for every node pair. The remaining failure scenarios are also processed similarly. It is to be noted that at any iteration, the spare resources computed for all the previous failure scenarios are taken into consideration.

The independent and coordinated design methods were evaluated through simulation for three restoration schemes: full reconfiguration, path-based, and link-based. In a full reconfiguration scheme, not only the failed lightpaths, but all the lightpaths are allowed to choose a backup lightpath. The performance was measured in terms of fiber requirements. The full reconfiguration scheme yields the best performance followed by path-based and-link based, in that order. Furthermore, the coordinated design method results in roughly 20 percent savings in fiber requirements over independent design.

ILP and Heuristic Solutions — Provisioning multifiber networks was considered in [7] for both wavelength converting and wavelength selective networks. Three proactive restoration methods were proposed. These methods include path-based failure-independent and failure-dependent, and link-based methods. ILP formulations were presented for all methods. The ILP formulations were solved for a small network, and the performance of all the schemes was compared. The results show that the spare capacity requirements in wavelength convertible networks are less than in wavelength selective networks. They also show that the spare capacity requirements are least in failure-dependent path-based restoration followed by failure-independent path-based and link-based restoration, in that order. Since ILP formulations are computationally prohibitive for larger networks, heuristic algorithms were proposed. For path-based restoration in wavelength selective networks, two methods were considered. While method 2 allows the backup lightpath to use any wavelength independent of its primary lightpath, method 1 selects the same wavelength for both lightpaths. The heuristic algorithms use path-based restoration.

For method 1, paths and wavelengths are selected separately. Backup and primary paths are selected using a cost criterion, ignoring wavelength assignment. Wavelengths are then assigned to each of the primary-backup lightpath pairs. The paths are ranked in decreasing order of length, and the longest ones are assigned wavelengths first. For each node pair, the wavelength that requires the fewest fibers to be added along both the primary and backup lightpaths is chosen.

The heuristic algorithm for method 2 and convertible networks works as follows. It starts with the primary network. A link is chosen randomly and is set as failed. All the node pairs whose primary lightpaths have failed are ranked in decreasing order of hop count. For each node pair, the backup lightpath that requires the fewest fibers to be added is chosen. This is repeated for each link failure.

The performance of the proposed heuristics was evaluated on a wide range of networks. The performance trends are similar to the results obtained by ILP solutions. The effect of network connectivity was also studied analytically and experimentally. The experimental results show that the spare capacity required decreases as connectivity increases. This is because the effectiveness of backup multiplexing is greater in densely connected networks.

The effect of traffic growth on the performance of wavelength convertible and wavelength selective networks was studied in terms of utilization of deployed capacity. While routing a new primary-backup lightpath pair, all existing primary and backup lightpaths are maintained. The results show that the utilization gain in wavelength convertible networks is greater than in wavelength selective networks, and the difference increases when the number of wavelengths per fiber increases.

Comparison of the Design Methods — The above design methods can be compared in respect to spare resources required, cost of cross-connects, and algorithm complexity. Both wavelength converting and wavelength selective networks were considered in [1, 7]. While wavelength converting networks require less spare resources, they are very expensive. While the design method developed in [1] is less complex than that in [7], its performance is slightly poorer. The coordinated design method developed in [3] yields better performance than the other methods since it effectively makes use of spare resources while the algorithm is in progress. However, it is more complex, and the order in which the failure scenarios are processed has an effect on the spare resources required. The design methods developed in [3] are more flexible than others in the sense that they can handle multiple component failures. They also optionally allow the unaffected lightpaths to be rerouted to reduce the spare resource requirements.

Provisioning Restorable Single-Fiber Networks

In this section we describe the design methods available in the literature for provisioning restorable single-fiber networks.

Wavelength Convertible Network Design

In [4], the problem of designing the restoration network for a given set of demands for wavelength convertible networks was considered. It is assumed that every node is capable of converting a wavelength to any wavelength; thus, the wavelength continuity constraint is not considered. The problem is formulated as an integer programming problem. The objective function is to minimize the weighted number of wavelengths required. The links are weighted by capacity consumption per wavelength. Single link and node failures were considered. Necessary constraints were provided to ensure the following:

- A wavelength is used by a backup lightpath on a link only when the primary lightpath fails.
- The primary and backup lightpaths for a demand are chosen to be link-disjoint.
- Failure-independent path-based restoration is used.

Because ILP formulation is computationally complex, heuristic algorithms are proposed. The problem is divided into several independent subproblems, one for each demand, by removing the constraint on link capacity and using Lagrangian multipliers. Two algorithms were presented. They work on the following idea. A cost value is assigned to a primary-backup pair for a demand based on the subproblem. The first algorithm starts with an initial feasible solution for a demand. Different candidate routes are considered for possible use as primary paths in increasing order of cost. For each path as a primary path, the best path is chosen as the backup path. The procedure continues until the cost value of the pair cannot be improved further. The above procedure is used for each demand. They will collectively form an initial solution to the original problem. Then an iterative procedure is used to compute the best feasible result. The second algorithm works as follows. It starts with the pair of shortest disjoint paths for every demand. For every demand, the primary path is fixed and the best backup path found. Then, for every demand, the backup path is fixed and the best primary path found.

The performance of these algorithms was evaluated through simulation on different random networks. The results show that the first algorithm performs slightly better than the second one, and both perform better than a shortest-path pair-based algorithm. It has also been established that the performance improvement achieved by backup multiplexing is significant when compared to dedicated backup reservation.

Wavelength Selective Network Design

In [9], ILP formulations were developed for three different proactive restoration methods: dedicated backup reservation, path-based restoration using backup multiplexing, and link-based restoration using backup multiplexing. A single-fiber wavelength selective network was considered. The objective here is to minimize the number of wavelengths used on the links. Constraints were provided for wavelength continuity, exclusive reservation or sharing of a wavelength channel by different backup lightpaths, and path-based and link-based restoration. The performance of all three methods was evaluated for a representative interconnected ring network by solving ILPs. The results show that the path-based method performs significantly better than the other two. It also shows that the link-based method performs slightly poorer than even the dedicated backup reservation method.

Lightpath Restoration for Dynamic Traffic

Unlike static traffic demand, dynamic traffic demand requires computationally simpler algorithms. Since lightpath demands arrive one by one, the objective of a dynamic routing algorithm is to select the best primary-backup lightpath pair to improve network blocking performance. Blocking performance is measured in terms of blocking probability of arrived demands. Ideally, we expect the algorithm to route as many restorable lightpaths as possible. We now describe two algorithms developed in [6, 8] for routing dynamic traffic. The algorithms are proactive and use the failure-independent path-based approach.

Backup-Multiplexing-Based Routing

The algorithm presented in [8] uses backup multiplexing. It basically uses an alternate routing method, wherein a set of candidate routes for every source-destination pair is precomputed. The candidate routes of an s-d pair are chosen to be link-disjoint. In response to a new request, a minimum-cost primary-backup lightpath pair is chosen. The key idea here is to choose the pair that requires the minimum free wavelength channels. A channel is free if it is used by neither any primary lightpath nor any backup lightpath. If a wavelength channel is already used by one or more backup lightpaths, it can be used by a new backup lightpath (if allowed) at no extra cost. Thus, the algorithm ensures that at the time of routing a new lightpath pair, the network is taken to a new state to maximize the total number of free channels in the network.

For every wavelength channel on a link, the algorithm maintains a list of links whose failure will lead to the use of the channel by a backup lightpath of a failed primary lightpath. In other words, the list associated with a channel consists of those links used by all the primary lightpaths whose backup lightpaths use the channel. A new backup lightpath can use a wavelength channel only if the corresponding primary lightpath does not use any of the links in the list associated with the channel. Two methods of wavelength assignment were discussed. In the first method, the primary and backup lightpaths use the same wavelength. In the second method, there is no such restriction on wavelength usage by the primary and backup lightpaths. While the first method is computationally simpler than the second, it results in poor blocking performance.

The performance of the algorithm has been verified through simulation experiments on ARPA-2 and mesh-torus networks. The connectivity of a mesh-torus network is denser than that of an ARPA-2 network. The results show that for a given blocking probability, both ARPA-2 and mesh-torus networks are able to carry more traffic load when the proposed algorithm is used than with the dedicated backup reservation method. The factor by which the carried load increases is about 3 for mesh-torus and about 0.8 for ARPA-2. The results show that the usefulness of backup multiplexing increases as network connectivity increases. This is because in a densely connected network, the candidate routes are usually shorter, and the number of possible link-disjoint candidate routes is greater.

Primary-Backup Multiplexing-Based Routing

The algorithm presented in [6] uses primary-backup multiplexing. It also uses alternate routing and a proactive path-based approach. The objective of this algorithm is to improve blocking performance while allowing an acceptable reduction in restoration guarantee. Here, a wavelength channel is allowed to be shared by a primary lightpath and one or more backup lightpaths. Because of this, the lightpaths that correspond to the backup lightpaths on this channel lose their restoration capability. A new lightpath pair may cause an increase in the average number of nonrestorable lightpaths per link. A lightpath pair is admissible only if its establishment does not take the network to a state where the average number of nonrestorable lightpaths per link exceeds a predefined threshold value. For a lower threshold value, the restoration guarantee is higher and the blocking performance lower. By appropriately choosing a threshold value, a desired trade-off can be achieved between restoration guarantee and network blocking performance.

The key idea of the algorithm is to choose a minimum cost lightpath pair among those that are admissible. The cost of a primary-backup lightpath pair computed by this algorithm is the number of free channels used by the pair plus the number of primary-backup multiplexed channels traversed times a penalty factor. A sufficiently high value is chosen as the penalty factor so that a pair which traverses a minimum number of primary-backup multiplexed channels is chosen, and in case of a tie, the number of free channels used is considered. An important issue here is how to compute the number of nonrestorable lightpaths created by routing a lightpath pair. When a new backup lightpath traverses a channel currently used by a primary lightpath, only the new primary lightpath becomes nonrestorable. However, it is nontrivial to compute the count of nonrestorable lightpaths created by routing a new primary lightpath over the channels reserved for some other backup paths. A straightforward solution is to keep the identity and restorability status of every primary lightpath that corresponds to backup lightpaths on each of the channels. This requires large storage and also a more complex algorithm to compute the count and update the restorability status. To overcome the above shortcomings, a computationally simple heuristic method was proposed to estimate the count of nonrestorable lightpaths created by routing a new primary lightpath. This heuristic need know only the number of backup lightpaths multiplexed on a channel and the number of links used by their primary lightpaths. It also keeps track of the number of backup lightpaths that continue to the next link.

The performance of the algorithm was verified through simulation experiments on ARPA-2 and mesh-torus networks. The performance improvement over the backup-multiplexing-based algorithm [8] was studied. The performance is measured using two metrics: relative performance gain and reduction in restoration guarantee. The blocking performance of the network when no lightpath is provided with a backup lightpath is taken as the lower limit, and the performance when only backup multiplexing is used [8] as the upper limit. The performance of the proposed algorithm using primary-backup multiplexing is measured with relation to the above two limits. The results show that the performance gain is attractive enough to allow some reduction in restoration guarantee. In particular, under light load conditions, more than 90 percent performance gain is achieved at the expense of less than 10 percent guarantee reduction. They also show that the performance improves as network connectivity increases.

Distributed Control Protocols

Many existing algorithms for selecting routes and wavelengths for primary and backup lightpaths assume centralized control. Such algorithms are useful for small networks and are not scalable to large networks. For simplicity and scalability purposes, distributed control protocols are desired. In a distributed implementation, up-to-date network state information is not known to any of the nodes. Also, there is a possibility of reservation conflicts between several simultaneous search attempts. This may result in poor channel utilization. Another important issue in reactive methods is to reduce restoration time since the search for a backup path is initiated only after failure.

A Distributed Control Protocol for Pro-active Methods

In [4], a distributed control protocol for selecting primary and backup lightpaths is presented. It basically uses proactive path-based restoration. It assumes that every routing node is equipped with WIXC (i.e., capable of performing wavelength conversion optically). The same method can be used for point-to-point WDM networks, where every node is capable of performing wavelength conversion by means of electro-optical conversion and buffering. However, this protocol requires suitable changes when the routing nodes are equipped with WSXCs.

The procedure for finding a backup path for a given primary path is explained here. It is to be noted that this procedure can be extended to finding both the primary and backup paths simultaneously. Since the nodes possess conversion capability, a path need not use the same wavelength on its links. Every link has a fixed capacity. Link capacity is measured in terms of wavelength channels. A demand requires a backup path to be established between a source and destination node, for a given primary path between the same nodes.

The protocol considers single link and node failures. At every node, link capacity control tables are maintained for each of the outgoing links. A table associated with a link has information about the possible (node or link) failures in the network. The following discussion pertains to one link or table. Associated with every possible failure is a list of demands whose backup paths will require the link upon failure. Each demand in a list requires one wavelength on the link. The maximum over the size of lists in a link capacity control table gives the number of wavelengths used on the link thus far. In other words, this protocol employs backup multiplexing. From the size of the list associated with a failure, the number of free wavelengths available (residual capacity) on the link upon occurrence of failure can be determined. A source node initiates the search for a backup path independent of searches by other source nodes. It sends control messages carrying the identity of the primary path on the supervisory channel. A distributed breadth-first search method is used to select a backup path. When a node receives such a probe message, it examines the table that corresponds to the outgoing link. In the table, it looks up the residual capacity associated with the failures that correspond to the nodes and links traversed by the primary path of the demand. By doing so, it determines if a free wavelength is available for the backup path on the outgoing link.

It is possible for the search process to occur for multiple demands simultaneously. Because of this, resource (wavelength channel) contention may occur on a link, leading to a deadlock situation. This protocol prevents resource contention as follows. Two primary paths that do not share a link or node (except source and destination) will never contend for a wavelength on any link. Therefore, a source node, before initiating a search for a backup path, locks the links and nodes on its primary path. Once the backup path search procedure terminates, the locks are released.

Because of the distributed nature and lack of global coordination of the search procedure, the resulting routes may need more wavelength channels than a centralized algorithm is able to achieve. In order to improve channel utilization, an optimization procedure is used. This procedure is invoked either periodically or whenever a search fails. Optimization is carried out by changing the existing backup paths. Since the backup paths are used as cold-standby, the traffic on the network is not disrupted.

The effectiveness of the distributed protocol was evaluated using a simplified simulation model since the complete simulation model is extremely complex. The simplified simulation model captures only the dynamics of demand contention locking. For a sample random network with 28 nodes, 48 links, and a set of 50 demands, the number of demands blocked by the distributed protocol is compared to that for an optimal solution. The simulation results show that the performance of the distributed protocol is very close to that of the optimal solution.

A Distributed Control Protocol for Reactive Methods

A distributed control protocol for reactive methods was proposed in [5]. Upon a link failure, this protocol searches for backup lightpaths for the failed lightpaths. Both link-based and path-based restoration were considered. It uses a *two-phase* process to search for a backup path between a pair of nodes on a given wavelength. These nodes are the end nodes of the failed link in link-based restoration and the end nodes of the failed path in path-based restoration. The two-phase process is described below.

The source node sends broadcast messages on all outgoing links. A message reserves (locks) the wavelength on a link while traversing it. This message carries the hop count of the path it has traversed so far. When a broadcast message reaches the destination node, it sends a confirm message along the path to the source node. Upon receiving a confirm message, intermediate nodes configure their cross-connects. When the source node receives a confirm message, backup path establishment is done. All the wavelengths reserved during the broadcast phase are released by cancel messages sent by a node upon a timeout or receipt of a message with a hop - count exceeding a predefined value.

Link-based restoration requires that the backup path use the same wavelength used by the corresponding primary path. Since the failed paths are on different wavelengths, the distributed path selection for the failed paths is carried out in parallel without any conflict. In path-based restoration, the free wavelengths on the failed link are partitioned into sets of wavelengths, one set for each failed path. This results in parallel path search on different wavelengths, avoiding conflicts between them.

The performance of both methods was evaluated on an interconnected ring topology. The performance metrics considered were restoration efficiency and restoration time. Restoration efficiency is defined as the ratio between the number of lightpaths restored and the total number of lightpaths failed. Suitable values for propagation delay and message processing delay were assumed. The results show that the link-based method has shorter restoration time than the path-based method. This can be attributed to two reasons. First, no control messages need be sent to the end nodes of the failed path in a link-based method. Second, backup paths around the failed link are likely to be shorter. The results show that the path-based method has better restoration efficiency, because a path-based method can use any wavelength for the backup path. It has been further observed that restoration time for both methods is on the order of milliseconds only. Also, restoration efficiency decreases with increasing load. This is because at a high load, fewer channels are free, and there are more reservation conflicts as many paths fail.

WDM Self-Healing Rings

SONET self-healing rings (SHRs) have been very successful mainly due to simple control and fast service restoration. For the same reasons, ring networks are a promising architecture for applying WDM technology [10]. The SONET SHR concept can be adapted to WDM ring networks. Like SONET SHRs, WDM rings can also be classified into several categories such as *two-fiber unidirectional rings* (UR-2), *four-fiber bidirectional rings* (BR-4), and *two-fiber bidirectional rings* (BR-2). The above architectures can be extended to multifiber ring networks where a link can have a number of fibers in multiples of 2 or 4.

In a UR-2 system, there are two unidirectional rings in opposite directions. Under normal conditions, the message traffic flows in one ring, called a *working ring*. When failure occurs, the other ring, called a *protection ring*, will be used to reroute the affected traffic. The protection mechanism employed in UR-2 is either *line-switched* or *path-switched*. A line switched UR-2 basically uses a *loopback* method. When a link between node i and node $i + 1$ fails, node i loops the traffic from the working ring to the protection ring, and node $i + 1$ loops the traffic from the protection ring to the working ring. The newly formed ring is unidirectional and carries the traffic in one direction. In a path-switched UR-2, primary paths use resources on the working ring, and backup paths use resources on the protection ring. Messages are simultaneously transmitted on both the primary and backup lightpaths. The destination node of a connection receives good signals from both the working and protection rings during normal operation, and that from the working ring is used. When a link used by the primary lightpath fails, the destination node receives a bad signal from the working ring and a good signal from the protection ring. The good signal is then used by the destination node. Both link and path protection switching mechanisms can be used to handle single node failures also. When node i fails, a line-switched UR-2 performs loopback at node $i - 1$ and node $i + 1$. In a path-switched UR-2, all paths except those for which node i is a source or destination are protected.

In a BR-4 system, there are two pairs of rings with each pair consisting of two unidirectional rings in opposite directions. One pair of rings carries message traffic during normal conditions. Thus, unlike in UR-2, message traffic flows in both directions in BR-4. When failure occurs, the affected traffic is rerouted over the other pair of rings. This is accomplished by loopback. In a BR-2 system, there are two unidirectional rings in opposite directions. In each ring, half the capacity is used for message traffic during normal operation, and the remaining capacity is used for rerouting the affected traffic when a component fails.

Conclusions

In this article, we present a survey on lightpath restoration schemes for WDM optical networks available in the literature. We explained the algorithms used by these schemes and discussed their performance. The restoration time for the reactive methods is longer, and restoration is not guaranteed when compared to proactive methods. However, in the absence of failures, resource utilization is more efficient in reactive methods. While link-based methods have shorter restoration time, they do not utilize resources efficiently when compared to path-based methods. Failure-dependent proactive path-based methods utilize resources more efficiently than failure-independent methods; however, they are more complex. Employing backup multiplexing results in significant performance improvement over dedicated backup reservation. In a dynamic traffic environment, proactive methods employing primary-backup multiplexing yield significant improvement over backup multiplexing, at the expense of reduced restoration guarantee. WDM rings are more promising architectures from the perspective of fast failure restoration and simple control. SONET's self-healing ring concepts can be applied to WDM rings to provide fault tolerance.

References

[1] N. Nagatsu, S. Okamoto, and K. Sato, "Optical Path Cross-Connect System Scale Evaluation Using Path Accommodation Design for Restricted Wavelength Multiplexing," *IEEE JSAC*, vol. 14, no. 5, June 1996, pp. 893–902.

[2] ITU-T Rec. G.872, "Architecture of Optical Transport Networks," Feb. 1999.

[3] M. Alanyali and E. Ayanoglu, "Provisioning Algorithms for WDM Optical Networks," *IEEE/ACM Trans. Net.*, vol. 7, no. 5, Oct. 1999, pp. 767–78.

[4] B. T. Doshi *et al.*, "Optical Network Design and Restoration," *Bell Labs Tech. J.,* Jan.–Mar. 1999, pp. 58–84.

[5] S. Ramamurthy and B. Mukherjee, "Survivable WDM Mesh Networks, Part II — Restoration," *Proc. ICC '99*, 1999.

[6] G. Mohan and A. K. Somani, "Routing Dependable Connections With Specified Failure Restoration Guarantees in WDM Networks," *Proc. IEEE INFOCOM 2000*, Mar. 2000.

[7] S. Baroni *et al.*, "Analysis and Design of Resilient Multifiber Wavelength-Routed Optical Transport Networks," *IEEE/OSA J. Lightwave Tech.*, vol. 17, no. 5, May 1999, pp. 743–58.

[8] G. Mohan and C. S. R. Murthy, "Routing and Wavelength Assignment for Establishing Dependable Connections in WDM Networks," *Proc. IEEE Int'l Symp. Fault-Tolerant Comp.*, June 1999, pp. 94–101.

[9] S. Ramamurthy and B. Mukherjee, "Survivable WDM Mesh Networks, Part I — Protection," *Proc. IEEE INFOCOM '99*, 1999, pp. 744–51.

[10] R. Ramaswami and K. Sivarajan, *Optical Networks: A Practical Perspective*, Morgan Kaufmann, 1998.

Biographies

GURUSAMY MOHAN (elegm@nus.edu.sg) received his Ph.D. degree in computer science and engineering from Indian Institute of Technology (IIT), Madras in 2000. From January 1991 to January 2000, he served as a faculty member in the Department of Computer Science and Engineering, Regional Engineering College, Tiruchirappalli, India. From January to June 2000, he worked as a senior project officer in the Department of Computer Science and Engineering, IIT, Madras. Since June 2000 he has been working as a faculty member in the Department of Electrical and Computer Engineering, National University of Singapore. He held a visiting position at Iowa State University from January to June 1999. His research interests include optical networks, and parallel and distributed computing.

C. SIVA RAM MURTHY (murthy@iitm.ernet.in) received his Ph.D. degree in computer science from Indian Institute of Science (IISc), Bangalore, in 1988. From March to September 1988 he worked as a scientific officer in the Supercomputer Education and Research Centre at IISc. He subsequently joined IIT Madras as a lecturer of computer science and engineering. He became an assistant professor in August 1989 and is currently a professor at the same place. He has held visiting positions at German National Research Centre for Information Technology (GMD), Sankt Augustin, Germany, University of Washington, Seattle, and University of Stuttgart, Germany. He is a recipient of the Seshagiri Kaikini Medal for the best Ph.D. thesis and also of the Indian National Science Academy Medal for Young Scientists. He has co-authored three books in the areas of parallel processing, parallel algorithms, and real-time systems.