

12. Routing and Wavelength (Spectrum) Assignment

Jane M. Simmons, George N. Rouskas

Routing a connection from its source to its destination is a fundamental component of network design. The choice of route affects numerous properties of a connection, most notably cost, latency, and availability, as well as the resulting level of congestion in the network. This chapter addresses various algorithms, strategies, and tradeoffs related to routing.

At the physical optical layer, connections are assigned a unique wavelength on a particular optical fiber, a process known as wavelength assignment (WA). Together with routing, the combination of these two processes is commonly referred to as RWA. In networks based on all-optical technology, WA can be challenging. It becomes more so when the physical properties of the optical signal need to be considered. This chapter covers several WA algorithms and strategies that have produced efficient designs in practical networks.

A recent development in the evolution of optical networks is flexible networking, where the amount of spectrum allocated to a connection can be variable. Spectrum assignment is analogous to, though more complex than, wavelength assignment; various heuristics have been proposed as covered in this chapter. Flexible (or elastic) networks are prone to more contention issues as compared to traditional optical networks. To maintain a high degree of capacity efficiency, it is likely that spectral defragmentation will be needed in these networks; several design choices are discussed.

12.1	Terminology	449
12.2	Shortest-Path Routing Algorithms	451
12.2.1	Dijkstra Algorithm	451
12.2.2	Breadth-First-Search Algorithm	451
12.2.3	Constrained Shortest-Path Routing	451
12.2.4	K-Shortest-Paths Routing Algorithms	452
12.2.5	Shortest-Distance Versus Minimum-Hop Routing	452
12.3	Disjoint-Path Routing for Protection	453
12.3.1	Disjoint-Path Routing with Multiple Sources and/or Destinations	455
12.3.2	Shared-Risk Link Groups	455
12.4	Routing Strategies	456
12.4.1	Fixed-Path Routing	456
12.4.2	Alternative-Path Routing	457
12.4.3	Dynamic Routing	457
12.5	Routing Order	458
12.6	Multicast Routing	458
12.6.1	Manycast Routing	459
12.7	Wavelength Assignment	460
12.7.1	Interaction Between Regeneration and Wavelength Assignment	460
12.8	Wavelength Assignment Algorithms	462
12.8.1	First-Fit Algorithm	464
12.8.2	Most-Used Algorithm	464
12.9	One-Step RWA	464
12.9.1	Topology Pruning	464
12.9.2	Reachability-Graph Transformation ...	465
12.9.3	Flow-Based Methods	466
12.9.4	ILP-Based Ring RWA	466
12.10	Impairment-Aware Routing and Wavelength Assignment	467
12.10.1	Mixed Line-Rate Systems	468
12.11	Flexible (Elastic) Optical Networks	469
12.12	Routing and Spectrum Assignment in Elastic Optical Networks	471
12.12.1	Routing	471
12.12.2	Spectrum Assignment	472
12.12.3	ILP-Based RSA	474
12.12.4	Distance-Adaptive RSA	474
12.12.5	Routing, Spectrum and Core Assignment	475
12.12.6	Multipath Routing	475
12.12.7	Defragmentation	476
12.13	Conclusion	478
	References	479

Routing a connection (or circuit) from its source to its destination is a fundamental component of network design. The choice of route affects numerous properties of a connection, most notably cost, latency, and availability. Furthermore, from a more macro perspective, routing determines the resulting level of congestion in a network, which may impact the amount of traffic that ultimately can be carried. This chapter addresses various algorithms, strategies, and tradeoffs related to routing. While much of the discourse applies to networks in general, the emphasis is on routing in the physical optical layer. Relevant network terminology and some of the fundamentals of optical systems are introduced in Sect. 12.1.

Network routing is a well-researched topic. Several commonly used routing algorithms have been developed to find an optimal path through the network with respect to a particular property. These algorithms are reviewed in Sect. 12.2. For a connection of relatively low criticality, it is sufficient to calculate a single path from source to destination. However, a failure that occurs along that path will bring the connection down. To guard against such failures, more critical connections are established with some level of protection. Some commonly used protection schemes rely on calculating two or more diverse paths between the source and destination. Optimal disjoint-path routing algorithms exist for this purpose as well, as covered in Sect. 12.3.

As noted above, the routing strategy that is utilized affects the congestion level in the network. Poor routing may lead to unbalanced load in the network, with some sections heavily utilized whereas others are lightly loaded. Such imbalances result in premature blocking in the network where it may not be possible to accommodate a new connection request due to insufficient capacity in portions of the network. Various routing strategies and their effect on network load are discussed in Sect. 12.4.

If multiple connections are being routed at once, then the order in which they are processed may also affect the blocking level. For example, it may be desirable to give precedence to geographically longer connections or connections that must traverse congested areas of the network, as it is likely more challenging to find a suitable path for these connections. Selecting the routing order is, in general, an arbitrary procedure; with rapid design runtimes, several orderings can be tested and the best result selected. Routing order is covered in Sect. 12.5.

Historically, connections in telecommunications networks have been unicast, with a single source and a single destination. However, with services such as video distribution growing in importance, connections directed from one source to multiple destinations are

desirable. Establishing a multicast tree from the source to the set of destinations is more capacity efficient than utilizing multiple unicast connections. Furthermore, the optical layer is well suited for delivering multicast services. Section 12.6 discusses routing algorithms that find the optimal, or near-optimal, multicast tree.

The end-to-end path is the high-level view of how a connection is carried in the network. When examined more closely at the physical optical layer, it is necessary to assign a portion of the electromagnetic spectrum to carry the connection at each point along its path. Optical networks typically employ wavelength-division multiplexing (WDM) technology, where connections are assigned a unique wavelength on a particular optical fiber and the various wavelengths are multiplexed together into a single WDM signal on that fiber. Selecting the wavelengths to assign to each connection is the process known as *wavelength assignment*; this topic is introduced in Sect. 12.7. The combination of the routing and wavelength assignment processes is commonly referred to as RWA.

The underlying optical system technology determines the level of difficulty involved with wavelength assignment. With one class of technology, known as optical-electrical-optical (O-E-O), where connections repeatedly go from the optical domain to the electrical domain and back, the wavelengths assigned to connections on one link have no impact on the wavelengths that can be assigned on any other link. With such technology, wavelength assignment is a simple process. O-E-O technology, however, does have numerous drawbacks, as is outlined in Sect. 12.1.

In contrast, optical-optical-optical (O-O-O), or *all-optical*, technology potentially enables connections to remain in the optical domain from source to destination. While there are many advantages to this technology, it does lead to challenges in the wavelength assignment process. In practice, most networks fall somewhere in between pure O-E-O and O-O-O, where connections enter the electrical domain at intermediate points of the path, but do so infrequently. It is common practice to refer to these networks as all-optical, despite this term being a misnomer; this chapter makes use of this naming convention as well. The respective properties of O-E-O and all-optical networks drive the routing and wavelength assignment processes and are discussed further in Sect. 12.1.

Poor wavelength assignment in all-optical networks can lead to blocking, where the capacity exists to carry a new connection but no wavelength is free to assign to that connection. This is known as *wavelength contention*. Various wavelength assignment heuristics that are simple to implement yet effective in minimizing

wavelength contention have been developed, as outlined in Sect. 12.8.

Another important design decision is whether routing and wavelength assignment should be treated as separate processes or whether they should be combined into a single operation. This dichotomy is discussed in Sect. 12.9.

In an all-optical network, there are numerous practical challenges of wavelength assignment that are tied to the physical optical system. Some of the more important effects are introduced in Sect. 12.10.

A recent development in the evolution of optical networks is *flexible networking*, introduced in Sect. 12.11. In traditional optical networks, there is a fixed number of wavelengths that can be carried on each optical fiber, and each such wavelength occupies a fixed amount of spectrum. In contrast, with flexible (or elastic) networks, the amount of spectrum allocated to a connection can be variable, potentially leading to a more efficient use of the fiber capacity. Selecting the portion of the spectrum to assign to a given connection is known as *spectrum assignment*. As elucidated in Sect. 12.12, spectrum assignment is analogous to, though more complex than, wavelength assignment. Various heuristics have been proposed to specifically address the nuances of spectrum assignment.

Flexible networks tend to lead to more contention issues as compared to traditional optical networks. This is especially true if the network is dynamic, such that connections are frequently established and then torn

down. Dynamism is expected to result in fragmented spectrum utilization where many small blocks of unutilized spectrum exist as opposed to large contiguous blocks of free spectrum. To maintain a high degree of capacity efficiency, it is likely that *spectral defragmentation* would need to be performed on a periodic basis, where connections are shifted to different regions of the spectrum (and possibly rerouted) to create larger chunks of available spectrum. Defragmentation techniques are discussed in Sect. 12.12.

Network design can be performed on different time scales. In *long-term* network planning, there is sufficient time between the planning and provisioning processes such that any additional equipment required by the plan can be deployed. The planning emphasis is on determining the optimal strategy for accommodating a set of demands; the runtime of the design algorithms is not critical. In *real-time* operation, there is little time between planning and provisioning. It is assumed that the traffic must be accommodated using whatever equipment is already deployed in the network, even if that sacrifices optimality. Any design calculations must be performed quickly (milliseconds to minutes time scale, depending on the requirements of the network). The long-term and real-time dichotomy is also referred to in the literature as *offline* and *online* design, respectively. Methodologies for both time scales are discussed in this chapter.

A number of algorithms are introduced in this chapter. Implementation of many of these algorithms, using the C programming language, can be found in [12.1].

12.1 Terminology

Network terminology that is relevant to the topics of this chapter is defined here. Additionally, a high-level comparison of O-E-O and all-optical networks is provided, with an emphasis on the aspects that affect routing and wavelength assignment.

Network nodes are the sites in the network that source, terminate, and/or switch traffic. Sites that serve only to amplify the optical signal are not considered nodes. Network nodes are depicted as circles in the figures contained here. *Network links* are the physical optical fibers that run between the nodes, with the fibers typically deployed in pairs. Links are almost always bidirectional, with one fiber of the pair carrying traffic in one direction and the other fiber carrying traffic in the opposite direction. Network links are depicted in figures as a single line. In discussions on routing, links are often referred to as *hops*.

The interconnection pattern among the nodes represents the *network topology*. The *nodal degree* is the number of fiber-pairs incident on a node. The nodal degree may also be defined as the number of links incident on a node. A link may be populated with multiple fiber-pairs, such that the two definitions are not always equivalent.

In a network that employs O-E-O technology (i.e., an O-E-O network), the connection signal is carried in the optical domain on the links but is converted to the electrical domain at each node that is traversed by the connection. An end-to-end path looks as depicted in Fig. 12.1a. The process of O-E-O conversion at each node cleans up the signal; more precisely, the signal is re-amplified, re-shaped, and re-timed. This is known as *3R regeneration*. (Note that passing a signal through a line optical-amplifier is considered *1R regen-*




Fig. 12.1a–c The connection path from node A to node Z is shown by the *dotted line*. (a) In an O-E-O network, regeneration occurs at each node along the path of the connection. (b) In a true all-optical network, there is no regeneration along the path. The connection remains in the optical domain from end to end. (c) With an optical reach of 2000 km and the link distances as shown, two regenerations are required for the connection. (© Monarch Network Architects LLC)

eration, as it only re-amplifies the signal.) Regeneration equipment is needed for each individual signal that traverses a node; i.e., the WDM signal is demultiplexed into its constituent wavelengths and each one of the wavelengths that is routed through the node is passed through a regenerator. (Terminal equipment, as opposed to a regenerator, is used for a wavelength that is destined for the node as opposed to being routed *through* the node; a regenerator often comprises back-to-back terminal equipment, as discussed further in Sect. 12.7.1.) Thus, a network node may be equipped with hundreds of regenerators. (Limited *multiwavelength regeneration* has been demonstrated but is still in the early stages of research.)

The large amount of equipment required for regeneration poses drawbacks in terms of cost, power consumption, reliability, and space requirements. In O-E-O networks, regeneration is the major contributor to these factors in the optical layer. Minimizing the number of nodes that a signal must traverse is desirable, as it minimizes the number of regenerations.

Conversely, a true all-optical signal that remains in the optical domain is illustrated in Fig. 12.1b. No re-

generations are present along the end-to-end path. The network nodes are equipped with elements such as *reconfigurable optical add/drop multiplexers* (ROADMs) that allow the signal to remain in the optical domain. This is known as *optical bypass* of the node.

An important property of all-optical networks is the *optical reach*, which is the distance that an optical signal can travel before it needs to be regenerated. (In practice, there are numerous factors other than distance that affect where regeneration is required [12.1]. For simplicity, we only consider distance here.) Most optical-bypass-enabled systems have an optical reach on the order of 1500–3000 km. If the length of the end-to-end path is longer than the optical reach, then the signal needs to be regenerated, possibly more than once, as illustrated in Fig. 12.1c. This is more often the case in core networks (also known as backbone networks, long-haul networks, or cross-country networks), where some of the end-to-end paths may be a few thousand kilometers in extent.

Regeneration is almost always performed in a network node as opposed to at an intermediate point along a link. Thus, regeneration typically occurs prior to the signal traveling exactly a distance equal to the optical reach. This may lead to more regeneration being required than predicted by the connection distance. For example, in Fig. 12.1c, the total end-to-end path distance is 3900 km, but with an optical reach of 2000 km, two regenerations are required, not one. Furthermore, a study performed on several realistic backbone networks indicated that the shortest path was not the path with the fewest number of required regenerations for roughly 1% of the source/destination pairs in those networks [12.1].

Taking advantage of extended optical reach and ROADMs to remove all, or at least most, of the required regeneration provides significant advantages in terms of cost, power consumption, reliability, and space. The drawback is that if a signal traverses a node in the optical domain, then it must be carried on the same wavelength into and out of the node. This is known as the *wavelength continuity constraint*. (Wavelength conversion in the optical domain is possible; however, it is costly and complex and has not been widely commercialized.) Thus, the presence of optical bypass creates an interdependence among the links of a network, where the wavelength assigned to a connection on one link affects the wavelengths that can be assigned to connections on other links. This leads to wavelength assignment being a critical aspect of network design in all-optical networks. Note that such an interdependence does not exist in O-E-O networks where regeneration at every node along a path allows the wavelengths to be assigned independently on each link.

12.2 Shortest-Path Routing Algorithms

We first consider algorithms that find a single path that is the shortest between a given source and destination [12.2]. If multiple paths are tied for the shortest, then one of the paths is found. Depending on the algorithm used, ties may be broken differently.

The class of shortest-path algorithms discussed here assumes that each link is assigned a cost metric and that the metric is additive. With an additive metric, the metric for a path equals the sum of the metrics of each link composing that path. Furthermore, there can be no cycles in the network where the sum of the link metrics around the cycle is negative. With these assumptions, a shortest-path algorithm finds the path from source to destination that minimizes the total cost metric.

These algorithms can be applied whether or not the links in the network are bidirectional. A link is bidirectional if traffic can be routed in either direction over the link. Different metrics can be assigned to the two directions of a bidirectional link. However, if the network is bidirectionally symmetric, such that the traffic flow is always two-way and such that the cost metric is the same for the two directions of a link, then a shortest path from source to destination also represents, in reverse, a shortest path from destination to source. In this scenario, which is typical of telecommunications networks, it does not matter which endpoint is designated as the source and which is designated as the destination.

There are numerous metrics that are useful for network routing. If the link metric is the geographic distance of the link, then the shortest-path algorithm does, indeed, find the path with the shortest distance. If all links are assigned a metric of unity (or any constant positive number), then the shortest-path algorithm finds the path with the fewest hops. Using the negative of the logarithm of the link availability as the metric produces the most reliable path, assuming that the failure rate of a path is dominated by independent link failures. (Note that the logarithm function can be used in general to convert a multiplicative metric to an additive metric.) Other useful metrics for routing in an all-optical network are the link noise figure and the link optical signal-to-noise ratio (OSNR) [12.1].

In spite of the link metric not necessarily being associated with distance, these algorithms are commonly still referred to as shortest-path algorithms.

12.2.1 Dijkstra Algorithm

The best known shortest-path algorithm is the Dijkstra algorithm [12.2]. The algorithm works by tracking the shortest path discovered thus far to a particular node,

starting with a path length of zero for the source node. It considers the resulting length if the path is extended to a neighboring node. If the path to the neighboring node is shorter than any previously discovered paths to that node, then the path to that neighbor is updated. Dijkstra is classified as a greedy algorithm because it makes the optimal decision at each step without looking ahead to the final outcome. Unlike many greedy algorithms, it is guaranteed to find the optimal solution; i.e., the shortest path from source to destination.

12.2.2 Breadth-First-Search Algorithm

An alternative shortest-path algorithm is breadth-first search (BFS) [12.3]. It works by discovering nodes that are one hop away from the source, then the nodes that are two hops away from the source, then the nodes that are three hops away from the source, etc., until the destination is reached.

As with Dijkstra, BFS produces the shortest path from source to destination. Additionally, if there are multiple paths that are tied for the shortest, BFS breaks the tie by finding the one that has the fewest number of hops. This can be useful in all-optical networks where the likelihood of wavelength contention increases with the number of hops in a path. Note that Dijkstra does not have a similar tie-breaking property; its tie-breaking is somewhat more arbitrary (e.g., it may depend on the order in which the link information is stored in the database).

12.2.3 Constrained Shortest-Path Routing

A variation of the shortest-path problem arises when one or more constraints are placed on the desired path; this is known as the *constrained shortest path* (CSP) problem. Some constraints are straightforward to handle. For example, if one is searching for the shortest path subject to all links of the path having at least N wavelengths free, then prior to running a shortest-path algorithm, all links with fewer than N free wavelengths are removed from the topology. As another example, the intermediate steps of the BFS shortest-path algorithm can be readily used to determine the shortest path subject to the number of path hops being less than H , for any $H > 0$ (similar to [12.4]). However, more generally, the CSP problem can be difficult to solve; for example, determining the shortest path subject to the availability of the path being greater than some threshold, where the availability is based on factors other than distance. Various heuristics have been proposed to address the CSP problem [12.5]. Some heuristics have been developed

to specifically address the scenario where there is just a single constraint; this is known as the *restricted shortest path* (RSP) problem. Additionally, a simpler version of the multiconstraint problem arises when *any* path satisfying all of the constraints is desired, not necessarily the shortest path; this is known as the *multiconstrained path* (MCP) problem. An overview, including a performance comparison, of various heuristics that address the RSP and MCP problems can be found in [12.6].

12.2.4 *K*-Shortest-Paths Routing Algorithms

Routing all connection requests between a particular source and destination over the shortest path can lead to network congestion, as is discussed more fully in Sect. 12.4. It is generally advantageous to have a set of paths to choose from for each source/destination pair. A class of routing algorithms that is useful for finding a set of *K* possible paths between a given source and destination is the *K-shortest-paths algorithm*. Such an algorithm finds the shortest path, the second shortest path, the third shortest path, etc., up until the *K*-th shortest path or until no more paths exist. A commonly used *K*-shortest-paths algorithm is Yen's algorithm [12.7]. As with shortest-path algorithms, the link metrics must be additive along the path.

It is important to note that the *K* paths that are found are not necessarily completely diverse with respect to the links that are traversed; i.e., some links may appear in more than one of the *K* paths. *K*-shortest-paths routing is typically used for purposes of load balancing, not for protection against failures. Explicitly routing for failure protection is covered in Sect. 12.3.

12.2.5 Shortest-Distance Versus Minimum-Hop Routing

As stated earlier, a variety of link metrics can be used in a shortest-path algorithm. Two of the most common metrics in a telecommunications network are link distance and unity (i.e., 1), producing the path of shortest geographic distance and the path of fewest hops, respectively. The most effective metric to use depends on the underlying network technology.

With O-E-O networks, the cost of a connection in the optical layer is dominated by the number of required regenerations. Regeneration occurs at every node that is traversed by the connection. Thus, utilizing 1 as the metric for all links in order to find the path with the least number of hops (and hence the fewest traversed nodes) is advantageous from a cost perspective. This is illustrated in Fig. 12.2 for a connection between nodes A and Z. Path 1 is the shortest-distance path at




Fig. 12.2 Path 1, A–B–C–D–Z, is the shortest-distance path between nodes A and Z, but Path 2, A–E–Z, is the fewest-hops path. In an O-E-O network, where the signal is regenerated at every intermediate node, Path 2 is typically the lower-cost path (© Monarch Network Architects LLC)

900 km but includes four hops. Path 2, though it has a distance of 1200 km, is typically lower cost in an O-E-O network because it has only two hops and, thus, requires fewer regenerations.

Selecting a metric to use with all-optical networks is not as straightforward. Assuming that the number of regenerations required for a connection is dominated by the path distance, then searching for the shortest-distance path will typically minimize the number of regenerations. However, wavelength assignment for a connection becomes more challenging as the number of path hops increases. For all-optical networks, one effective strategy is to generate candidate paths by invoking a *K*-shortest-paths algorithm twice, once with distance as the link metric and once with 1 as the link metric. Assuming that link load is not a concern, the most critical factor in selecting one of the candidate




Fig. 12.3 Assume that this is an all-optical network with an optical reach of 2000 km. Path 1, A–B–C–Z, has the fewest hops but requires one regeneration. Path 2, A–D–E–F–G–Z, and Path 3, A–H–I–J–Z, require no regeneration. Of these two lowest-cost paths, Path 3 is preferred because it has fewer hops (© Monarch Network Architects LLC)

paths is typically cost. Thus, the candidate path that requires the fewest number of regenerations is selected. If there are multiple paths tied for the fewest regenerations, then the one with the fewest number of hops is selected.

This strategy is illustrated in Fig. 12.3. Three of the possible paths between nodes A and Z are: Path 1: A–B–C–Z of distance 2500 km and three hops; Path 2: A–D–E–F–G–Z of distance 1500 km and five hops; and

Path 3: A–H–I–J–Z of distance 1600 km and four hops. Assume that the optical reach is 2000 km, such that there can be no all-optical segment that is longer than this distance. With this assumption, Path 1 requires one regeneration, whereas Paths 2 and 3 do not require any regeneration. Of the latter two paths, Path 3 has fewer hops and is, thus, more desirable despite it being longer than Path 2 (again, assuming that link load is not a factor).

12.3 Disjoint-Path Routing for Protection

Network customers desire a certain level of *availability* for each of their connections. Availability is defined as the probability of being in a working state at a given instant of time. The desired availability is typically specified contractually in the *service level agreement* (SLA) with the network provider. A common cause of connection failure is the failure of one or more links in the end-to-end path. Link failures are typically caused by fiber cuts or optical amplifier failures. Being able to route around the failed link allows the connection to be restored to the working state.

In some protection schemes, it is necessary to first identify which link has failed; the detour route around that failed link is then utilized. Determining the location of a failure can be time consuming. To expedite the restoration process and improve availability, network providers often utilize protection schemes where the same backup path is utilized regardless of where the failure has occurred in the original (i.e., primary) path. This allows the protection process to commence prior to the completion of the fault location process. In order to implement such a failure-independent protection

scheme, also known as *path protection*, it is necessary that the backup path be completely link-disjoint from the primary path. If a very high level of availability is desired for a connection, then it may be necessary to protect against node failures in addition to link failures. In this scenario, the backup path must be completely link and node disjoint from the primary path. However, it should be noted that node failures occur much less frequently than link failures, such that node disjointness is often not required.

Shortest-path algorithms find the single shortest path from source to destination. With path protection, it is desirable to find two disjoint paths where the overall sum of the link metrics on the two paths is minimized. One of the paths is used as the primary path (typically the shorter one) with the other serving as the backup path.

It may seem reasonable to find the desired two paths by invoking a shortest-path algorithm twice. After the first invocation, the links (and intermediate nodes if node disjointness is also required) composing the first path are removed from the network topology that is

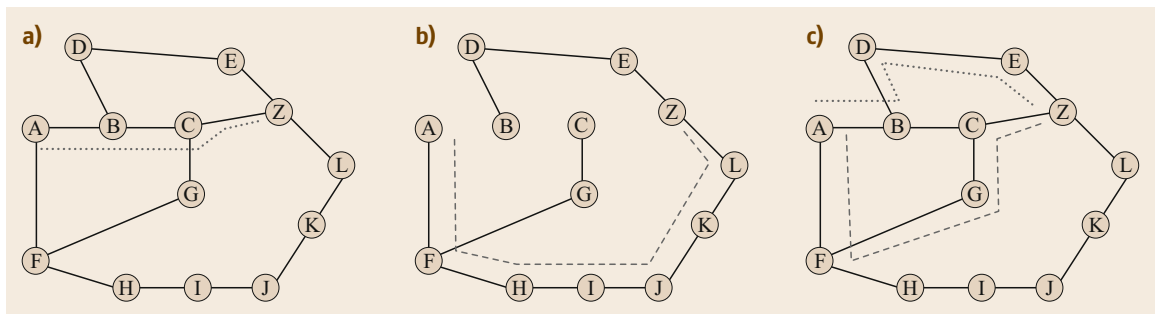


Fig. 12.4a–c The pair of disjoint paths with the fewest number of total hops is desired between nodes A and Z. (a) The first call to the shortest-path algorithm returns the path shown by the *dotted line*. (b) The network topology after pruning the links composing the shortest path. The second call to the shortest-path algorithm finds the path indicated by the *dashed line*. The total number of hops in the two paths is ten. (c) The shortest pair of disjoint paths between nodes A and Z is shown by the *dotted* and *dashed lines*; the total number of hops in these two paths is only eight (© Monarch Network Architects LLC)




Fig. 12.5a-c The shortest pair of disjoint paths is desired between nodes A and Z. (a) The first call to the shortest-path algorithm returns the path shown by the *dotted line*. (b) The network topology after pruning the links composing the shortest path. The second call to the shortest-path algorithm fails, as no path exists between nodes A and Z in this pruned topology. (c) The shortest pair of disjoint paths between nodes A and Z, shown by the *dotted* and *dashed lines* (© Monarch Network Architects LLC)

used in the second call of the algorithm. Clearly, if a path is found by the second invocation, then that path is guaranteed to be disjoint from the first path that was found.

Although this strategy may appear to be reasonable, there are potential pitfalls. First, it inherently assumes that the shortest disjoint pair of paths includes the shortest path. As shown in Fig. 12.4, this is not always the case. In this example, the source is node A and the destination is node Z. Assume that this is an O-E-O network, with all links assigned a metric of 1. The shortest path (i.e., the fewest-hops path) is A-B-C-Z, as shown in Fig. 12.4a. If the links of this path are removed and the shortest-path algorithm is invoked again, then the resulting path is A-F-H-I-J-K-L-Z; see Fig. 12.4b. The sum of the metrics of these two paths is 10. However, the shortest disjoint pair of paths is actually A-B-D-E-Z and A-F-G-C-Z, as illustrated in Fig. 12.4c; this pair of paths has a total metric of 8. Note that the shortest path, A-B-C-Z, is not part of the optimal solution.

Not only may the two-invocation method produce a suboptimal result, it may fail completely. This is illustrated in Fig. 12.5. Assume that this is an all-optical network, and the metric is the physical distance as shown next to each link. In this figure, the shortest path from A to Z is A-D-C-Z, as shown in Fig. 12.5a. If the links of this path are removed from the topology, as shown in Fig. 12.5b, then no other paths exist between nodes A and Z. Thus, the process fails to find a disjoint pair of paths despite the existence of such a pair, as shown in Fig. 12.5c: A-B-C-Z and A-D-E-Z. This type of scenario, where the two-invocation methodology fails despite the existence of disjoint paths, is called a *trap topology*.

Rather than using the two-invocation method, it is recommended that an algorithm specifically designed to find disjoint paths be used. The two most commonly used shortest-disjoint-paths algorithms are the Suurballe algorithm [12.8, 9] and the Bhandari algorithm [12.3]. Both algorithms make use of a shortest-path algorithm; however, extensive graph transforma-

tions are performed as well to ensure that the shortest pair of disjoint paths is found, assuming such a pair exists. Both algorithms require that the link metric be additive and both algorithms are guaranteed to produce the optimal result. They can be utilized to find the shortest pair of link-disjoint paths or the shortest pair of link-and-node-disjoint paths. The runtimes of the Suurballe and Bhandari algorithms are about the same; however, the latter may be more easily adapted to various network routing applications. We focus on its use here.

The Bhandari algorithm is readily extensible. For example, a mission-critical connection may require a high level of availability such that three disjoint paths are needed, one primary with two backups. If a failure occurs on the primary path, the first backup path is used. If a failure occurs on the backup path prior to the primary path being repaired, then the connection is moved to the second backup path. The Bhandari algorithm can be used to find the shortest set of three disjoint paths, if they exist. More generally, the Bhandari algorithm can be used to find the shortest set of N disjoint paths, for any N , assuming N such paths exist. (In most optical networks, however, there are rarely more than just a small number of disjoint paths between a given source and destination, especially in a backbone network.)




Fig. 12.6 There is no completely disjoint pair of paths between nodes A and Z. The set of paths shown by the *dotted* and *dashed lines* represents the shortest maximally-disjoint pair of paths. The paths have nodes D and G, and the link between them, in common (© Monarch Network Architects LLC)




Fig. 12.7 (a) A disjoint path is desired between one source (node A) and two destinations (nodes Y and Z). (b) A dummy node is added to the topology and connected to the two destinations via links that are assigned a metric of zero. A shortest-disjoint-paths algorithm is run between node A and the dummy node to implicitly generate the desired disjoint paths, as shown by the *dotted line* and the *dashed line* (© Monarch Network Architects LLC)

In some scenarios, completely disjoint paths between the source and destination do not exist. In such cases, the Bhandari algorithm can be utilized to find the shortest *maximally-disjoint* set of paths. The resulting paths minimize the number of links (and optionally nodes) that are common to multiple paths. This is illustrated in Fig. 12.6, where the shortest maximally-disjoint pair of paths is A–C–D–G–H–Z and A–E–F–D–G–I–J–K–Z. Note that link DG is common to both paths.

12.3.1 Disjoint-Path Routing with Multiple Sources and/or Destinations

An important variation of the shortest-disjoint-paths routing problem exists when there is more than one source and/or destination, and each of the source/destination paths must be mutually disjoint for protection purposes. (Note that this is different from multicast routing, where the goal is to create a tree from one source to multiple destinations.) This protection scenario arises, for example, when backhauling traffic to multiple sites, utilizing redundant data centers in cloud computing, and routing through multiple gateways in a multidomain topology [12.1]. The example of Fig. 12.7a is used to illustrate the basic strategy. We assume that this is a backhauling example, where node A backhauls its traffic to two diverse sites, Y and Z. (Backhauling refers to the general process of transporting traffic from a minor site to a major site for further distribution.) Additionally, it is required that the paths from A to these two sites be disjoint.

In order to apply the shortest-disjoint-paths algorithm, a dummy node is added to the network topology as shown in Fig. 12.7b. Links are added from both Y and Z to the dummy node, and these links are assigned

a metric of 0. The shortest-disjoint-paths algorithm is then run with node A as the source and the dummy node as the destination. This implicitly finds the desired disjoint paths as shown in the figure.

A similar strategy is followed if there are D possible destinations, with $D > 2$, and disjoint paths are required from the source to M of the destinations, with $M \leq D$. Each of the D destinations is connected to a dummy node through a link of metric 0. An extensible shortest-disjoint-paths algorithm such as the Bhandari algorithm is invoked between the source and the dummy node to find the desired M disjoint paths. Note that this procedure implicitly selects M of the D destinations that produce the shortest such set of disjoint paths.

If the scenario is such that there are multiple sources and one destination, then the dummy node is connected to each of the sources via links with a metric of 0. The shortest-disjoint-paths algorithm is then run between the dummy node and the destination. If there are both multiple sources and multiple destinations, then two dummy nodes are added, one connected to the sources and one connected to the destinations. The shortest-disjoint-paths algorithm is then run between the two dummy nodes. The algorithm does not allow control over which source/destination combinations will result.

12.3.2 Shared-Risk Link Groups

One challenge of routing in practical networks is that the high-level network topology may not reveal interdependencies among the links. Consider the network topology shown in Fig. 12.8a and assume that it is desired to find the shortest pair of disjoint paths from node A to node Z. From this figure, it appears that the paths A–B–Z and A–Z are the optimal solution. However, the fiber-level depiction of the network, shown in Fig. 12.8b, indicates that links AB and AZ are not fully




Fig. 12.8 (a) In the link-level view of the topology, links AB and AZ appear to be disjoint. (b) In the fiber-level view, these two links lie in the same conduit exiting node A and are, thus, not fully diverse. A single cut to this section of conduit can cause both links to fail. (c) Graph transformation to account for the SRLG extending from node A (© Monarch Network Architects LLC)

disjoint. The corresponding fibers partially lie in the same conduit exiting node A such that a rupture to the conduit likely causes both links to fail. These two links are said to constitute a *shared-risk link group* (SRLG). Links that are a member of the same SRLG are not fully disjoint. It is desirable to avoid a solution where an SRLG link is in the primary path and another link in that same SRLG is in the backup path.

For the SRLG configuration depicted in Fig. 12.8b, it is straightforward to find truly disjoint paths. First, a graph transformation is performed as shown in Fig. 12.8c, where a dummy node is added and each link belonging to the SRLG is modified to have this dummy node as its endpoint instead of node A. A link of metric 0 is added between node A and the dummy

node. The shortest-disjoint-paths algorithm is run on the modified topology to find the desired solution: paths A–B–Z and A–C–D–Z. (While it is not necessary in this example, for this type of SRLG graph transformation, the shortest-disjoint-paths algorithm should be run with both link and node disjointness required [12.1].)

The scenario of Fig. 12.8 is one class of SRLG, known as the *fork configuration*. There are several other SRLG configurations that appear in practical networks [12.1, 3]. There are no known computationally efficient algorithms that are guaranteed to find the optimal pair of disjoint paths in the presence of any type of SRLG. In some cases, it may be necessary to employ heuristic algorithms that are not guaranteed to find the optimal (shortest) set of disjoint paths [12.10].

12.4 Routing Strategies

Monitoring network load is an essential aspect of network design. If too much traffic is carried on a small subset of the links, it may result in an unnecessarily high blocking rate, where future connections cannot be accommodated despite the presence of free capacity on most network links. The routing strategy clearly affects the network load. We discuss three of the most common routing strategies here.

12.4.1 Fixed-Path Routing

The simplest routing strategy is known as *fixed-path routing*. One path is calculated for each source/destination pair, and that path is utilized for every connection request between those two nodes. If any link along that path has reached its capacity (e.g., 80 connections are already routed on a link that supports 80 wavelengths), then further requests between the source/destination pair are blocked.

This strategy is very simple to implement, as it requires no calculations to be performed on an ongoing

basis; all calculations are performed up front. Routing is a simple binary decision: either all links of the precalculated path have available capacity for the new connection, or they do not.

The drawback to fixed-path routing is that it is completely nonadaptive. The same path is selected regardless of the load levels of the links along that path. Such load-blind routing typically leads to uneven load distribution among the links, which ultimately may lead to premature blocking.

Despite this drawback, fixed-path routing is employed by many (if not most) network service providers, where the shortest-distance path is the one path selected for each source/destination pair. The rationale for this strategy is that it minimizes the latency of the connection. (Latency is the propagation delay from source to destination; in the optical layer, latency is typically dominated by the distance of the path.) While there are some network customers where latency differences on the order of microseconds can be critical (e.g., enterprises involved with electronic financial trad-

ing [12.11]), for most applications, a tolerance of a few extra milliseconds is acceptable. In most realistic networks, it is possible to find alternative paths that are very close in distance to the shortest path such that the increase in latency is acceptably small. By not considering such alternative paths, fixed-path routing is overly restrictive and can lead to poor blocking performance.

12.4.2 Alternative-Path Routing

A second routing strategy is *alternative-path routing*. In contrast to fixed-path routing, a set of candidate paths is calculated for each source/destination pair. At the time of a connection request, one of the candidate paths is selected, typically based on the current load in the network. For example, the candidate path that results in the lowest load on the most heavily loaded link of the path is selected. This enables adaptability to the current load levels in the network.

It is well known that alternative-path routing can lead to lower blocking levels as compared to fixed-path routing [12.12, 13]. One study performed by a service provider on its own large backbone network demonstrated one to two orders of magnitude lower blocking probability using alternative-path routing, even when limiting the length of the alternative paths for purposes of latency [12.14].

Furthermore, a large number of candidate paths is not required to achieve significant blocking improvements. Generating on the order of three candidate paths is often sufficient to provide effective load balancing. The key to the success of this strategy lies in the choice of the candidate paths. There is typically a relatively small set of links in a network that can be considered the *hot links*. These are the links that are likely to be in high demand and, thus, likely to become congested. An effective strategy to determine the expected hot links is to perform a design for a typical traffic profile where all connections in the profile follow the shortest path. The links that are the most heavily loaded in this exercise are generally the hot links.

With knowledge of the likely hot links, the candidate paths for a given source/destination pair should be selected such that the paths exhibit diversity with respect to these links. It is not desirable to have a particular hot link be included in all of the candidate paths, if possible, as this will not provide the opportunity to avoid that link in the routing process. Note that there is no requirement that the candidate paths for a given source/destination be *completely* diverse. While total disjointness is necessary for protection, a less severe diversity requirement is sufficient for purposes of load balancing [12.15].

For completeness, we mention one variation of alternative-path routing known as *fixed alternative-path routing*. In this scheme, a set of candidate paths is generated and ordered. For each connection request, the first of the candidate paths (based on the fixed ordering) that has free capacity to accommodate the request is selected. Thus, load balancing is implemented only after one or more links in the network are full, thereby providing limited benefit as compared to fixed-path routing.

12.4.3 Dynamic Routing

With dynamic routing, no routes are precalculated. Rather, when a connection request is received by the network, a search is performed for a path at that time. Typically, path selection is based on cost and/or load. For example, consider assigning each link in an O-E-O network a value of $LARGE + L_j$, where $LARGE$ is a very large constant and L_j is a metric that reflects the current load level on link j . (Any load-related metric can be used, as long as the metric is additive.) Running a shortest-path algorithm with these link values will place the first priority on minimizing the number of hops in the path, due to the dominance of the $LARGE$ component. In an O-E-O network, minimizing the number of hops also minimizes the number of regenerations required. The second priority is selecting a path of minimal load, as defined by the metric. Once the network contains several full or close-to-full links, the priorities may shift such that each link j is assigned a metric of simply L_j . With this assignment, the path of minimal load is selected regardless of the number of hops.

Dynamic routing clearly provides the most opportunity for the path selection process to adapt to the current state of the network. This may appear to be the optimal routing strategy; however, there are important ramifications in an all-optical network. By allowing any path to be selected as opposed to limiting the set of candidate paths to a small set of precalculated paths, connections between a given source/destination pair will tend to follow different paths. This has the effect of decreasing the network *interference length*, which can potentially lead to more contention in the wavelength assignment process. The interference length is the average number of links shared by two paths that have at least one link in common [12.16]. In addition, if the all-optical network makes use of wavebands, where groups of wavelengths are treated as a single unit, then the diversity of paths produced by a purely dynamic strategy can be detrimental from the viewpoint of efficiently packing the wavebands. Note that wavebands may be required with some of the technology that has been proposed for

achieving large fiber capacity in future networks, as is discussed in Sect. 12.12.5 [12.17, 18].

Another drawback to dynamic routing is the additional delay inherent in the scheme. This is especially true in a centralized architecture where any route request must be directed to a single network entity (e.g., a *path computation element* (PCE) [12.19, 20]). Delays incurred due to communication to and from the PCE, as well as possible queuing delays within the PCE, may add tens of milliseconds to the route calculation pro-

cess. For dynamic applications that require very fast connection setup, this additional delay may be untenable.

To summarize this section on routing strategies, studies have shown that alternative-path routing provides a good compromise between adaptability and calculation time, and is more amenable to wavelength assignment in an all-optical network. If all candidate paths between a source/destination pair are blocked, then dynamic routing can be invoked.

12.5 Routing Order

In real-time network operation, connection requests are processed as soon as they are received. Thus, the routing process typically consists of searching for a path for just one connection (if batched routing is employed, then there may be a small number of connections to be routed at one time). In contrast, with long-term network design, routing is likely to be performed on a large traffic matrix comprising hundreds of connections. The order in which the connections are routed may have a significant impact on the ultimate level of network loading and blocking.

An effective ordering approach is to give routing priority to the source/destination pairs for which finding a path is likely to be more problematic. Thus, factors such as the relative locations of the two endpoints, whether or not a disjoint protection path is required for the connection, and whether or not the candidate paths for the connection contain a large number of *hot links* should be considered when assigning the routing order. For example, finding a path for a source/destination pair that is located at opposite ends of a network is likely to be more difficult than finding a path between an adjacent source and destination. It is typically advantageous to route such a connection early in the process when the links are relatively lightly loaded to minimize the

constraints. Similarly, requiring disjoint paths for purposes of protection is already a significant constraint on the possible number of suitable paths. Giving priority to routing connections that require protection may enhance the likelihood that a feasible pair of disjoint paths can be found.

One enhancement of this strategy is to combine it with a metaheuristic such as *simulated annealing* [12.21]. A baseline solution is first generated based on the priority ordering method described above. The ordering that generated the baseline result is passed to the simulated annealing process. In each step of simulated annealing, the orderings of two connections are swapped, and the routing process rerun. If the result is better (e.g., less blocking or lower levels of resulting load), this new ordering is accepted. If the result is worse, the new ordering is accepted with some probability (this allows the process to extricate itself from local minima). The probability threshold becomes lower as the simulated annealing process progresses. If simulated annealing is run long enough, the overall results are likely to improve as compared to the baseline solution.

Other ordering schemes can be found in [12.1]. With rapid design runtimes, several orderings can be tested and the best result selected.

12.6 Multicast Routing

Video distribution is one of the major drivers of traffic growth in networks. Such services are characterized by a single source delivering the same traffic to a set of destinations. One means of delivering these services is to establish multiple unicast connections between the source and each of the destinations. Alternatively, a single multicast tree that includes each of the destinations can be established [12.22]. These two options

are illustrated in Fig. 12.9. The multicast tree eliminates duplicate routing on various links of the network (e.g., link QR in Fig. 12.9) and is thus more capacity efficient. One study in a realistic backbone network showed that multicast provides a factor of roughly two to three improvement in capacity as compared to multiple unicast connections, where capacity was measured as the average number of wavelengths required on a link [12.1].




Fig. 12.9 (a) Four unicast connections are established between the source Q and the destinations W, X, Y and Z. (b) One multicast connection is established between node Q and the four destinations (© Monarch Network Architects LLC)

A tree that interconnects the source to all of the destinations is known as a *Steiner tree* (where it is assumed that all links are bidirectionally symmetric; i.e., two-way links with the same metric in both directions). The *weight* of the tree is the sum of the metrics of all links that compose the tree. As with point-to-point connections, it is often desirable to find the routing solution that minimizes the weight of the tree. Finding the Steiner tree of minimum weight is, in general, a difficult problem to solve unless the source is *broadcasting* to every node in the network. However, several heuristics exist to find good approximate solutions to the problem [12.23]. Two such heuristics are *minimum spanning tree with enhancement* (MSTE) [12.24, 25] and *minimum paths* (MP) [12.26]. Examples of these heuristics, along with their C-code implementation, can be found in [12.1]. Based on studies utilizing realistic backbone networks, MP tends to produce better results, although the relative performance of MSTE improves as the number of destinations increases [12.1]. In the case of broadcast, an algorithm such as Prim's or Kruskal's can be used to optimally find the tree of minimum weight [12.2].

The multicast algorithms enumerated above provide a single path from the source to any of the destinations. If a failure occurs along that path, connectivity with that destination is lost. Furthermore, due to the tree topology, a link failure along the tree often disconnects several destinations. Providing protection for a multicast tree, where connectivity with all destinations is maintained regardless of any single link failure, can be cumbersome and may require a large amount of protection resources.

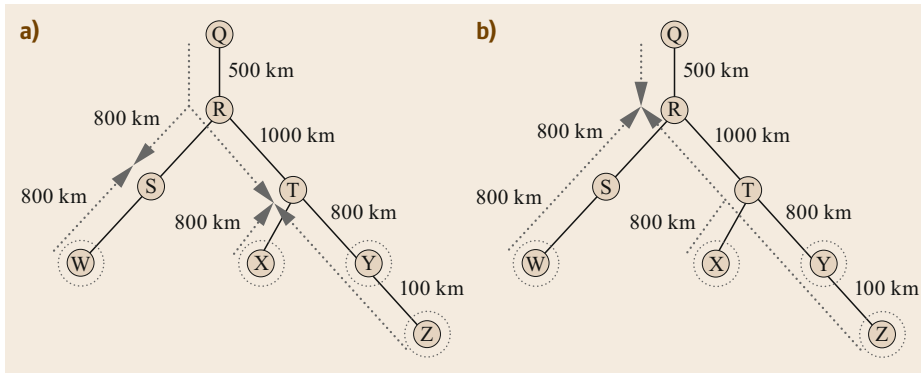
Various strategies have been devised for providing multicast protection [12.27–29]. One approach common to several of these strategies is to make use of segment-based protection, where the multicast tree is conceptually partitioned into segments, and each segment is protected separately. A very different concept that has been applied to multicast protection is *network coding* [12.30–32]. With this approach, which has applications beyond just multicast protection, the

destinations receive independent, typically linear, combinations of various optical signals rather than the individual optical signals that originated at the source. Processing at one or more nodes is required to create these signal combinations. The processing is preferably performed in the optical domain, but electrical processing may be required to generate more complex signal combinations. With proper processing of the received data, the destination can recreate the original signals. For protection purposes, the transmissions are sent over diverse paths and the signal combinations are such that if one signal is lost due to a failure, it can be recovered (almost) immediately from the other signals that are received. To mine the full benefits of network coding, there must be multiple signals that can be advantageously combined, as is often the case when routing a large amount of multicast traffic. Network coding *may* result in a more efficient use of network resources; however, even if the amount of required capacity is approximately the same as in a more conventional shared-mesh restoration approach, the recovery time is typically much faster.

Another important design aspect of multicast routing in an all-optical network is the selection of regeneration sites, where judicious use of regeneration may reduce the cost of the multicast tree. For example, it may be advantageous to favor the branching points of a multicast tree for regeneration. Refer to the tree of Fig. 12.10, where the source is node Q, and the multicast destinations are nodes W, X, Y, and Z. Assume that the optical reach is 2000 km. In Fig. 12.10a, the signal is regenerated at the furthest possible node from the source without violating the optical reach. This results in regenerations at nodes S and T. If, however, the regeneration occurs at the branching point node R as in Fig. 12.10b, then no other regeneration is needed, resulting in a lower cost solution.

12.6.1 Manycast Routing

In one variation of multicast routing, only N of the M destinations must be reached by the multicast tree,

**Fig. 12.10a,b**

Assume that the optical reach is 2000 km. **(a)** Regeneration at both nodes S and T. **(b)** Regeneration only at node R (© Monarch Network Architects LLC)

where $N < M$. The routing goal is still to produce the lowest-weight multicast tree; however, the routing algorithm must incorporate the selection of the N destinations as well. This scenario is known as *manycast*.

Manycast is useful in applications such as distributed computing. For example, there may be M processing centers distributed in the network, but an end user requires the use of only N of the processors. The goal of the end user is typically to minimize latency; it

does not have a preference as to which N data centers are utilized. Optimal, or near-optimal, manycast routing, where the link metric is distance, is useful for such a scenario.

Various heuristic algorithms have been proposed for manycast routing, where the goal is to find the manycast tree of lowest weight. One effective heuristic in particular is a variation of the MP multicast algorithm [12.33].

12.7 Wavelength Assignment

Routing is one critical aspect of network design. Another important component is wavelength assignment, where each routed connection is assigned to a portion of the spectrum centered on a particular wavelength. Each time a connection enters the electrical domain, the opportunity exists (typically) to change the wavelength to which the connection has been assigned. The wavelength assignment process must satisfy the constraint that no two connections can be assigned the same wavelength on a given fiber. If this constraint were to be violated, the connections would occupy the same portion of the spectrum and interfere with each other. Note that when a link is populated with multiple fiber pairs, there can be multiple connections carried on the same wavelength on the link as long as each of the connections is routed on a different fiber.

12.7.1 Interaction Between Regeneration and Wavelength Assignment

As discussed in Sect. 12.1, all traffic that is routed through a node in an O-E-O network is regenerated. Regeneration is most commonly accomplished through the use of two back-to-back *transponders*, as shown in Fig. 12.11. A signal enters and exits a transponder in the optical domain but is converted to the electrical

domain internally. The transponders interface to each other on a common wavelength (typically 1310 nm) through what is known as the *short-reach interface*. At the opposite end, the transponder receives and transmits a WDM-compatible wavelength, where the wavelength is generally in the 1500 nm range. The WDM-compatible wavelengths that are received/transmitted by the two transponders (i.e., λ_j and λ_k in Fig. 12.11) can be selected independently, thereby enabling wavelength conversion. (Wavelengths are also referred to as *lambdas*, with wavelength j represented by λ_j .)

Because of the flexibility afforded by regeneration, wavelengths can be assigned arbitrarily in an O-E-O network as long as each connection on a fiber is




Fig. 12.11 Regeneration through the use of back-to-back transponders. The WDM-compatible signals associated with the two transponders, λ_j and λ_k , do not have to be the same, thereby enabling wavelength conversion (© Monarch Network Architects LLC)




Fig. 12.12 A representative node in an O-E-O network. Wavelengths can be assigned independently to the connections on each link as long as each wavelength assigned on a fiber is unique. The lower-most connection is carried on λ_1 on the West link and carried on λ_4 on the East link (© Monarch Network Architects LLC)

assigned a unique wavelength. This is illustrated in Fig. 12.12, where the lower most connection enters the node from the *West* link on wavelength 1 and exits the node on the *East* link on wavelength 4. In the reverse direction, wavelength 4 is converted to wavelength 1. (By convention, the nodal fibers at a degree-two node are referred to as *West* and *East*; it may have no correlation to the actual geography of the node.) There is no requirement that the wavelength must be changed. As shown in Fig. 12.12, another connection enters and exits the node on λ_2 .

In contrast, signals in an all-optical network are not regenerated at every node along the path. It is only required that a signal be regenerated prior to it traveling a distance that is longer than the system’s optical reach.

If a signal is not regenerated at a node (i.e., it traverses the node in the optical domain), then it is carried on the same wavelength into and out of the node. More generally, if the signal traverses N consecutive nodes in the optical domain, then it is carried on the same wavelength on $N + 1$ links. The wavelength assignment process consists of finding a single wavelength that is available on each of these links. We use the term *all-optical segment* to refer to a portion of a connection that rides in the optical domain without any conversion to the electrical domain.

The difficulty of assigning wavelengths clearly depends on two factors: the number of links in an all-optical segment and the utilization level on those links. As the number of links in an all-optical segment increases, the difficulty in finding a wavelength that is free on each one of the links typically increases as well. Similarly, as the utilization level of a link increases, fewer wavelengths are available, thus making it less likely that a free wavelength can be found on the entire all-optical segment.

This presents an interesting tradeoff. Each regeneration requires the deployment of two transponders, adding to the cost, power consumption, and failure rate




Fig. 12.13 (a) The entire end-to-end connection is carried in the optical domain. The same wavelength must be available on all seven links traversed by the connection. (b) The connection is regenerated at node D, creating two all-optical segments: AD and DZ. These two segments can be assigned a different wavelength. (c) Assume that an available wavelength cannot be found from A to D, but λ_j is available from A to C, and λ_k is available from C to D. An additional regeneration can be added at node C to make wavelength assignment feasible for this connection (© Monarch Network Architects LLC) ◀

of the associated connection. However, wavelength assignment becomes simpler the more regeneration that is present. This is illustrated in Fig. 12.13. Figure 12.13a shows a true all-optical connection that remains in the optical domain from node A to node Z. In order to find a suitable wavelength to assign to this connection, the same wavelength must be available on all seven links that compose this end-to-end path. Figure 12.13b shows that same connection with an intermediate regeneration at node D. The presence of regeneration simplifies the wavelength assignment process. One needs to find a wavelength that is free on the three links from node A to node D and a wavelength that is free on the four links from node D to node Z. There is no requirement that these two wavelengths be the same.

If wavelength assignment fails for a connection, where a suitable wavelength cannot be found for one or more of the all-optical segments composing the connection's path, then at least four design options exist. First, the connection request can be blocked. Second, it may be possible to select different regeneration locations for the connection, thereby producing a different set of all-optical segments. Wavelength assignment may be feasible on this alternate set of segments. For example, in Fig. 12.1c (this figure appeared in Sect. 12.1), the two required regenerations are chosen to be at nodes C and F, producing the all-optical segments AC, CF, and FZ. Alternatively, the two regenerations could be placed at nodes B and E without violating the optical-reach constraint (there are other alternatives as well). This choice produces a completely different set of all-optical segments: AB, BE, and EZ. If moving the regenerations does not produce a feasible wavelength assignment, or if there are no regenerations for the connection, then a third option is to route the connection on a different path and re-attempt the wavelength assignment process on the new path. As a fourth option, the connection can remain on the same path, but one or more regenerations can be added even though they are not required due to optical-reach concerns, thus incurring greater cost. This last option is illustrated in Fig. 12.13c. Assume that a wavelength is available along the DZ segment, but no single wavelength is available on all of the links of the AD segment. Furthermore, assume that λ_j is available on the links from node A to node C, and λ_k

is available on the link between node C and node D. By adding a regeneration at node C (and its attendant costs), wavelength assignment for this connection becomes feasible.

Numerous studies have been performed to study the level of blocking that results due to a failure of the wavelength assignment process (assuming extra regenerations cannot be added to alleviate wavelength contention). The consensus of the majority of these studies is that *sparse* regeneration provides enough opportunities for wavelength conversion, resulting in a relatively low level of blocking due to wavelength contention [12.1, 12, 34, 35]. In a continental-scale network, the regeneration that is required based on optical reach is minimal (e.g., three regenerations in a connection that extends from the East coast to the West coast in a United States backbone network) but sufficient to achieve low levels of wavelength contention. In networks of smaller geographic extent, e.g., metro networks, the optical reach may be longer than any end-to-end connection such that regeneration is not required. However, there are other network functionalities that limit the extent of any all-optical segment. For example, low-data-rate traffic may need to be processed periodically by a grooming switch or router to make better use of the fiber capacity. Currently, grooming devices operate in the electrical domain such that the grooming process concurrently regenerates the signal. Thus, sparse grooming translates to sparse regeneration, which, in turn, allows sparse wavelength conversion. The net effect is that wavelength contention can remain low in a range of networks. Furthermore, it has been shown that just a small amount of extra regeneration effectively eliminates wavelength contention in a typical network [12.35].

It should be noted that some studies appear to indicate that wavelength contention is a major problem that results in excessive blocking. Further investigation of the details of these studies may reveal flaws in the underlying assumptions. For example, the study may not take advantage of regeneration as an opportunity to convert the wavelength. Another possible weakness is that effective wavelength assignment algorithms may not be employed. Wavelength assignment algorithms are covered in the next section.

12.8 Wavelength Assignment Algorithms

In this section, we assume that routing and wavelength assignment are two separate steps; Sect. 12.9 considers integrated approaches. It is assumed that one or more new connection requests are passed to the net-

work design process. In the first step, each connection in the set is routed. Once a path has been selected for a connection, the required regeneration locations along that path are determined. This yields a set of

all-optical segments (i.e., the portions of the paths that lie between an endpoint and a regeneration or between two regenerations). The list of all-optical segments is passed to the wavelength assignment algorithm, which is responsible for selecting a feasible wavelength for each segment. In order for a wavelength assignment to be feasible, any two all-optical segments that are routed on the same fiber at any point along their respective paths must be assigned different wavelengths. Furthermore, the wavelength assignment must be compatible with the technology of the underlying optical system. An optical system supports a limited number of wavelengths on a fiber, thereby placing a bound on the number of different wavelengths that can be used in the assignment process. For example, a backbone network system may support 80 wavelengths on a fiber, whereas a metro network deployment may support only 40 wavelengths (there are typically fewer wavelengths needed in a metro network and 40-wavelength technology is of lower cost than 80-wavelength technology).

The wavelength assignment problem is analogous to the graph coloring problem, where each node of a graph must be assigned a color subject to the constraint that any two nodes that are adjacent in the graph topology must be assigned different colors. The objective is to color the graph using as few colors as possible. To elucidate the analogy with wavelength assignment, let each all-optical segment in a network design correspond to a node in the coloring graph. Links are added between any two nodes of the coloring graph if the paths of the two corresponding all-optical segments have any fibers in common. The resulting graph is known as the *conflict graph*, which is illustrated for a small example with four all-optical segments in Fig. 12.14. (This graph is also called the *auxiliary*

graph.) Solving the graph coloring problem on this graph produces the wavelength assignment; i.e., the color assigned to a node represents the wavelength assigned to the all-optical segment corresponding to that node. The resulting wavelength assignment is such that the minimum number of wavelengths is used.

There are no known polynomial-time algorithms for optimally solving the graph coloring problem for general instances of the problem. This implies that there are no corresponding efficient algorithms that can optimally solve general instances of the wavelength assignment problem. Thus, heuristic algorithms are typically used. The heuristics encompass two aspects: (1) generating the order in which the all-optical segments are assigned a wavelength and (2) selecting a wavelength for each of the segments. As with routing order, there are many strategies that can be used for ordering the all-optical segments. Some of the strategies developed for ordering the nodes in a graph coloring can readily be extended to this problem, most notably the *Dsatur* strategy [12.36].

Developing heuristics to select which wavelength to assign to an all-optical segment is a well-researched topic, and numerous such heuristics have been proposed [12.37]. They differ in factors such as complexity and the amount of network-state information that needs to be monitored. In spite of the array of proposals, two of the simplest heuristics, both proposed in the very early days of optical-network research, remain the algorithms most commonly used for wavelength assignment. These heuristics are first-fit and most-used [12.38], described in further detail below. Both of these algorithms are suitable for any network topology and provide relatively good performance in realistic networks. For example, wavelength contention does not generally become an issue until there are at least a few links in the network with roughly 85% of the wavelengths used. Whether first-fit or most-used performs better for a particular network design depends on the network topology and the traffic. In general, the differences in performance are small. One advantage of first-fit is that, in contrast to most-used, it does not require any global knowledge, making it more suitable for distributed implementation.

Either of the schemes can be applied whether there is a single fiber-pair or multiple fiber-pairs on a link. Note that there are wavelength assignment schemes specifically designed for the multiple fiber-pair scenario, most notably the *least-loaded* scheme [12.12]. This has been shown to perform better than first-fit and most-used when there are several fiber-pairs per link [12.37]. As fiber capacities have increased, however, systems with several fiber-pairs on a link have become a less common occurrence.




Fig. 12.14 (a) Four all-optical segments, each one numbered, are routed as indicated by the *dotted lines*. It is assumed that there is one fiber-pair on each link. (b) The resulting conflict graph, where each node represents one of the all-optical segments. A link exists between two nodes if the corresponding all-optical segments are routed on the same fiber on any link in the original graph

12.8.1 First-Fit Algorithm

In the first-fit algorithm, each wavelength is assigned an index from 1 to W , where W is the maximum number of wavelengths supported on a fiber. No correlation is required between the order in which a wavelength appears in the spectrum and the index number assigned. The indices remain fixed as the network design evolves. Whenever wavelength assignment is needed for an all-optical segment, the search for an available wavelength proceeds in an order from the lowest index to the highest index. The first wavelength that is available on all links that compose the all-optical segment is assigned to the segment. First-fit is simple to implement and requires only that the status of each wavelength on each link be tracked.

Due to the interdependence of wavelength assignment across links, the presence of failure events, and the presence of network churn (i.e., the process of connections being established and then later torn down), the *indexing ordering* does not guarantee the actual *assignment order* on a particular link. Thus, relying on the indexing scheme to enforce a particular assignment

ordering on a link for performance purposes is not prudent; for more details, see [12.1].

12.8.2 Most-Used Algorithm

The most-used algorithm is more adaptive than first-fit but requires more computation. Whenever a wavelength needs to be assigned to an all-optical segment, a wavelength order is established based on the number of link-fibers on which each wavelength has already been assigned in the network. The wavelength that has been assigned on the most link-fibers already is given the lowest index, and the wavelength that has been assigned on the second-most link-fibers is given the second lowest index, etc. Thus, the indexing order changes depending on the current state of the network. After the wavelengths have been indexed, the assignment process proceeds as in first-fit. The motivation behind this scheme is that a wavelength that has already been assigned on many fibers will be more difficult to use again. Thus, if a scenario arises where a heavily-used wavelength can be used, it should be assigned.

12.9 One-Step RWA

When routing and wavelength assignment are treated as separate steps in network design, it is possible that the routing process produces a path on which the wavelength assignment process fails (assuming extra regenerations are not added to alleviate the encountered wavelength contention). Alternatively, one can consider integrated routing and wavelength assignment methodologies such that if a path is selected, it is guaranteed to be feasible from a wavelength assignment perspective as well.

Various one-step RWA methodologies are discussed below, all of which impose additional processing and/or memory burdens. When the network is not heavily loaded, implementing routing and wavelength assignment as independent steps typically produces feasible solutions. Thus, under these conditions, the multistep process is favored, as it is usually faster. However, under heavy load, using a one-step methodology can provide a small improvement in performance [12.1]. Furthermore, under heavy load, some of the one-step methodologies may be more tractable, as the scarcity of free wavelengths should lead to lower complexity.

12.9.1 Topology Pruning

One of the earliest advocated one-step algorithms starts with a particular wavelength and reduces the network

topology to only those links on which this wavelength is available. The routing algorithm (e.g., a shortest-path algorithm) is run on this pruned topology. If no path can be found, or if the path is too circuitous, another wavelength is chosen and the process run through again on the corresponding pruned topology. The process is repeated with successive wavelengths until a suitable path is found. With this approach, it is guaranteed that there will be a free wavelength on any route that is found. If a suitable path cannot be found after repeating the procedure for all of the wavelengths, the connection request is blocked.

In a network with regeneration, using this combined routing and wavelength assignment procedure makes the problem unnecessarily more difficult because it implicitly searches for a wavelength that is free along the whole extent of the path. As discussed earlier, it is necessary to find a free wavelength only along each all-optical segment, not along the end-to-end connection. One variation of the scheme is to select ahead of time where the regenerations are likely to occur for a connection and apply the combined routing and wavelength assignment approach to each expected all-optical segment individually. However, the route that is ultimately found could be somewhat circuitous and require regeneration at different sites than where was predicted,

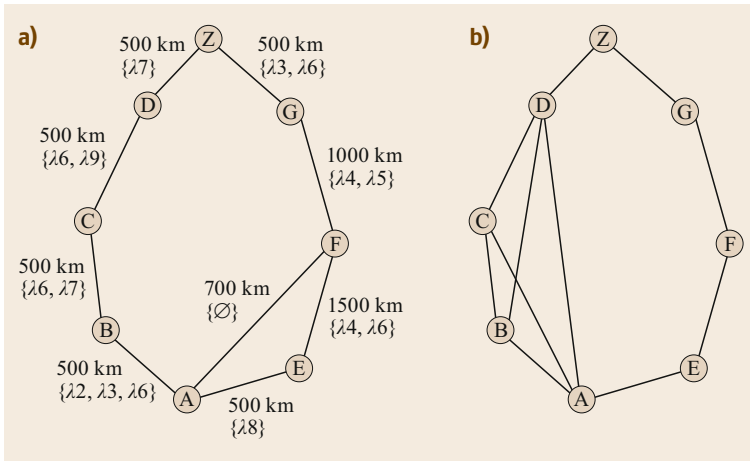


Fig. 12.15 (a) The true network topology, where it is assumed that the optical reach is 2000 km. The wavelengths listed next to each link are the wavelengths that are assumed to be free on the link. (b) The reachability graph, where a link is added between a node pair if there is a regeneration-free path between the nodes with at least one available wavelength along the path (© Monarch Network Architects LLC)

so that the process may need to be run through again. Overall, this strategy is less than ideal.

12.9.2 Reachability-Graph Transformation

A more direct integrated RWA approach is to create a transformed graph whenever a new connection request needs to be routed and assigned a wavelength, where the newly-formed graph represents feasible all-optical segments; i.e., feasible with respect to both optical reach and available wavelengths [12.1, 39]. The process to form this graph, specifically selecting the nodes that are added to this graph, depends on the timeframe of the network design. In long-term network planning, it is assumed that there is time to deploy any equipment that may be required for a design. Under these planning conditions, every network node of the real network topology appears in the transformed graph. In contrast, with real-time network operation, the required connection setup times do not allow for equipment to be deployed; only equipment that is already available in the network can be utilized for the new traffic. In this scenario, only nodes with available regeneration equipment, plus the source and the destination of the new connection request, are added to the transformed graph.

After adding the nodes to the transformed graph, a link is added between a pair of nodes in this graph only if there exists a regeneration-free path between the nodes in the true topology *and* there exists at least one wavelength that is available along the path. Even if there are multiple regeneration-free paths between a node pair, or multiple wavelengths free on a path, at most one link is added between a node pair. This transformed graph is referred to as the *reachability graph*.

An example of such a graph transformation is shown in Fig. 12.15. The true topology is shown in

Fig. 12.15a, where the wavelengths that are assumed to still be available on a link are shown. The connection request is assumed to be between nodes A and Z, and the optical reach is assumed to be 2000 km. Furthermore, it is assumed that this is a long-term network design exercise and regeneration is permitted at any node. Thus, all nodes appear in the reachability graph. With these assumptions, the corresponding reachability graph is shown in Fig. 12.15b. All of the original links appear in this graph, except for link AF, which has no available wavelengths. In addition, links AC, AD, and BD are added because the respective associated paths, A-B-C, A-B-C-D, and B-C-D, are less than 2000 km and have a free wavelength (i.e., on each of these paths λ_6 is free). No link is added to represent the path E-F-G, even though λ_4 is available on this path because the path distance is 2500 km, which is longer than the optical reach.

In a real network with many nodes and wavelengths, creating the reachability graph can potentially be time consuming. A list of all node pairs where the shortest path between the nodes is less than the optical reach is maintained. These node pairs represent the possible all-optical-segment endpoints. At the time of a new connection request, a search is performed for each node pair in this list to find a regeneration-free path between the two nodes where some wavelength is available along the whole path. This determines whether a link is added between the two nodes in the reachability graph.

One strategy to do this is the topology-pruning approach described in Sect. 12.9.1, where the true topology is reduced to just those links that have a particular wavelength free. A shortest-path algorithm between the pair of nodes (i.e., the potential endpoints of an all-optical segment) is run on the pruned topology. If the distance of the resulting path is less than the optical reach, then a suitable regeneration-free path has been

found, and a link is added in the reachability graph between the node pair. If the path distance is greater than the optical reach, or no path can be found, then the process must be repeated for another wavelength. Many (if not all) of the wavelengths may need to be considered.

An alternative strategy is to run a K -shortest-paths algorithm on the true topology, where K is large enough such that *all* regeneration-free paths between a node pair of interest are found, and the resulting list is stored. At the time of a new connection request, the paths are checked for a free wavelength. If any of the paths have a free wavelength along the whole path, then a link is added between the node pair in the reachability graph. (To speed up the process, one could consider just a subset of the possible regeneration-free paths between a node pair; there may be some loss of optimality, but it is likely to be small.)

Once the reachability graph is formed, a shortest-path algorithm is run from the connection source to the connection destination to find the path in the reachability graph with the fewest hops (i.e., all link metrics are set to 1). Each hop corresponds to an all-optical segment in the true topology. If a path is found, then it is guaranteed to have the fewest number of feasible regenerations, and each resulting all-optical segment is guaranteed to have an available wavelength. In the example of Fig. 12.15, path A–D–Z is found, which corresponds to the all-optical segments A–B–C–D and D–Z in the true topology. These segments are assigned λ_6 and λ_7 , respectively. Further subtleties of the reachability-graph methodology can be found in [12.1].

12.9.3 Flow-Based Methods

Global optimization techniques, such as *integer linear programming* (ILP), can be applied to the one-step RWA problem as well. ILPs typically consider the whole solution space to find the optimal solution. However, ILP methodologies often have a long runtime and are impractical except for small networks with little traffic. A more practical approach is to use efficient *linear programming* (LP) techniques (e.g., the Simplex algorithm), combined with strategies that drive the solution to integer values. Relaxing the integrality constraints enables more rapid convergence. Various techniques are applied to ultimately produce a (possibly nonoptimal) integer solution.

For example, routing a set of traffic connections can be formulated as a *multicommodity flow* (MCF) problem, where each source/destination pair in the traffic set can be considered a different commodity that needs to be carried by the network [12.40–42]. Additional variables and constraints are needed to enforce wavelength continuity. An integer solution to the problem

is typically desired, which corresponds to routing each connection over just one path, using a single wavelength on a link. The integrality constraints are relaxed in the LP approach to make the problem more tractable. Despite not enforcing integer solutions, the LP can be combined with various perturbation and rounding techniques to improve the likelihood that an integer solution is found [12.41, 42]. However, even with LP relaxation techniques, there may be a greater-than-linear increase in runtime as the number of connections increases [12.42].

One approach to speed up the process is to input a set of possible paths that can be utilized by a connection between any given source and destination. This is analogous to calculating a set of candidate paths for alternative-path routing. Restricting the LP to a set of candidate paths, as opposed to allowing the LP to freely select the paths, may result in a less than optimal solution; however, with a good choice of candidate paths, the effect should be small. Another benefit to preselecting the paths is that the regeneration sites can be selected up front. This allows the wavelength continuity constraint to be specified on a per-all-optical-segment basis rather than requiring that wavelength continuity be enforced end-to-end.

Ideally, the cost function that is used in the LP encourages load balancing. Also, it is preferable if the cost function is input as a piecewise linear function with integer breakpoints as another means of pushing the LP towards an integer solution.

As suggested earlier, using a one-step RWA approach such as an LP methodology may be more expedient when adding connections to a highly loaded network. At that stage, there are few available wavelengths on each link, such that the solution space is much smaller. This should allow the LP to converge more quickly.

It is interesting to compare the results of the one-step LP-based RWA approach to those of a multistep approach, where an LP methodology is used just for the routing portion and a commonly-used graph coloring algorithm is used for wavelength assignment. The performances have been shown to be similar (depending on the cost functions used in the LPs), indicating that good results can be obtained using the simpler multistep approach [12.42]. The runtime of the multistep approach was an order of magnitude faster.

12.9.4 ILP-Based Ring RWA

Although ILP formulations have generally been considered too slow for practical RWA, a scalable ILP methodology has been proposed for ring topologies [12.43]. This methodology includes a decomposi-

tion approach that is optimal, fast for any reasonably sized ring, with a runtime that is essentially independent of the amount of traffic on the ring. This ILP methodology is, thus, a scalable one-step RWA ap-

proach for realistic ring problem instances. Further research is needed to determine whether the decomposition procedure can be extended to arbitrary mesh topologies.

12.10 Impairment-Aware Routing and Wavelength Assignment

The transmission of an optical signal is subject to an array of impairments that affect the signal quality and where the signal must be regenerated. There has been a large research effort in the area of *impairment-aware routing and wavelength assignment* (IA-RWA), leading to a number of proposed link metrics and design methodologies that account for various impairments during the routing and regeneration processes [12.44, 45].

One of the major impairments that an optical signal encounters is accumulated noise. The strength of the signal compared to the level of the noise is captured by the signal's OSNR, where signals with lower OSNR are more difficult to receive without errors. Many other optical impairments arise from the physical properties of light propagating in a fiber. For example, the propagation speed of light within a fiber depends on the optical frequency. This causes the optical signal pulses, which have a finite spectral width, to be distorted as they propagate along a fiber. This phenomenon is known as *chromatic dispersion*, or simply *dispersion*. Dispersion is a linear impairment that can typically be managed (e.g., through the use of dispersion-compensating fiber or the use of coherent technology) such that its detrimental effect is limited. Furthermore, the presence of some level of dispersion can be helpful in mitigating the effect of other optical impairments [12.46–48].

There are numerous nonlinear optical impairments that are more difficult to manage [12.49–52]. Several of these optical effects arise as a result of the fiber refractive index being dependent on the optical intensity. (The refractive index governs the speed of light propagation in a fiber.) As the optical signal power is increased, these nonlinearities become more prominent. One such nonlinearity is *self-phase modulation* (SPM), where the intensity of the light causes the phase of the optical signal to vary with time. This potentially interacts with the system dispersion to cause significant pulse distortion. *Cross-phase modulation* (XPM) is a similar effect, except that it arises from the interaction of two signals, which is more likely to occur when signals are closely packed together in the spectrum. Another nonlinear effect is *four-wave mixing* (FWM). This arises when signals carried on three particularly spaced

optical frequencies interact to yield a stray signal at a fourth frequency, or two frequencies interact to generate two stray signals. These stray signals can potentially interfere with the desired signals at or near these frequencies.

In many all-optical networks, the transmission system is designed such that the power levels are low enough, or the dispersion levels are high enough, so that impairments due to adjacently propagating wavelengths are relatively small. However, there may be transmission systems where relatively high power levels are required, leading to scenarios where populating adjacent, or nearly adjacent, wavelengths in the spectrum produces non-negligible nonlinear impairments, most notably XPM. In such systems, the quality of transmission (QoT) for a given connection may depend on which other wavelengths are in use on the same fibers.

There are two methods for dealing with this scenario. The first strategy is to ensure that connections are established with enough system margin to tolerate the *worst-case* impairments that could possibly arise from populating adjacent wavelengths with other connections. This allows wavelengths to be assigned to connections without concern over interwavelength impairments. If a particular connection is deemed feasible at the time of its establishment, it should remain feasible regardless of what other connections may later be added.

In the second strategy, the effects of interwavelength impairments are calculated more precisely. The optical reach of a particular available wavelength along a given path is determined at the time a connection request is received, based on the state of the adjacent wavelengths. Consider assigning wavelength i to a new connection on a given path. If wavelengths $i - 1$ and $i + 1$ are not being used on the fibers that compose this path, then wavelength i may have additional optical reach, as compared to the case where a worst-case reach assumption is used. This could lead to fewer required regenerations for the new connection. The drawback is that if future connections populate wavelength $i - 1$ and/or $i + 1$, the performance of wavelength i may degrade below an acceptable QoT, forcing the associated connection to be assigned to a different wavelength or be rerouted, which is undesirable. If such a modifica-

tion to an existing connection is not permitted, then the strategy of maximizing the reach of wavelength i could result in future connections being blocked from using wavelengths $i - 1$ or $i + 1$.

The decision as to which of the two strategies to use may depend on how regenerations are handled. If regeneration is permitted in the network, then the impact of using a worst-case impairment assumption will likely be extra regenerations, because the system optical reach will effectively be reduced. However, as long as the optical reach is still reasonably long, small reductions in the reach (e.g., 2500 km versus 2800 km) do not have a large impact on the amount of regeneration required and the overall network cost [12.1].

If, however, the system requires that connections be *truly* all-optical, with no regeneration, then the policy for handling impairments may have an impact on blocking. For example, the end-to-end path distance of a new connection may be very close to the nominal optical reach. Establishing the new connection on a wavelength that is distant from any populated wavelengths may allow the connection to be successfully deployed, whereas the worst-case impairment assumption would dictate that it be blocked. This effect was examined more fully in [12.53] for a backbone network of relatively small geographic extent, where no regeneration was permitted. The two strategies outlined above were compared; i.e., either assume worst-case interwavelength impairments or calculate the interwavelength impairments more accurately based on the actual network state. In either strategy, moving an existing connection to a different path or wavelength was not permitted. The results indicated that when interwavelength impairments were more precisely calculated, the blocking rates were reduced by about an order of magnitude, due to there being a larger set of feasible paths from which to choose. However, this type of pure all-optical scenario would not arise in a network of large geographic extent, because some regeneration is needed regardless of how interwavelength impairments are treated. Interwavelength impairments are also unlikely to be an issue in a metro network, where the optical reach, even with worst-case assumptions, is typically longer than any path. Thus, the benefit of more precisely calculating interwavelength effects may not be significant in many practical networks.

Nevertheless, strategies have been developed to take interwavelength impairments into account when performing RWA for a new connection request; for example, using a cost-vector approach to routing [12.54]. Various per-wavelength components are included in the cost vector that is used for shortest-path routing, where the vector captures interwavelength impairments such as XPM and FWM. For each available wavelength on

a link, the cost component for that wavelength-link combination is calculated based on the wavelengths that are already populated on that link. A modified Dijkstra routing algorithm is run with the cost vector, using the principle of dominated paths. (A path between two nodes that has all of its cost metrics higher than those of another path between the same two nodes is considered *dominated*.) Multiple nondominated paths from source to intermediate nodes are tracked, and any dominated paths are eliminated from further consideration. When the routing algorithm terminates, a scalar-generating function is applied to the final cost vector for each remaining feasible path/wavelength combination to determine which one to use.

Machine learning (as part of a *cognitive network*) has also been proposed for estimating the QoT of a new connection [12.55–57]. In one such approach, a database is maintained for a set of paths for which the QoT is known (through prior analysis, experimentation, and/or performance monitoring of live connections in the network). Each of these paths is characterized by a set of metrics (e.g., path distance, assigned wavelength, modulation format). When a potential new connection (or all-optical segment) is being evaluated, the paths in the database that are most similar to it are used to determine whether the QoT will meet the system threshold. Maintaining a proper-sized database is important to achieve the proper balance between accuracy and computation time.

12.10.1 Mixed Line-Rate Systems

An important scenario that may warrant accounting for interwavelength impairments more precisely is when multiple line rates co-propagate on a single fiber. For example, a single fiber may carry wavelengths that have been assigned to a combination of 10 and 40 Gb/s connections. The different modulation formats that are typically used for these connection rates may have a negative impact on each other. Experiments have shown that 10 Gb/s signals may have an especially detrimental effect on near-by co-propagating 40 Gb/s signals due to XPM [12.58, 59]. Furthermore, the performance penalties are severe enough that leaving enough system margin to account for the worst-case XPM would be too detrimental to the system reach. (The performance penalty induced by 10 Gb/s signals on near-by co-propagating 100 Gb/s signals is less severe; similarly, co-propagating 40 and 100 Gb/s connections are not problematic.)

To deal with this situation, a *soft* partitioning can be enforced in the wavelength assignment process, where the 40 Gb/s wavelengths are assigned from one end of the spectrum and the 10 Gb/s wavelengths are as-

signed starting at the other end. Additionally, relatively short 40 Gb/s connections, which can tolerate the performance penalty of adjacent 10 Gb/s wavelengths, can be proactively assigned wavelengths from the buffer area between the two portions of the spectrum. Note that a *fixed* partitioning of resources between the line rates is not advocated, as fixed partitioning generally leads to more blocking.

12.11 Flexible (Elastic) Optical Networks


WDM systems have historically utilized a standardized wavelength grid alignment. For example, since the 2000 timeframe, WDM *backbone* networks have typically employed a grid where the wavelengths are spaced at 50 GHz intervals, and each channel is assigned a fixed 50 GHz of spectrum, as shown in Fig. 12.16a. While the spacing has remained fixed at 50 GHz, advancements in technology have enabled the capacity of a wavelength (i.e., the line rate) to increase, rising from 2.5 to 100 Gb/s over an approximately 20-year span.

A fixed wavelength grid in combination with higher line rates has resulted in a corresponding increase in network capacity. However, it has also resulted in a greater disparity between the line rate and the rate of the client services being carried on the wavelengths. For example, there may be a significant amount of Gigabit (1 Gb/s) Ethernet services that must be carried in a network that utilizes 100 Gb/s wavelengths. In order to efficiently utilize the bandwidth of each wavelength, it is necessary to carry multiple services on one wavelength. The packing of services onto a wavelength is typically accomplished through a process known as *grooming* [12.60, 61]. Grooming is generally performed in the electrical domain using, for example, synchronous optical network/synchronous digital hierarchy (SONET/SDH) switches, optical transport network (OTN) switches, or internet protocol (IP) routers. While effective at packing the wavelengths, electronic grooming switches and routers pose major challenges in cost, size, and power consumption.

Fig. 12.16 (a) Fixed grid with 50 GHz spectral widths and spacing. **(b)** Gridless architecture with arbitrary spectral widths and spacing. In a practical implementation, the spectrum cannot be partitioned arbitrarily. Rather, the spectrum is likely to be logically divided into fixed-sized fine-granularity slots. Each optical channel is allocated the number of slots that it requires, as opposed to a completely unquantized amount of spectrum. This *minigrid* approach is illustrated in Fig. 12.17 ►

This wavelength-assignment strategy will have a tendency to segregate the conflicting line rates to minimize the performance penalties. As the network fill rate increases, and the high and low spectral ranges approach each other, the cost-vector RWA approach discussed above could be used to capture the penalties associated with adding a particular wavelength of a particular rate to a given link.

In order to eliminate or reduce the need for electronic grooming, various optical-domain grooming schemes have been proposed, such as optical packet switching (OPS) [12.62] and optical burst switching (OBS) [12.63]. However, these schemes have challenges of their own that have prevented their widespread adoption. One optical-domain grooming scheme that has gained traction is the *spectrum-sliced elastic* (SLICE) optical path architecture, originally proposed in 2008 [12.64–66]. In contrast to most grooming schemes, SLICE grooms in the frequency domain, not the time domain, thereby avoiding problematic time-based con-



attention issues (electronic versus optical, and time versus frequency are separate dichotomies). The central tenet of the SLICE approach is that network client services are allocated the amount of optical spectrum that they require; i.e., the capacity of a fiber is *sliced* arbitrarily to match the requirements of the clients. (Practical limitations to this vision are discussed below.) For example, a 40 Gb/s client service is assigned 20 GHz of spectrum (assuming a spectral efficiency of 2 b/s/Hz). This approach eliminates the notion of wavelengths with fixed line rate located on a fixed grid. Networks that adhere to the SLICE approach are considered flexible, or gridless, networks. An example of this flexibility is illustrated in Fig. 12.16b, where *optical channels* of various spectral widths have been allocated.

Note that while an optical channel may be allocated less than 50 GHz worth of spectrum in order to efficiently carry a low-rate client service, it may also be allocated more than 50 GHz of spectrum. Thus, this technology can efficiently carry high-rate services as well, e.g., 400 Gb/s or 1 Tb/s clients. (It is expected that rates of 400 Gb/s and greater are likely to require more than 50 GHz of spectrum. Thus, the standardized wavelength grid was updated in 2012 to offer more flexibility. The grid plan supports any mix of wavelength spacings on one fiber, as long as each wavelength aligns with a 6.25 GHz grid, and the bandwidth assigned to each wavelength is a multiple of 12.5 GHz [12.67].)

Another component of the SLICE architecture is the ability to increase or decrease the amount of spectrum allocated to an optical channel. (An increase in spectral width requires that there be free spectrum available for expansion.) For example, a SONET/SDH-based client may dynamically adjust its service rate through the use of the *link capacity adjustment scheme* (LCAS) [12.68]. A SLICE-based network can correspondingly adjust the amount of spectrum allocated to that service. Overall, the flexibility engendered by the SLICE approach has resulted in such networks being referred to as *elastic optical networks* (EONs), which is the terminology that is used in the remainder of this chapter.

Various enabling technologies are needed to implement an EON, as briefly introduced here. First, the transmission technology must enable the deployment of tightly-packed, variable-sized optical channels. One transmission technique that has emerged as a leading candidate for EONs is the optical analog of *orthogonal frequency-division multiplexing* (OFDM) [12.69–73]. With OFDM, the optical signal is carried on a number of low-rate carriers. By increasing or decreasing the number of carriers, the bandwidth of the aggregate optical signal is modified accordingly. One advantageous property of OFDM is that, due to the lower rate of the constituent carriers, there is a greater tolerance to many

fiber-based impairments. Additionally, the speed of the underlying electronics can be lower. There are other possible transmission techniques suitable for EONs, including *Nyquist-WDM* [12.74, 75].

The transmit/receive technology must also be capable of handling variable-rate, fine-granularity optical channels [12.76]. Software-controlled *bandwidth-variable transponders* (BVTs) have been developed for this purpose. Additionally, as analyzed in [12.1], the number of BVTs required in an EON network is potentially large, especially if there are numerous narrow-bandwidth optical channels routed in the network. One proposal to address this is a BVT that can be sliced into several *virtual transponders*, each of which serves one optical channel [12.76].

The network switches that route the optical channels must be compatible with the EON model. More specifically, the filter technology of such switches must be compatible with variable-granularity optical channels, where the filter shape and bandwidth can be set remotely via software [12.77–79]. This can be implemented with, for example, liquid crystal on silicon (LCoS) technology [12.80–82].

The limitations of the filtering technology impose practical implementation restrictions on EONs. First, the bandwidth of an optical channel cannot be arbitrarily fine. There is a minimum granularity bandwidth on which filters can efficiently operate (the narrower the filter, the more it deviates from the ideal sharp-edged shape). Thus, practical EONs are likely to be quasi-gridless, not totally gridless. For example, the spectrum may be divided into 320 frequency *slots*, where each slot is 12.5 GHz in bandwidth (as opposed to, for example, 80 wavelengths, each with 50 GHz of bandwidth). Each optical channel is allocated the number of slots that it requires, as opposed to a completely unquantized amount of spectrum. (This type of system is sometimes referred to as having a *mini-grid* [12.83].) One implication is that some amount of electronic grooming is still required to efficiently fill the slots.

A second limitation is that the optical channels cannot be immediately adjacent to each other. A *guardband* (likely one slot of bandwidth) is needed between the optical channels to allow the switch filters to operate on a single channel without affecting the neighboring channels. The need for guardbands, which represent wasted bandwidth, negatively impacts the utilization efficiency of the fiber capacity.

With these two limitations, an EON spectrum assignment might look as shown in Fig. 12.17. The portion of the spectrum shown has been partitioned into 30 slots. There are 4 allocated optical channels, of size 5, 3, 8, and 4 slots (left to right). A one-slot guardband is assumed.




Fig. 12.17 In this depiction, the spectrum is divided into 30 spectral slots. Four optical channels are assigned on the fiber, as indicated by the *shaded boxes*; *G* indicates a guardband slot (© Monarch Network Architects LLC)

Even with these practical issues, EONs have the *potential* to: reduce, but not eliminate, the amount of electronic grooming; use bandwidth more efficiently, assuming the number of required guardbands is not

excessive; and be better suited to carry a wide range of services; i.e., both low-rate and high-rate services. EONs are being aggressively pursued as an enabling technology for future networks.

12.12 Routing and Spectrum Assignment in Elastic Optical Networks

Just as routing and wavelength assignment (RWA) is a key component of the network design process in traditional grid-based wavelength networks, *routing and spectrum assignment* (RSA) is fundamental to the design of EONs [12.84, 85]. Spectrum assignment corresponds to assigning a particular set of slots to an optical channel. The restrictions are analogous to those of wavelength assignment. First, a slot on a fiber can be assigned to only one optical channel at any point in time. Second, the slots assigned to an all-optical segment must be the same along that segment; i.e., there is a *continuity constraint*. Additionally, the slots assigned must be contiguous; this *contiguity constraint* does not exist for traditional wavelength assignment (this constraint is relaxed when multipath routing is supported in the EON, as is detailed in Sect. 12.12.6).

Algorithm scalability is of greater concern with EONs because of the increased complexity. For example, there may be 320 slots to track in an EON as opposed to 80 wavelengths in a grid-based network. Furthermore, the contiguity constraint poses a significant additional challenge. Finding a path with enough available bandwidth to carry a new optical channel is not sufficient; the bandwidth must be contiguous (i.e., the slots must be consecutive). The notion of spectral fragmentation becomes an important metric that must be monitored. Because of the additional complexity, treating routing, regeneration, and spectrum assignment as separate steps is likely the approach to be used, although single-step methodologies have been proposed as well, as noted below.

With long-term network planning in an EON, a number of connection requests undergo the RSA process at one time. As with long-term RWA (Sect. 12.5), the order in which connections are routed and/or assigned spectrum can affect the ultimate network performance. For example, it is typically better to start the

routing process with the connections where selecting a route is more challenging; e.g., connections that require longer paths and/or more bandwidth, or where the possible paths for a connection include more hot links. Furthermore, as in RWA, metaheuristics, such as simulated annealing, can be used to adjust the ordering to improve upon the solution [12.86, 87].

12.12.1 Routing

With respect to routing, the same options exist as for RWA, i.e., fixed-path routing, alternative-path routing, and dynamic routing. Fixed-path routing is generally undesirable due to the resulting load imbalances. Dynamic routing typically leads to several different paths being chosen between a given source and destination. This is especially problematic with spectrum assignment (SA), where greater freedom in selecting paths is likely to lead to more spectral fragmentation; i.e., it is preferable to assign spectrum along the same link sequences, so that contiguous blocks of spectrum remain free on the links.

Thus, alternative-path routing is typically favored for RSA, as it is for RWA. A set of candidate paths is calculated for each relevant source/destination pair, where the paths in the set provide diversity with respect to the expected hot links of a network. With alternative-path routing in RWA, link load is typically used to select one of the candidate paths for a new connection request. With RSA, link load does not tell the whole story. The amount of fragmentation that results from the selection of a particular route may be more important [12.88].

Two types of fragmentation arise in EONs: *vertical fragmentation* corresponds to the spectral fragmentation of a single link, whereas *horizontal* (or *spatial*) *fragmentation* considers the alignment of the available spectrum on adjacent links [12.89, 90]. In contrast, tra-

ditional wavelength-based networks suffer from just horizontal fragmentation.

Various measures have been proposed to capture both types of fragmentation. One such metric is *utilization entropy* (UE) [12.91]. Let a slot that is occupied by an existing optical channel be represented by a 1 and an unoccupied slot be represented by a 0 . For any two consecutive slots on a fiber, a status change occurs when a 1 is followed by a 0 , or vice versa. The UE for a link is defined as (total number of status changes on the link)/(total number of slots on the link -1) yielding a normalized value between 0.0 and 1.0. Higher UE indicates a greater degree of fragmentation. UE can also be calculated for a path, where the average number of status changes per slot across any two consecutive links of the path is considered.

UE essentially captures the number of spectral gaps; however, it does not explicitly consider the size of those gaps. Larger spectral gaps are more useful as they can accommodate traffic of higher bandwidth. One metric to address this compares the size of the largest available gap to the sum of the sizes of all available gaps [12.92]. To capture available gap size more formally, a Shannon-entropy fragmentation metric has been proposed [12.93]

$$H_{\text{frag}} = - \sum_{i=1}^N \frac{D_i}{D} \ln \frac{D_i}{D},$$

where D is the total number of slots on a link, and D_i is the number of slots in the i -th block (a block is a consecutive sequence of slots that are all 1s or all 0s; i.e., a block is either fully utilized or fully available). Higher values of H_{frag} indicate more fragmentation. This metric can be extended to paths by performing a bitwise-OR of the occupancy value of a slot along each link of the path. Interestingly, simulations showed that considering the metric for each link of a path produced better results than the path-based metric [12.93]. This is largely due to the coarse nature of the bitwise-OR operation (i.e., a single 1 in a slot along the path results in a 1 for that slot).

Another metric that takes into account horizontal fragmentation considers the alignment of available slots across link pairs, where each pair is composed of one link on the path, and one link that is adjacent to that link [12.88]. The metric captures how many optical channels of size S slots can be accommodated on each such link pair, for all possible S , and weights this count according to the expected number of optical channels of that size. Other fragmentation metrics, similar in spirit to those mentioned above, are proposed in [12.89, 90].

When alternative-path routing is used in a multistep RSA approach, the candidate path with, for example, the lowest path-based fragmentation metric could be

selected for a new optical connection. In an example of a one-step approach where routing and SA are handled together, the increase in the fragmentation metric is calculated for each feasible candidate path/SA combination. The path and assignment that results in the smallest increase in fragmentation is then chosen. While likely to produce improved performance as compared to the multistep approach, evaluating the various route/SA combinations could be time consuming.

12.12.2 Spectrum Assignment

In the multistep RSA approach, once a route has been selected for an optical channel, the next step is to determine where regeneration, if any, is required. Regeneration divides the path of the optical channel into a set of all-optical segments. Spectrum can be assigned to each segment independently; i.e., regeneration in EONs allows *spectrum conversion*, just as regeneration in wavelength-based networks allows wavelength conversion.

The WA problem was mapped to an instance of graph coloring, as is detailed in Sect. 12.8. This analogy can be extended to SA as well. As with WA, a conflict graph is constructed, where each vertex corresponds to an all-optical segment, and two vertices are connected by an edge if the corresponding segments have at least one network fiber in common. Additionally, for SA, each vertex is weighted by the number of slots that are required by the optical channel to which the segment belongs. The graph is colored using a *weighted* graph coloring algorithm. Any solution to the weighted graph coloring problem can be mapped to a solution for the corresponding SA problem, where the slots assigned to each vertex must be contiguous in order to enforce the spectral contiguity constraint [12.94].

Alternatively, SA can be mapped to the problem of scheduling tasks on a multiprocessor system [12.95]. Consider a set of tasks, each associated with a fixed set of processors and a required processing time. The task scheduling problem involves minimizing the *schedule length* (i.e., the time by which all tasks have been completed) subject to the following: a processor can work on at most one task at any given time; each task must be processed simultaneously by all processors in its associated set; and preemptions are not permitted [12.96–98]. By mapping each all-optical segment that needs to be assigned spectrum to a task and letting the set of links over which the segment is routed correspond to the fixed set of processors assigned to each task, the SA problem becomes an instance of task scheduling. The assigned start and end times of a task correspond to the assigned start and end slots of a segment. The objective of minimizing the schedule length corresponds to min-




Fig. 12.18 A new optical channel, requiring a total of four slots, is routed all-optically on three links. The *solid shaded slots* indicate spectrum that has previously been assigned. Of the three existing gaps, the best-fitting one is gap 1. The new optical channel, represented by the *hatched slots*, is shown assigned to this gap. This strands bandwidth on link 1 (slot 3) and on links 2 and 3 (slot 8) (© Monarch Network Architects LLC)

imizing the highest numbered slot needed to carry any of the optical channels.

Most WA heuristics can be extended for use in SA. For example, in the first-fit heuristic, the lowest-numbered spectral gap that is wide enough to accommodate the new optical channel (or all-optical segment) is selected [12.99, 100]. In the most-used heuristic, the spectral gap that maximizes the usage of the slots in that gap across the network is selected. Despite the additional information that is considered in most-used, first-fit yielded a slightly lower blocking probability in various studies [12.86, 101].

In one proposed variation, first-fit is combined with alternative-path routing, such that the candidate path is selected that has the lowest-numbered spectral gap of sufficient width along the whole path [12.102]. This is a one-step RSA algorithm; i.e., it selects the route and spectrum assignment together. This scheme is more appropriate for networks that do not require regeneration because it inherently assumes that the same spectrum assignment is utilized along the whole end-to-end path.

Variations of first-fit have also been proposed to better address fragmentation issues. For example, in first-last-fit [12.103], optical channels that require an odd number of slots are assigned spectrum using first-fit, whereas optical channels that require an even number of slots are assigned spectrum using last-fit (i.e., the highest-indexed spectral gap is selected). This *soft* partitioning of odd and even sized optical channels results in a small improvement in blocking probability [12.104].

Another SA heuristic that attempts to minimize fragmentation is best-fit. It assigns a new optical channel to the smallest spectral gap that is large enough to accommodate it. This heuristic is illustrated in Fig. 12.18, where the new optical channel is routed on links 1, 2, and 3, and requires four slots, including the guardband slot. The three existing gaps that span all three links are as shown. Gap 2 (two slots) is too small. Of the remaining two gaps, gap 1 (five slots) is selected over gap 3 (seven slots) because it is closer in size to the required four slots. The motivation behind best-fit is appealing, yet it has been shown to yield higher blocking as compared to first-fit [12.1]. The drawback of the scheme is that the residual unused portion of the selected gap (e.g., slot 8 on links 2 and 3) is typically very small, and effectively becomes stranded bandwidth. If gap 3 had been selected, starting at slot 18, a three-slot gap would remain on all three links. (Another potential advantage of first-fit is that it tends to group the utilized slots at one end of the spectrum, leaving larger available gaps at the other end of the spectrum. However, for a dynamic network, this quasi-orderly assignment pattern will likely dissipate over time.)

As compared to best-fit, a small performance improvement can be attained with exact-fit. This heuristic preferentially assigns an optical channel to a spectral gap that is of equal size. If such a gap does not exist, the scheme reverts to first-fit.

Another challenge that arises with SA is fairness with respect to blocking. As the spectrum becomes more fragmented, leaving only relatively small unoc-

cupied spectral gaps, finding a suitable spectrum assignment for high-bandwidth optical channels becomes increasingly more difficult. One can employ various partitioning schemes to address SA fairness [12.89]. For example, the spectrum can be divided into three regions, one each for low-rate, medium-rate, and high-rate connections. If the amount of traffic that falls within each one of these classes is not known, better performance may be obtained with a soft partitioning rather than a strict fixed partitioning.

Additional SA heuristic algorithms can be found in [12.84, 85]. Simulation results indicate only small differences in performance among the various heuristics considered (with the exception of the relatively poor-performing random assignment strategy, which assigns any spectral gap of sufficient size with equal probability) [12.85, 104].

12.12.3 ILP-Based RSA

A number of ILPs have been proposed for the RSA problem, where optimal (or near-optimal) routing and spectrum assignment is desired for a set of connection requests [12.100, 105]. As with RWA, such an approach is feasible only for small RSA instances. The ILP formulations need to capture the unique assignment of spectrum on a given fiber, the continuity constraint, and the contiguity constraint. Various objective functions can be considered, including: minimizing the maximum number of slots utilized on any link; minimizing the highest-numbered slot utilized on any link; and minimizing the total number of slots utilized across the network. If sufficient bandwidth is not available to carry all of the offered traffic, then the objective is typically to minimize the amount of requested bandwidth that is blocked.

Similar to ILPs used for RWA, by restricting the possible number of paths that can be utilized between a given source and destination (referred to as a *path-based* approach), the computation time can be markedly reduced with little loss of optimality [12.86, 106]. Furthermore, by inputting the set of possible slots that can be utilized by a particular connection (referred to as a *channel-based* approach), the number of required ILP constraints can be reduced significantly, thus speeding up the runtime (although it is still not fast enough to be able to handle large-sized RSA instances) [12.107]. Most of the proposed ILPs handle routing and spectrum assignment in one step. Alternatively, the problem can be decomposed into two ILPs (R + SA) to improve runtimes, though sacrificing some amount of optimality [12.86].

Alternatively, the RSA problem can be formulated as a *mixed integer program* (MIP). For example,

in [12.108], a MIP that utilizes both the path-based approach (to limit the number of paths to consider) and the channel-based approach (to limit the number of possible spectrum assignment choices) is combined with a number of advanced optimization techniques (e.g., branch-and-bound, column generation). It was shown that optimal results can be found for networks of moderate size.

12.12.4 Distance-Adaptive RSA

In order to enable higher bit rates, optical transmission modulation schemes have become more complex, thereby requiring the use of advanced digital signal processing (DSP) in the transmit/receive (TxRx) equipment. The presence of the DSP also affords the opportunity to provide greater TxRx flexibility, where transmission characteristics such as modulation format, error correction coding, and symbol rate can be modified remotely through software.

This flexibility enables important system design tradeoffs, including spectral width vs. optical reach [12.99, 102, 109, 110]. By utilizing a more spectrally efficient modulation format, the amount of spectrum allocated to a signal can be reduced, while the signal bit-rate is maintained. The disadvantage is that the optical reach of the signal is decreased. For example, a 100 Gb/s TxRx could be capable of either 3000 km optical reach with 75 GHz bandwidth, or 2000 km optical reach with 50 GHz bandwidth. For relatively short connections (or, more generally, connections that are afflicted by fewer impairments), the decrease in optical reach may not lead to any extra regeneration. Even for longer connections, the tradeoff (i.e., extra regeneration) may be worthwhile if capacity is at a premium. EONs, with their ability to dynamically adjust the spectral width assigned to an optical signal, are especially well suited to take advantage of this TxRx flexibility.

The selection of the modulation format can be considered in the network design process in order to better optimize the network with respect to factors such as cost or capacity utilization. This has given rise to distance-adaptive RSA (DA-RSA), also known as the routing, modulation level, and spectrum-assignment (RMLSA) problem, where the spectral width is tailored to the path over which an optical signal is routed. Note that the presence of a regeneration typically affords the opportunity to change the modulation scheme, just as it allows spectrum conversion [12.1, 111].

One potential drawback to a flexible modulation scheme is that there will likely be a mix of modulation formats co-propagating on one fiber. As discussed in Sect. 12.10.1, when certain combinations of modulation formats are used in adjacent regions of the spectrum,

the performance of the carried optical signals may degrade. This may reduce the optical reach or require that extra guardbands be utilized between optical channels, where the latter negates some of the spectral benefit of utilizing a more efficient modulation format. Mixed modulation formats were investigated in the context of the tradeoff between bit rate and optical reach, and a reduction in spectral benefits was noted [12.112].

Numerous algorithms that address the RMLSA problem in EONs have been proposed, e.g., [12.99, 102]. Many of these studies consider only true all-optical networks, where regeneration is not permitted. Thus, the approach often utilized is to generate a set of candidate paths and then associate each path with the most spectrally efficient modulation format that yields a sufficient optical reach. A slightly more restrictive version of this strategy considers only those candidate paths that can make use of the most spectrally efficient modulation format suitable for the shortest possible path [12.113]. This allows some amount of path diversity without sacrificing spectral utilization. An RMLSA algorithm that favors higher path OSNR, greater spectral efficiency, shorter path distance, and less slot utilization along the path (in order from the highest-priority criterion to the lowest one) is considered in [12.114]. ILPs that solve RMLSA in one step or two steps (RML + SA) have been proposed as well [12.102]. The MIP approach can also be extended to the RMLSA problem [12.108].

12.12.5 Routing, Spectrum and Core Assignment

As the capacity limits of *conventional* fiber are close to being reached [12.17, 115], several *space-division-multiplexed* (SDM) solutions have been proposed, where signals are carried on spatially separate channels. Each signal can still be a WDM signal, thus yielding a multiplicative effect in capacity. For example, in contrast to conventional single-core fiber, *multicore fiber* (MCF) supports multiple cores in a single fiber [12.116]. With C cores, it is expected that the capacity of a single fiber would increase by a factor of C .

Having multiple cores presents interesting tradeoffs with respect to wavelength and spectrum assignment. More attention has been given to spectrum assignment, as the presence of multiple cores may ease some of the challenges. The associated problem has been termed *routing, spectrum and core assignment* (RSCA) [12.117, 118]. Assuming that ROADMs allow for signals to be routed through a node on any of the cores, then the presence of multiple cores may improve the likelihood of finding contiguous available spec-

trum to support a new all-optical segment. Additionally, some of the cores could be restricted to supporting optical channels of a particular bandwidth. This would reduce the amount of stranded bandwidth and improve the alignment of available resources in adjacent links to provide more opportunity for optical bypass.

One drawback of the MCF approach, however, is that depending on how closely the cores are packed in the fiber, crosstalk may be present between signals in different cores [12.119]. If the crosstalk is severe enough, the wavelengths in different cores become coupled. For example, with C coupled cores, all $C\lambda$ s would need to be routed or added/dropped as a single block (essentially the C wavelengths form a *waveband*); the individual wavelengths cannot be routed independently. Thus, algorithms would be needed to optimize this coarser granularity. Note that coupling occurs (to an even greater degree) with *few-mode fiber* (FMF), which is an alternative SDM approach that parallels MCF (i.e., a small number of modes are supported on a fiber in contrast to the single mode that is present in conventional fiber in a long-haul or metro network) [12.116].

12.12.6 Multipath Routing

In the above discussion, it was assumed that a single optical channel is established to carry a connection. If sufficient contiguous spectral slots are not available to carry a new connection along a particular route, then either another route must be considered or the new connection request is blocked. EONs offer an alternative approach, where multiple smaller optical channels are utilized to carry a single connection [12.92]. For example, a new connection may require 45 GHz of spectrum, but no single path from the source to the destination exists with adequate available spectral gaps. Rather than blocking the connection request, it is possible to split it into two or more connections, each with spectral requirements of less than 45 GHz, in order to take advantage of whatever spectrum is available. Thus, the 45 GHz of spectrum required could be provided, for example, via three spectral regions of size 20, 15, and 10 GHz. The destination must be capable of reconstituting the original signal.

There are two strategies that can be utilized in this multiple-optical-channel approach. First, the optical channels can be limited to following the same path. With this restriction, the latency from source to destination will be approximately the same for all constituent channels (small variances arise due to slight differences in the speed of light in fiber as a function of the spectral frequency utilized), thereby avoiding the challenge of having to provide large buffers at the des-

mination in order to synchronize the received signals. This methodology is especially useful as the spectrum becomes more fragmented, where there exists free spectrum on a fiber, but it is divided into relatively small spectral gaps. It also can be advantageous for routing high-bandwidth connections, where it is more challenging to find a spectral gap of appropriate size.

A more aggressive approach allows splitting a connection across optical channels that are routed on different paths. This clearly provides greater flexibility. The drawback is the *delay differential* among the selected paths, chiefly arising from the differences in the end-to-end distances of the paths. The greater the differential, the larger the required destination buffers that are used to reconstitute the original signal. Various algorithms exist to find a set of paths where the delay differential between the shortest and longest paths is below an acceptable threshold [12.120–124]. In one approach [12.123], a graph is created, where the nodes represent possible paths between the source and destination. Links are added between a pair of nodes if the differential delay between the corresponding two paths is below the acceptable threshold. A search is performed to find the minimum clique on this graph such that the total bandwidth that can be carried by the set of paths represented by the clique nodes is sufficient to carry the connection.

Regardless of whether the constituent optical channels are routed along the same path or different paths, one important drawback stems from the requirement of needing guardbands between optical channels. Thus, partitioning a connection across multiple optical channels results in more guardbands being required. This effectively limits the usage of this tactic; otherwise the benefits gained by the greater spectrum-assignment flexibility will be mitigated by the additional wasted bandwidth assigned to guardbands. To address this, one can impose a minimum bandwidth size for each of the constituent optical channels [12.123]. Another strategy is to limit multipath routing to those connections requiring a bandwidth greater than some threshold. This would limit the number of excess guardbands and alleviate some of the fairness concerns regarding high-rate connections.

Another potential drawback is the number of transponders required to support the constituent optical channels. For example, a connection that is split among three optical channels would require at least six transponders (one transponder at either end of an optical channel), and possibly more for regeneration. As noted in Sect. 12.11, the challenge of requiring numerous transponders may be addressed by BVTs that can be sliced into several *virtual transponders*, each of which serves one optical channel.

Multipath routing provides benefits beyond just capacity management. It can also be incorporated as part of a protection strategy [12.125, 126]. For example, a connection may ideally require a rate of R but under failure conditions is satisfied with R' , for some $R' < R$. The connection can be split across N diverse paths, where the total bandwidth of the paths is R , and the total bandwidth of any $(N - 1)$ of the paths is R' . If a failure occurs such that one of the paths is brought down, the remaining paths provide adequate bandwidth. This is an example of *bandwidth squeezing restoration*, where the bandwidth may be reduced under failure conditions [12.127]. In another scheme, which utilizes shared restoration, one optical channel is used for the working path, but multipath routing is permitted for the backup resources. The added flexibility in routing the backup traffic typically results in a greater amount of sharing of the restoration resources, thus reducing the amount of required spare capacity in the network. Note that most multipath-based protection/restoration schemes rely on algorithms that specifically find *diverse* paths that satisfy the maximum acceptable differential delay [12.128].

Multipath routing is supported in many standards, e.g., *virtual concatenation* (VCAT) in SONET/SDH and OTN [12.129, 130].

12.12.7 Defragmentation

The quasi-gridless nature of EONs, combined with increasing network dynamism where connections are established and torn down on shorter time scales, produces significant spectral fragmentation. As indicated in Sect. 12.12.1, fragmentation can be present on a single link, where the available spectrum consists of narrow, noncontiguous spectral gaps. Larger gaps are preferable, as they offer the opportunity to carry a greater range of new traffic (it is assumed here that single-path routing is employed). Fragmentation can also be problematic across links, where two consecutive links in a path may have spectral gaps of sufficient size to carry new traffic, but the gaps do not align. This necessitates regeneration of the optical connection for purposes of spectrum conversion. This is illustrated in Fig. 12.19a, where it is desired to route a new 25 GHz optical channel over links 1 and 2 (explicit guardbands are not shown). While both links have an available 25 GHz gap, the gaps are not spectrally aligned. Thus, regeneration would be required at the ROADM. In a network where regeneration is not supported, the new connection would be blocked (unless another route is utilized).

Fragmentation is likely to occur in EONs even if fragmentation-aware RSA policies are utilized. This

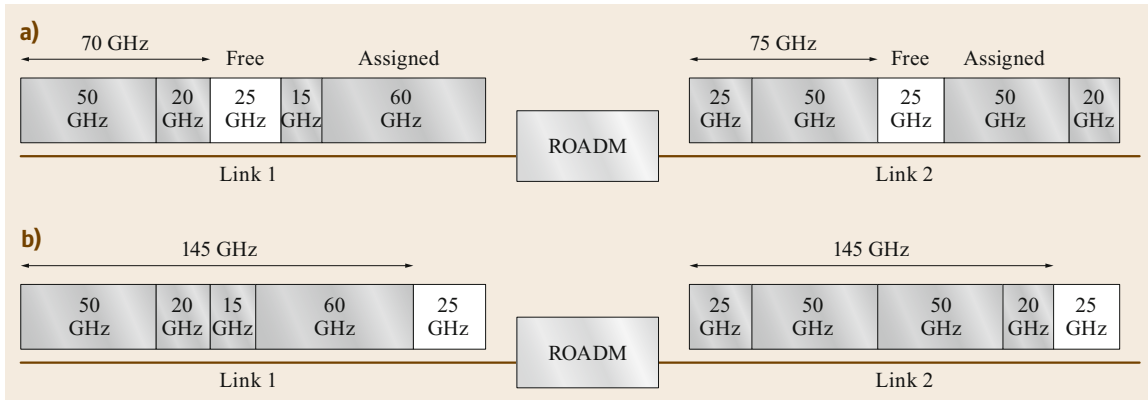


Fig. 12.19a,b Assume that a new optical channel requiring 25 GHz of spectrum needs to be routed across links 1 and 2. **(a)** Both links have an available 25 GHz gap; however, the gaps are not aligned. The optical channel cannot be routed all-optically through the ROADM, but instead would be regenerated to accomplish spectrum conversion. **(b)** After defragmentation, the assigned spectrum forms a contiguous block, leaving the available spectrum aligned (© Monarch Network Architects LLC)

mandates that defragmentation be implemented in such networks for purposes of efficient capacity management and reduced cost (less regeneration). Defragmentation involves the shifting or moving of spectrum that has already been assigned in order to aggregate the free spectrum into larger contiguous blocks on a link and align the free spectrum on adjacent links. This is shown in Fig. 12.19b, where the spectrum utilized on both links 1 and 2 is pushed into one contiguous region, leaving the available spectrum aligned at the high end of the spectrum. A new 25 GHz-channel can now be routed all-optically through the ROADM.

It should be noted that most defragmentation techniques have a tendency to realign the spectrum assignment such that the spectral slots utilized on a fiber are contiguous. However, this impedes the ability of an optical channel to increase its assigned bandwidth. Thus, defragmentation that leaves small spectral gaps between optical channels that exhibit elastic behavior may be desirable [12.131]. These small gaps could potentially be used in a time-based multiplexing scheme. For example, an optical channel on one side of the gap may expand its spectral allotment during daytime hours, whereas the optical channel on the other side of the gap expands at night. If there are times when neither optical channel is expanded, that spectrum can be assigned to low-priority traffic that may eventually be bumped.

Defragmentation is a challenging operation, as it requires the adjustment of numerous network elements; furthermore, it is operating on live traffic. Careful timing is needed to effect hitless (or close to hitless) defragmentation. The TxRx equipment at either end of the optical channel being shifted needs to be si-

multaneously retuned to the new spectral region. At the same time, the filters of any ROADMs along the path of the optical channel must be reconfigured to accommodate the newly-assigned spectral region, without disrupting any other existing traffic that passes through the ROADM. To ease the timing requirements, many schemes advocate the use of a *make-before-break* mechanism, where the new optical channel is established prior to the original one being removed. This depends on there being extra TxRx equipment to temporarily support an extra optical channel (or virtual or multiflow TxRx's can be used that support multiple optical channels [12.76, 132, 133]).

There are several approaches that can be taken with respect to defragmentation. In a reactive approach, defragmentation is performed only when new connection requests are blocked due to spectral fragmentation [12.134] (or when too much excess regeneration is required to accommodate a new path). The selection of the route for the new channel can, at least in part, be based on minimizing the number of conflicts with existing channels [12.135, 136].

In a proactive approach, defragmentation is performed periodically, e.g., as part of routine network maintenance, or when a network-wide fragmentation metric rises above a set threshold. One implicit periodic approach to defragmentation takes advantage of the growth/contraction cycle of the optical channels that exhibit elastic behavior. Growth is accomplished by preferentially adding slots at the lower end of the channel, whereas contraction is performed from the upper end of the channel [12.1]. This has a tendency to shift the channels to the lower end of the spectrum, resulting in more closely packed channels.




Fig. 12.20a–c Push–pull defragmentation. (a) The original spectral partitioning. Assume that it is desired to shift the optical channel on slots 9 and 10 to slots 4 and 5 on that same link. (b) In the first phase, the channel is extended to encompass the original slots, the new slots, and any slots in between. (c) In the second phase, the channel is contracted to encompass just the new slots, 4 and 5

Another bifurcation relates to the aggressiveness of the approach. In one class of strategies, rerouting is permitted, such that both the path and the spectrum assignment may be modified for an existing optical connection. In a less aggressive approach, rerouting of existing connections is not permitted; defragmentation is limited to shifting spectrum on a link. An example of this is a defragmentation mechanism that takes advantage of channel elasticity [12.76, 137]. This push–pull mechanism is illustrated in Fig. 12.20. Assume that the channel to be moved is initially assigned to slots 9 and 10, as shown in Fig. 12.20a. The objective of the defragmentation is to move the channel to slots 4 and 5. To accomplish this, the channel is first expanded to encompass both the original slots and the new slots, and all slots in between, as shown in Fig. 12.20b. The channel is then contracted to occupy only the desired new slots, as is shown in Fig. 12.20c. These expansion and contraction operations must be performed concurrently on each ROADMs along the path of the optical channel, as well as at the TxRxs at the endpoints, so that end-to-end connectivity is never lost. Hitless operation of push–pull has been demonstrated [12.137].

While push–pull supports only a limited amount of defragmentation; i.e., spectrum can only be shifted

along the same fiber, and all slots between the old and new spectrum assignments must be unassigned, it can provide a notable improvement in performance. For example, simulations have shown that it can potentially reduce the level of blocking by one to two orders of magnitude when used reactively to avoid blocking [12.138].

A third design decision is whether the defragmentation process is limited to sequential moves, where no more than one optical channel is adjusted at a given time. Or, defragmentation can involve parallel operations, where multiple optical channels are spectrally shifted and/or rerouted at one time [12.139]. Parallel operation provides more flexibility and faster operation time, although it requires accurate time synchronization across the network.

While defragmentation of an EON has been the focus here, fragmentation may also arise in the wavelength-based flexible-grid architecture supported by the new standard [12.67]. However, due to the relatively small number of bandwidths that are likely to be utilized on one fiber (e.g., perhaps 50, 62.5, and 75 GHz), the level of fragmentation should not be as severe as in an EON. Nevertheless, it is possible that defragmentation could be warranted [12.140].

12.13 Conclusion

As illustrated in this chapter, algorithms are a fundamental component of network design. Efficient utilization of network resources, which is tightly coupled to

the economics of the network, relies on efficient algorithms. Furthermore, as new technologies are developed to improve network performance, new algorithms must

be developed to take full advantage of these advances. By extending previously developed design algorithms, relying to some extent on graph theory, and utilizing

common sense in design decisions, the growing complexity of managing an optical network has thus far been met.

References

- 12.1 J.M. Simmons: *Optical Network Design and Planning*, 2nd edn. (Springer, New York 2014)
- 12.2 T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein: *Introduction to Algorithms*, 3rd edn. (MIT Press, Cambridge 2009)
- 12.3 R. Bhandari: *Survivable Networks: Algorithms for Diverse Routing* (Kluwer, Boston 1999)
- 12.4 R. Guerin, A. Orda: Computing shortest paths for any number of hops, *IEEE/ACM Trans. Netw.* **10**(5), 613–620 (2002)
- 12.5 T. Korkmaz, M. Krunz: Multi-constrained optimal path selection. In: *Proc. IEEE INFOCOM, Anchorage, Apr. 22–26, Vol. 2* (2001) pp. 834–843
- 12.6 F. Kuipers, T. Korkmaz, M. Krunz, P. Van Mieghem: Performance evaluation of constraint-based path selection algorithms, *IEEE Network* **18**(5), 16–23 (2004)
- 12.7 J.Y. Yen: Finding the K shortest loopless paths in a network, *Manag. Sci.* **17**(11), 712–716 (1971)
- 12.8 J.W. Suurballe, R.E. Tarjan: A quick method for finding shortest pairs of disjoint paths, *Networks* **14**, 325–336 (1984)
- 12.9 J.W. Suurballe: Disjoint paths in a network, *Networks* **4**, 125–145 (1974)
- 12.10 D. Xu, Y. Xiong, C. Qiao, G. Li: Trap avoidance and protection schemes in networks with shared risk link groups, *J. Lightwave Technol.* **21**(11), 2683–2693 (2003)
- 12.11 A. Bach: The financial industry's race to zero latency and terabit networking. In: *Proc. Opt. Fiber Commun./Natl. Fiber Opt. Eng. Conf. (OFC/NFOEC), Los Angeles, Mar. 6–10* (2011), Service Provider Summit Keynote Address
- 12.12 E. Karasan, E. Ayanoglu: Effects of wavelength routing and selection algorithms on wavelength conversion gain in WDM optical networks, *IEEE/ACM Trans. Netw.* **6**(2), 186–196 (1998)
- 12.13 X. Chu, B. Li, Z. Zhang: A dynamic RWA algorithm in a wavelength-routed all-optical network with wavelength converters. In: *Proc. IEEE INFOCOM, San Francisco, Mar. 30–Apr. 3, Vol. 3* (2003) pp. 1795–1804
- 12.14 J.L. Strand: Integrated route selection, transponder placement, wavelength assignment, and restoration in an advanced ROADM architecture, *J. Opt. Commun. Netw.* **4**(3), 282–288 (2012)
- 12.15 J.M. Simmons: Diversity requirements for selecting candidate paths for alternative-path routing. In: *Proc. Opt. Fiber Commun./Natl. Fiber Opt. Eng. Conf., San Diego, Mar. 21–25* (2010), Paper NThA4
- 12.16 R.A. Barry, P.A. Humblet: Models of blocking probability in all-optical networks with and without wavelength changers, *IEEE J. Sel. Areas Commun.* **14**(5), 858–867 (1996)
- 12.17 J.M. Simmons: Technology and architectural approaches to address continued explosive growth in network traffic. In: *Proc. Int. Conf. Comput. Netw. Commun., Santa Clara, Jan. 28* (2017)
- 12.18 H. Hasegawa, S. Subramaniam, K. Sato: Node architecture and design of flexible waveband routing optical networks, *J. Opt. Commun. Netw.* **8**(10), 734–744 (2016)
- 12.19 A. Farrel, J.-P. Vasseur, J. Ash: A path computation element (PCE)-based architecture, Internet Engineering Task Force, Request for Comments 4655 (2006)
- 12.20 F. Paolucci, F. Cugini, A. Giorgetti, N. Sambo, P. Castoldi: A survey on the path computation element (PCE) architecture, *IEEE Commun. Surv. Tutor.* **15**(4), 1819–1841 (2013)
- 12.21 P.J.M. van Laarhoven, E.H.L. Aarts: *Simulated Annealing: Theory and Applications* (D. Reidel, Boston 1987)
- 12.22 L.H. Sahasrabudhe, B. Mukherjee: Multicast routing algorithms and protocols: a tutorial, *IEEE Network* **14**(1), 90–102 (2000)
- 12.23 S. Voss: Steiner's problem in graphs: heuristic methods, *Discrete Appl. Math.* **40**(1), 45–72 (1992)
- 12.24 L. Kou, G. Markowsky, L. Berman: A fast algorithm for Steiner trees, *Acta Inform.* **15**(2), 141–145 (1981)
- 12.25 B.M. Waxman: Routing of multipoint connections, *IEEE J. Sel. Areas Commun.* **6**(9), 1617–1622 (1988)
- 12.26 H. Takahashi, A. Matsuyama: An approximate solution for the Steiner problem in graphs, *Math. Jpn.* **24**(6), 573–577 (1980)
- 12.27 N.K. Singhal, L.H. Sahasrabudhe, B. Mukherjee: Provisioning of survivable multicast sessions against single link failures in optical WDM mesh networks, *IEEE/OSA J. Lightwave Technol.* **21**(11), 2587–2594 (2003)
- 12.28 X. Wang, L. Guo, L. Pang, J. Du, F. Jin: Segment protection algorithm with load balancing for multicasting WDM mesh networks. In: *Proc. 10th Int. Conf. Adv. Commun. Technol., Pyeongchang, Feb. 17–20, Vol. 3* (2008) pp. 2013–2016
- 12.29 T. Panayiotou, G. Ellinas, N. Antoniadis: Segment-based protection of multicast connections in metropolitan area optical networks with quality-of-transmission considerations, *J. Opt. Commun. Netw.* **4**(9), 692–702 (2012)
- 12.30 R. Koetter, M. Médard: An algebraic approach to network coding, *IEEE/ACM Trans. Netw.* **11**(5), 782–795 (2003)
- 12.31 R.C. Menendez, J.W. Gannett: Efficient, fault-tolerant all-optical multicast networks via network coding. In: *Proc. Opt. Fiber Commun./Natl. Fiber*

- Opt. Eng. Conf., San Diego, Feb. 24–28 (2008)*, Paper JThA82
- 12.32 E.D. Manley, J.S. Deogun, L. Xu, D.R. Alexander: All-optical network coding, *J. Opt. Commun. Netw.* **2**(4), 175–191 (2010)
- 12.33 N. Charbonneau, V.M. Vokkarane: Tabu search meta-heuristic for static manycast routing and wavelength assignment over wavelength-routed optical WDM networks. In: *Proc. IEEE Int. Conf. Commun., Cape Town, May 23–27 (2010)*, <https://doi.org/10.1109/ICC.2010.5502241>
- 12.34 S. Subramaniam, M. Azizoglu, A.K. Somani: All-optical networks with sparse wavelength conversion, *IEEE/ACM Trans. Netw.* **4**(4), 544–557 (1996)
- 12.35 J.M. Simmons: Analysis of wavelength conversion in all-optical express backbone networks. In: *Proc. Opt. Fiber Commun. (OFC), Anaheim, Mar. 17–22 (2002)*, Paper TuG2
- 12.36 D. Breaz: New methods to color the vertices of a graph, *Communications ACM* **22**(4), 251–256 (1979)
- 12.37 H. Zang, J.P. Jue, B. Mukherjee: A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks, *Opt. Netw. Mag.* **1**(1), 47–60 (2000)
- 12.38 I. Chlamtac, A. Ganz, G. Karmi: Purely optical networks for terabit communication. In: *Proc. IEEE INFOCOM, Ottawa, Apr. 23–27, Vol. 3 (1989)* pp. 887–896
- 12.39 O. Gerstel, H. Raza: Predeployment of resources in agile photonic networks, *J. Lightwave Technol.* **22**(10), 2236–2244 (2004)
- 12.40 D. Banerjee, B. Mukherjee: A practical approach for routing and wavelength assignment in large wavelength-routed optical networks, *IEEE J. Sel. Areas Commun.* **14**(5), 903–908 (1996)
- 12.41 A.E. Ozdaglar, D.P. Bertsekas: Routing and wavelength assignment in optical networks, *IEEE/ACM Trans. Netw.* **11**(2), 259–272 (2003)
- 12.42 K. Christodoulopoulos, K. Manousakis, E. Varvarigos: Comparison of routing and wavelength assignment algorithms in WDM networks. In: *Proc. IEEE Global Commun. Conf. (GLOBECOM), New Orleans, Nov. 30–Dec. 4 (2008)*, <https://doi.org/10.1109/GLOCOM.2008.ECP.510>
- 12.43 E. Yetginer, Z. Liu, G.N. Rouskas: Fast exact ILP decompositions for ring RWA, *J. Opt. Commun. Netw.* **3**(7), 577–586 (2011)
- 12.44 S. Azodolmolky, M. Klinkowski, E. Marin, D. Careglio, J. Pareta, I. Tomkos: A survey on physical layer impairments aware routing and wavelength assignment algorithms in optical networks, *Comput. Netw.* **53**(7), 926–944 (2009)
- 12.45 A.G. Rahbar: Review of dynamic impairment-aware routing and wavelength assignment techniques in all-optical wavelength-routed networks, *IEEE Commun. Surv. Tutor.* **14**(4), 1065–1089 (2012)
- 12.46 C. Kurtzke: Suppression of fiber nonlinearities by appropriate dispersion management, *IEEE Photonics Technol. Lett.* **5**(10), 1250–1253 (1993)
- 12.47 R.W. Tkach, A.R. Chraplyvy: Dispersion and nonlinear effects in lightwave systems. In: *Proc. 7th Annu. Meeting IEEE LEOS, Boston, Oct. 31–Nov. 3, Vol. 1 (1994)* pp. 192–193
- 12.48 R.W. Tkach, A.R. Chraplyvy, F. Forghieri, A.H. Gnauck, R.M. Derosier: Four-photon mixing and high-speed WDM systems, *J. Lightwave Technol.* **13**(5), 841–849 (1995)
- 12.49 F. Forghieri, R.W. Tkach, A.R. Chraplyvy: Fiber nonlinearities and their impact on transmission systems. In: *Optical Fiber Telecommunications, Vol. III A*, ed. by I. Kaminow, T. Koch (Academic Press, San Diego 1997)
- 12.50 A.H. Gnauck, R.M. Jopson: Dispersion compensation for optical fiber systems. In: *Optical Fiber Telecommunications, Vol. III A*, ed. by I. Kaminow, T. Koch (Academic Press, San Diego 1997)
- 12.51 C.D. Poole, J. Nagel: Polarization effects in lightwave systems. In: *Optical Fiber Telecommunications, Vol. III A*, ed. by I. Kaminow, T. Koch (Academic Press, San Diego 1997)
- 12.52 P. Bayvel, R. Killley: Nonlinear optical effects in WDM transmission. In: *Optical Fiber Telecommunications, Vol. IV B*, ed. by I. Kaminow, T. Li (Academic Press, San Diego 2002)
- 12.53 K. Christodoulopoulos, P. Kokkinos, K. Manousakis, E.A. Varvarigos: Cross layer RWA in WDM networks: Is the added complexity useful or a burden? In: *Proc. Int. Conf. Transpar. Opt. Netw., Ponta Delgada, Jun. 28–Jul. 2 (2009)*, Paper Tu.A3.3
- 12.54 K. Manousakis, P. Kokkinos, K. Christodoulopoulos, E. Varvarigos: Joint online routing, wavelength assignment and regenerator allocation in translucent optical networks, *J. Lightwave Technol.* **28**(8), 1152–1163 (2010)
- 12.55 T. Jimenez, J.C. Aguado, I. de Miguel, R.J. Duran, M. Angelou, N. Merayo, P. Fernandez, R.M. Lorenzo, I. Tomkos, E.J. Abril: A cognitive quality of transmission estimator for core optical networks, *J. Lightwave Technol.* **31**(6), 942–951 (2013)
- 12.56 L. Barletta, A. Giusti, C. Rottondi, M. Tornatore: QoT estimation for unestablished lighpaths using machine learning. In: *Proc. Opt. Fiber Commun. Conf. (OFC), Los Angeles, Mar. 19–23 (2017)*, Paper Th1J.1
- 12.57 S. Oda, M. Miyabe, S. Yoshida, T. Katagiri, Y. Aoki, T. Hoshida, J.C. Rasmussen, M. Birk, K. Tse: A learning living network with open ROADMs, *J. Lightwave Technol.* **35**(8), 1350–1356 (2017)
- 12.58 A. Bononi, M. Bertolini, P. Serena, G. Bellotti: Cross-phase modulation induced by OOK channels on higher-rate DQPSK and coherent QPSK channels, *J. Lightwave Technol.* **27**(18), 3974–3983 (2009)
- 12.59 O. Bertran-Pardo, J. Renaudier, G. Charlet, H. Mardoyan, P. Tran, M. Salsi, S. Bigo: Overlaying 10 Gb/s legacy optical networks with 40 and 100 Gb/s coherent terminals, *J. Lightwave Technol.* **30**(14), 2367–2375 (2012)
- 12.60 R. Dutta, G.N. Rouskas: Traffic grooming in WDM networks: past and future, *IEEE Network* **16**(6), 46–56 (2002)

- 12.61 R. Dutta, A.E. Kamal, G.N. Rouskas (Eds.): *Traffic Grooming for Optical Networks: Foundations, Techniques and Frontiers* (Springer, New York 2008)
- 12.62 D.J. Blumenthal: Optical packet switching. In: *Proc. 17th Annu. Meeting IEEE LEOS, Puerto Rico, Nov. 7–11* (2004), Paper ThU1
- 12.63 C. Qiao, M. Yoo: Optical burst switching (OBS)—a new paradigm for an optical internet, *J. High Speed Netw.* **8**(1), 69–84 (1999)
- 12.64 M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, T. Yoshimatsu, T. Kobayashi, Y. Miyamoto, K. Yonenaga, A. Takada, O. Ishida, S. Matsuoka: Demonstration of novel spectrum-efficient elastic optical path network with per-channel variable capacity of 40 Gb/s to over 400 Gb/s. In: *Proc. Eur. Conf. Opt. Commun. (ECOC), Brussels, Sep. 21–25* (2008), Paper Th.3.F.6
- 12.65 M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, S. Matsuoka: Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies, *IEEE Commun. Mag.* **47**(11), 66–73 (2009)
- 12.66 O. Gerstel, M. Jinno, A. Lord, S.J.B. Yoo: Elastic optical networking: a new dawn for the optical layer?, *IEEE Commun. Mag.* **50**(2), S12–S20 (2012)
- 12.67 International Telecommunication Union: *Spectral Grids for WDM Applications: DWDM Frequency Grid*, ITU-T Rec. G.694.1, 2.0 edn. 2012)
- 12.68 International Telecommunication Union: *Link Capacity Adjustment Scheme (LCAS) for Virtual Concatenated Signals*, ITU-T Rec. G.7042/Y.1305, 2006)
- 12.69 I.B. Djordjevic, B. Vasic: Orthogonal frequency division multiplexing for high-speed optical transmission, *Opt. Express* **14**(9), 3767–3775 (2006)
- 12.70 W. Shieh, H. Bao, Y. Tang: Coherent optical OFDM: theory and design, *Opt. Express* **16**(2), 841–859 (2008)
- 12.71 W. Shieh: OFDM for flexible high-speed optical networks, *J. Lightwave Technol.* **29**(10), 1560–1577 (2011)
- 12.72 G. Shen, M. Zukerman: Spectrum-efficient and agile CO-OFDM optical transport networks: architecture, design, and operation, *IEEE Commun. Mag.* **50**(5), 82–89 (2012)
- 12.73 G. Zhang, M. De Leenheer, A. Morea, B. Mukherjee: A survey on OFDM-based elastic core optical networking, *IEEE Commun. Surv. Tutor.* **15**(1), 65–87 (2013)
- 12.74 G. Bosco, A. Carena, V. Curri, P. Poggiolini, F. Forghieri: Performance limits of Nyquist-WDM and CO-OFDM in high-speed PM-QPSK systems, *IEEE Photonics Technol. Lett.* **22**(1), 1129–1131 (2010)
- 12.75 G. Gavioli, E. Torrenco, G. Bosco, A. Carena, V. Curri, V. Miot, P. Poggiolini, M. Belmonte, F. Forghieri, C. Muzio, S. Piciaccia, A. Brinciotti, A. La Porta, C. Lezzi, S. Savory, S. Abrate: Investigation of the impact of ultra-narrow carrier spacing on the transmission of a 10-carrier 1Tb/s super-channel. In: *Proc. Opt. Fiber Commun./Natl. Fiber Optic Eng. Conf. (OFC/NFOEC), San Diego, Mar. 21–25* (2010), Paper OTuD3
- 12.76 O. Gerstel: Flexible use of spectrum and photonic grooming. In: *Proc. Int. Conf. Photonics Switch., Monterey, Jul. 25–28* (2010), Paper PMD3
- 12.77 R. Ryf, Y. Su, L. Moller, S. Chandrasekhar, X. Liu, D.T. Nelson, C.R. Giles: Wavelength blocking filter with flexible data rates and channel spacing, *J. Lightwave Technol.* **23**(1), 54–61 (2005)
- 12.78 P. Colbourne, B. Collings: ROADM switching technologies. In: *Proc. Opt. Fiber Commun./Natl. Fiber Opt. Eng. Conf. (OFC/NFOEC), Los Angeles, Mar. 6–10* (2011), Paper OTuD1
- 12.79 D.M. Marom, D. Sinefeld: Beyond wavelength-selective channel switches: trends in support of flexible/elastic optical networks. In: *Proc. Int. Conf. Transpar. Opt. Netw. (ICTON), United Kingdom, Jul. 2–5* (2012), Paper Mo.B1.4
- 12.80 G. Baxter, S. Frisken, D. Abakoumov, H. Zhou, I. Clarke, A. Bartos, S. Poole: Highly programmable wavelength selective switch based on liquid crystal on silicon switching elements. In: *Proc. Opt. Fiber Commun./Natl. Fiber Opt. Eng. Conf. (OFC/NFOEC), Anaheim, Mar. 5–10* (2006), Paper OTuF2
- 12.81 S. Frisken: Advances in liquid crystal on silicon wavelength selective switching. In: *Proc. Opt. Fiber Commun./Natl. Fiber Opt. Eng. Conf. (OFC/NFOEC), Anaheim, Mar. 25–29* (2007), Paper OWV4
- 12.82 S. Frisken, G. Baxter, D. Abakoumov, H. Zhou, I. Clarke, S. Poole: Flexible and grid-less wavelength selective switch using LCOS technology. In: *Proc. Opt. Fiber Commun./Natl. Fiber Opt. Eng. Conf. (OFC/NFOEC), Los Angeles, Mar. 6–10* (2011), Paper OTuM3
- 12.83 G. Shen, Q. Yang: From coarse grid to mini-grid to gridless: How much can gridless help contention-less? In: *Proc. Opt. Fiber Commun./Natl. Fiber Opt. Eng. Conf. (OFC/NFOEC), Los Angeles, Mar. 6–10* (2011), Paper OTuI3
- 12.84 S. Talebi, F. Alam, I. Katib, M. Khamis, R. Salama, G.N. Rouskas: Spectrum management techniques for elastic optical networks: a survey, *Opt. Switch. Netw.* **13**, 34–48 (2014)
- 12.85 B.C. Chatterjee, N. Sarma, E. Oki: Routing and spectrum allocation in elastic optical networks: a tutorial, *IEEE Commun. Surv. Tutor.* **17**(3), 1776–1800 (2015)
- 12.86 K. Christodoulouopoulos, I. Tomkos, E.A. Varvarigos: Routing and spectrum allocation in OFDM-based optical networks with elastic bandwidth allocation. In: *Proc. IEEE Global Telecommun. Conf. (GLOBECOM), Miami, Dec. 6–10* (2010) pp. 1–6
- 12.87 A.N. Patel, P.N. Ji, J.P. Jue, T. Wang: A naturally-inspired algorithm for routing, wavelength assignment, and spectrum allocation in flexible grid WDM networks. In: *Proc. IEEE Global Commun. Conf. (GLOBECOM), Anaheim, Dec. 3–7* (2012) pp. 340–345
- 12.88 Y. Sone, A. Hirano, A. Kadohata, M. Jinno, O. Ishida: Routing and spectrum assignment algorithm maximizes spectrum utilization in optical networks. In: *Proc. Eur. Conf. Opt. Commun. (ECOC), Geneva, Sep. 18–22* (2011), Paper Mo.1.K.3

- 12.89 R. Wang, B. Mukherjee: Spectrum management in heterogeneous bandwidth optical networks, *Opt. Switch. Netw.* **11**(A), 83–91 (2014)
- 12.90 M. Zhang, W. Lu, Z. Zhu, Y. Yin, S.J.B. Yoo: Planning and provisioning of elastic O-OFDM networks with fragmentation-aware routing and spectrum assignment (RSA) algorithms. In: *Proc. Asia Commun. Photonics Conf., Guanzhou, Nov. 7–10* (2012), Paper Ath2D.3
- 12.91 X. Wang, Q. Zhang, I. Kim, P. Palacharla, M. Sekiya: Utilization entropy for assessing resource fragmentation in optical networks. In: *Proc. Opt. Fiber Commun. Conf./Natl. Fiber Opt. Eng. Conf. (OFC/NFOEC), Los Angeles, Mar. 4–8* (2012), Paper OTh1A.2
- 12.92 Z. Zhu, W. Lu, L. Zhang, N. Ansari: Dynamic service provisioning in elastic optical networks with hybrid single-/multi-path routing, *J. Lightwave Technol.* **31**(1), 15–22 (2013)
- 12.93 P. Wright, M.C. Parker, A. Lord: Minimum- and maximum-entropy routing and spectrum assignment for flexgrid elastic optical networking, *J. Opt. Commun. Netw.* **7**(1), A66–A72 (2015)
- 12.94 I. Popescu, I. Cerutti, N. Sambo, P. Castoldi: On the optimal design of a spectrum-switched optical network with multiple modulation formats and rates, *J. Opt. Commun. Netw.* **5**(11), 1275–1284 (2013)
- 12.95 S. Talebi, E. Bampis, I. Katib, G.N. Rouskas: The spectrum assignment (SA) problem in optical networks: a multiprocessor scheduling perspective, *J. Opt. Commun. Netw.* **6**(8), 754–763 (2014)
- 12.96 E. Bampis, M. Caramia, J. Fiala, A. Fishkin, A. Iovannella: Scheduling of independent dedicated multiprocessor tasks. In: *Proc. 13th Int. Symp. Algorithms Comput., Vancouver, Nov. 21–23* (2002) pp. 391–402
- 12.97 J.A. Hoogeveen, S.L. Van de Velde, B. Veltman: Complexity of scheduling multiprocessor tasks with prespecified processor allocations, *Discrete Appl. Math.* **55**, 259–272 (1994)
- 12.98 E. Bampis, A. Kononov: On the approximability of scheduling multiprocessor tasks with time dependent processing and processor requirements. In: *Proc. 15th Int. Parallel Distrib. Process. Symp., San Francisco, Apr. 23–27* (2001) pp. 2144–2151
- 12.99 M. Jinno, B. Kozicki, H. Takara, A. Watanabe, Y. Sone, T. Tanaka, A. Hirano: Distance-adaptive spectrum resource allocation in spectrum-sliced elastic optical path network, *IEEE Commun. Mag.* **48**(8), 138–145 (2010)
- 12.100 Y. Wang, X. Cao, Y. Pan: A study of the routing and spectrum allocation in spectrum-sliced elastic optical path networks. In: *Proc. IEEE INFOCOM, Shanghai, Apr. 10–15* (2011) pp. 1503–1511
- 12.101 R.J. Durán, I. Rodríguez, N. Fernández, I. de Miguel, N. Merayo, P. Fernández, J.C. Aguado, T. Jiménez, R.M. Lorenzo, E.J. Abril: Performance comparison of methods to solve the routing and spectrum allocation problem. In: *Proc. Int. Conf. Transpar. Opt. Netw. (ICTON), United Kingdom, Jul. 2–5* (2012), Paper Mo.C2.4
- 12.102 K. Christodoulopoulos, I. Tomkos, E.A. Varvarigos: Elastic bandwidth allocation in flexible OFDM-based optical networks, *J. Lightwave Technol.* **29**(9), 1354–1366 (2011)
- 12.103 W. Fadini, E. Oki: A subcarrier-slot partition scheme for wavelength assignment in elastic optical networks. In: *Proc. Int. Conf. High Perform. Switch. Routing (HPSR), Vancouver, Jul. 1–4* (2014) pp. 7–12
- 12.104 B.C. Chatterjee, E. Oki: Performance evaluation of spectrum allocation policies for elastic optical networks. In: *Proc. Int. Conf. Transpar. Opt. Netw. (ICTON), Budapest, Jul. 5–9* (2015), Paper Tu.D3.5
- 12.105 Y. Wang, X. Cao, Q. Hu, Y. Pan: Towards elastic and fine-granular bandwidth allocation in spectrum-sliced optical networks, *J. Opt. Commun. Netw.* **4**(11), 906–917 (2012)
- 12.106 M. Klinkowski, K. Walkowiak: Routing and spectrum assignment in spectrum sliced elastic optical path network, *IEEE Commun. Lett.* **15**(8), 884–886 (2011)
- 12.107 L. Velasco, M. Klinkowski, M. Ruiz, J. Comellas: Modeling the routing and spectrum allocation problem for flexgrid optical networks, *Photonic Netw. Commun.* **24**, 177–186 (2012)
- 12.108 M. Klinkowski, M. Zoltkiewicz, K. Walkowiak, M. Pioro, M. Ruiz, L. Velasco: Solving large instances of the RSA problem in flexgrid elastic optical networks, *J. Opt. Commun. Netw.* **8**(5), 320–330 (2016)
- 12.109 B. Kozicki, H. Takara, Y. Sone, A. Watanabe, M. Jinno: Distance-adaptive spectrum allocation in elastic optical path network (SLICE) with bit per symbol adjustment. In: *Proc. Opt. Fiber Commun./Natl. Fiber Opt. Eng. Conf. (OFC/NFOEC), San Diego, Mar. 21–25* (2010), Paper OMU3
- 12.110 B.T. Teipen, H. Griesser, M.H. Eiselt: Flexible bandwidth and bit-rate pro-grammability in future optical networks. In: *Proc. Int. Conf. Transpar. Opt. Netw. (ICTON), United Kingdom, Jul. 2–5* (2012), Paper Tu.C2.1
- 12.111 S. Yang, F. Kuipers: Impairment-aware routing in translucent spectrum-sliced elastic optical path networks. In: *Proc. 17th Eur. Conf. Netw. Opt. Commun., Villanova i la Geltru, June 20–22* (2012)
- 12.112 O. Rival, G. Villares, A. Morea: Impact of inter-channel nonlinearities on the planning of 25–100 Gb/s elastic optical networks, *J. Lightwave Technol.* **29**(9), 1326–1334 (2011)
- 12.113 T. Takagi, H. Hasegawa, K. Sato, Y. Sone, B. Kozicki, A. Hirano, M. Jinno: Dynamic routing and frequency slot assignment for elastic optical path networks that adopt distance adaptive modulation. In: *Proc. Opt. Fiber Commun. Conf./Natl. Fiber Opt. Eng. Conf. (OFC/NFOEC), Los Angeles, Mar. 6–10* (2011), Paper OTu17
- 12.114 H. Beyranvand, J. Salehi: A quality-of-transmission aware dynamic routing and spectrum assignment scheme for future elastic optical networks, *J. Lightwave Technol.* **31**(18), 3043–3054 (2013)

- 12.115 R.-J. Essiambre, G. Kramer, P.J. Winzer, G.J. Foschini, B. Goebel: Capacity limits of optical fiber networks, *J. Lightwave Technol.* **28**(4), 662–701 (2010)
- 12.116 T. Morioka, Y. Awaji, R. Ryf, P. Winzer, D. Richardson, F. Poletti: Enhancing optical communications with brand new fibers, *IEEE Commun. Mag.* **50**(2), 31–S42 (2012)
- 12.117 S. Fujii, Y. Hirota, H. Tode, K. Murakami: On-demand spectrum and core allocation for reducing crosstalk in multicore fibers in elastic optical networks, *J. Opt. Commun. Netw.* **6**(12), 1059–1071 (2014)
- 12.118 A. Muhammad, G. Zervas, D. Simeonidou, R. Forchheimer: Routing, spectrum and core allocation in flexgrid SDM networks with multi-core fibers. In: *Proc. Opt. Netw. Design Model., Stockholm, May 19–22* (2014) pp. 192–197
- 12.119 J.M. Fini, T. Taunay, B. Zhu, M. Yan: Low cross-talk design of multi-core fibers. In: *Proc. Conf. Lasers Electro-Opt. (CLEO), San Jose, May 16–21* (2010), Paper CTuAA3
- 12.120 A. Srivastava, S. Acharya, M. Alicherry, B. Gupta, P. Risbood: Differential delay aware routing for Ethernet over SONET/SDH. In: *Proc. IEEE INFOCOM, Miami, Mar. 13–17*, Vol. 2 (2005) pp. 1117–1127
- 12.121 S. Ahuja, M. Krunz, T. Korkmaz: Optimal path selection for minimizing the differential delay in Ethernet over SONET, *Comput. Netw.* **50**(13), 2349–2363 (2006)
- 12.122 X. Chen, Y. Zhong, A. Jukan: Multi-path routing in elastic optical networks with distance-adaptive modulation formats. In: *Proc. IEEE Int. Conf. Commun. (ICC), Budapest, Jun. 9–13* (2013)
- 12.123 W. Lu, X. Zhou, L. Gong, M. Zhang, Z. Zhu: Dynamic multi-path service provisioning under differential delay constraint in elastic optical networks, *IEEE Commun. Lett.* **17**(1), 158–160 (2013)
- 12.124 X. Chen, A. Jukan, A. Gumaste: Optimized parallel transmission in elastic optical networks to support high-speed Ethernet, *J. Lightwave Technol.* **32**(2), 228–238 (2014)
- 12.125 A. Castro, L. Velasco, M. Ruiz, J. Comellas: Single-path provisioning with multi-path recovery in flexgrid optical networks. In: *Proc. 4th Int. Workshop Reliab. Netw. Design Model., St. Petersburg, Oct. 3–5* (2012)
- 12.126 L. Ruan, N. Xiao: Survivable multipath routing and spectrum allocation in OFDM-based flexible optical networks, *J. Opt. Commun. Netw.* **5**(3), 172–182 (2013)
- 12.127 Y. Sone, A. Watanabe, W. Imajuku, Y. Tsukishima, B. Kozicki, H. Takara, M. Jinno: Highly survivable restoration scheme employing optical bandwidth squeezing in spectrum-sliced elastic optical path (SLICE) network. In: *Proc. Opt. Fiber Commun./Natl. Fiber Opt. Eng. Conf. (OFC/InfoECC), San Diego, Mar. 22–26* (2009), Paper OTh02
- 12.128 S. Huang, C.U. Martel, B. Mukherjee: Survivable multipath provisioning with differential delay constraint in telecom mesh networks, *IEEE/ACM Trans. Netw.* **19**(3), 657–669 (2011)
- 12.129 L. Choy: Virtual concatenation tutorial: enhancing SONET/SDH networks for data transport, *J. Opt. Netw.* **1**(1), 18–29 (2002)
- 12.130 G. Bernstein, D. Caviglia, R. Rabbat, H. Van Helvoort: VCAT–LCAS in a clamshell, *IEEE Commun. Mag.* **44**(5), 34–36 (2006)
- 12.131 S. Zhang, C. Martel, B. Mukherjee: Dynamic traffic grooming in elastic optical networks, *IEEE J. Sel. Areas Commun.* **31**(1), 4–12 (2013)
- 12.132 H. Takara, T. Goh, K. Shibahara, K. Yonenaga, S. Kawai, M. Jinno: Experimental demonstration of 400 Gb/s multi-flow, multirate, multi-reach optical transmitter for efficient elastic spectral routing. In: *Proc. Eur. Conf. Opt. Commun. (ECOC), Geneva, Sep. 18–22* (2011), Paper Tu.5.A.4
- 12.133 M. Jinno, H. Takara, K. Yonenaga, A. Hirano: Virtualization in optical networks from network level to hardware level, *J. Opt. Commun. Netw.* **5**(10), A46–A56 (2013)
- 12.134 T. Takagi, H. Hasegawa, K. Sato, Y. Sone, A. Hirano, M. Jinno: Disruption minimized spectrum defragmentation in elastic optical path networks that adopt distance adaptive modulation. In: *Proc. Eur. Conf. Opt. Commun. (ECOC), Geneva, Sept. 18–22* (2011), Paper Mo.2.K.3
- 12.135 A. Castro, L. Velasco, M. Ruiz, M. Klinkowski, J.P. Fernández-Palacios, D. Careglio: Dynamic routing and spectrum (re)allocation in future flexgrid optical networks, *Comput. Netw.* **56**(12), 2869–2883 (2012)
- 12.136 K. Wen, Y. Yin, D.J. Geisler, Sh Chang, S.J.B. Yoo: Dynamic on-demand lightpath provisioning using spectral defragmentation in flexible bandwidth networks. In: *Proc. Eur. Conf. Opt. Commun. (ECOC), Geneva, Sep. 18–22* (2011), Paper Mo.2.K.4
- 12.137 F. Cugini, F. Paolucci, G. Meloni, G. Berrettini, M. Secondini, F. Fresi, N. Sambo, L. Potì, P. Castoldi: Push-pull defragmentation without traffic disruption in flexible grid optical networks, *J. Lightwave Technol.* **31**(1), 125–133 (2013)
- 12.138 R. Wang, B. Mukherjee: Provisioning in elastic optical networks with non-disruptive defragmentation, *J. Lightwave Technol.* **31**(15), 2491–2500 (2013)
- 12.139 A.S. Muqaddas, M. Garrich, P. Giaccone, A. Bianco: Exploiting time synchronized operations in software-defined elastic optical networks. In: *Proc. Opt. Fiber Commun. Conf. (OFC), Los Angeles, Mar. 19–23* (2017), Paper W4J.6
- 12.140 A.N. Patel, P.N. Ji, J.P. Jue, T. Wang: Defragmentation of transparent flexible optical WDM (FWDM) networks. In: *Proc. Opt. Fiber Commun./Natl. Fiber Opt. Eng. Conf. (OFC/InfoECC), Los Angeles, Mar. 6–10* (2011), Paper OTuI8

Jane M. Simmons

Monarch Network Architects
Holmdel, NJ, USA
jsimmons@monarchna.com



Jane M. Simmons has been involved in the research and development of optical networks for 20 years. She is the founder of Monarch Network Architects. Prior to founding Monarch, she was the Chief Network Architect of Corvis Corp. and worked at Bell Labs/AT&T Labs Research. She received a BS, from Princeton University, and SM and PhD degrees from MIT, all in Electrical Engineering.

George N. Rouskas

Dept. of Computer Science
North Carolina State University
Raleigh, NC, USA
rouskas@ncsu.edu



George N. Rouskas is the Director of Graduate Programs in the Computer Science Department at North Carolina State University. He received a degree in Computer Engineering from the National Technical University of Athens (NTUA), Greece, and MS and PhD degrees in Computer Science from the College of Computing, Georgia Institute of Technology. His research interests include network architectures and protocols, optical networks, network design and optimization, and performance evaluation.