Optical Switching and Networking 13 (2014) 34–48

Contents lists available at ScienceDirect

Optical Switching and Networking

journal homepage: www.elsevier.com/locate/osn

Review

Spectrum management techniques for elastic optical networks: A survey ☆

Sahar Talebi [a], Furqan Alam [b], Iyad Katib [b], Mohamed Khamis [b], Reda Salama [b], George N. Rouskas [a,b,*]

[a] Operations Research and Department of Computer Science, North Carolina State University, Raleigh, NC 27695-8206, USA
[b] King Abdulaziz University, Jeddah, Saudi Arabia

ARTICLE INFO

Article history:
Received 2 December 2013
Received in revised form
4 February 2014
Accepted 9 February 2014
Available online 17 February 2014

Keywords:
Elastic optical networks
Spectrum management
Routing and spectrum assignment

ABSTRACT

In recent years, OFDM has been the focus of extensive research efforts in optical transmission and networking, initially as a means to overcome physical impairments in optical communications. However, unlike, say, in wireless LANs or xDSL systems where OFDM is deployed as a transmission technology in a *single link*, in optical networks it is being considered as the technology underlying the novel elastic network paradigm. Consequently, *network-wide spectrum management* arises as the key challenge to be addressed in network design and control. In this work, we review and classify a range of spectrum management techniques for elastic optical networks, including offline and online routing and spectrum assignment (RSA), distance-adaptive RSA, fragmentation-aware RSA, traffic grooming, and survivability.

Contents

☆ This work was supported in part by the National Science Foundation under Grant CNS-1113191, and in part by the Deanship of Scientific Research (DSR), King Abdulaziz University, under Grant no. 2-611-1434-HiCi.

* Corresponding author.
E-mail address: rouskas@ncsu.edu (G.N. Rouskas).

http://dx.doi.org/10.1016/j.osn.2014.02.003
1573-4277

1. Introduction

Optical networking technologies are crucial to the operation of the global Internet and its ability to support critical and reliable communication services. In response to rapidly growing IP traffic demands, 40 and 100 Gbps line rates over long distances have been deployed, while there is substantial research and development activity targeted to commercializing 400 and 1000 Gbps rates [1]. On the other hand, emerging applications, including IPTV, video-on-demand, and inter-datacenter networking, have heterogeneous bandwidth demand granularities that may change dynamically over time. Accordingly, mixed line rate (MLR) networks [2] have been proposed to accommodate variable traffic demands. Nevertheless, optical networks operating on a fixed wavelength grid [3] necessarily allocate a full wavelength even to traffic demands that do not fill its entire capacity [4]. This inefficient utilization of spectral resources is expected to become an even more serious issue with the deployment of higher data rates [5,6].

Elastic optical networks [7,8] have the potential to overcome the fixed, coarse granularity of existing WDM technology and are expected to support flexible data rates, adapt dynamically to variable bandwidth demands by applications, and utilize the available spectrum more efficiently [6]. The enabling technology for such an agile network infrastructure is orthogonal frequency division multiplexing (OFDM), and other efficient transmission techniques including Nyquist WDM and low-density parity-check (LDPC) based transmission [7]. OFDM, a modulation format that has been widely adopted in broadband wireless and copper-based communication systems, is a promising candidate for high-speed (i.e., beyond 100 Gbps) optical transmission [9]. Other key technologies include distance-adaptive modulation, bandwidth-variable transponders and flexible spectrum selective switches; for a recent survey of optical OFDM and related technologies, and how they impact network and control algorithm design, we refer the reader to [9].

OFDM is a multiple-carrier modulation scheme that splits a data stream into a large number of sub-streams [10]. Each data sub-stream is carried on a narrowband sub-channel created by modulating a corresponding carrier with a conventional scheme such as quadrature amplitude modulation (QAM) or quadrature phase shift keying (QPSK). The modulated signals are further multiplexed by frequency division multiplexing to form what is referred to as multicarrier transmission. The composite signal is a broadband signal that is more immune to multipath fading (in wireless communications) and inter-symbol interference. The main feature of OFDM is the orthogonality of subcarriers that allows data to travel in parallel, over sub-channels constituted by these orthogonal subcarriers, in a tight frequency space without interference from each other. Consequently, OFDM has found many applications, including in ADSL and VDSL broadband access, power line communications, wireless LANs (IEEE 802.11 a/g/n), WiMAX, and terrestrial digital TV systems.

In recent years, OFDM has been the focus of extensive research efforts in optical transmission and networking, initially as a means to overcome physical impairments in optical communications [11,12]. However, unlike, say, in wireless LANs or xDSL systems where OFDM is deployed as a transmission technology in a *single link*, in optical networks it is being considered as the technology underlying the novel elastic network paradigm [6]. Consequently, in the quest for a truly agile, resource-efficient optical infrastructure, *network-wide spectrum management* arises as the key challenge to be addressed in network design and control.

In this work we review and classify recent work in spectrum management in elastic optical networks. The paper is organized as follows. In Section 2, we discuss the advantages of, and challenges associated with, elastic optical networks. In Section 3 we define the offline routing and spectrum assignment (RSA) problem, and we discuss its complexity and various ILP formulations and heuristics. In Section 4 we describe and classify algorithms for the online RSA problem, and discuss performance modeling techniques. In Sections 5 and 6 we examine and categorize solution approaches based on distance-adaptive RSA (DA-RSA) and fragmentation-aware RSA (FA-RSA), respectively. In Section 7 we present traffic grooming techniques for RSA, in Section 8 we review survivability mechanisms for elastic optical networks, and in Section 9 we discuss multi-path variants of RSA. We conclude the paper in Section 10.

2. OFDM-based elastic optical networks

OFDM technology is the foundation of the elastic optical network (EON) concept [7], also referred to as "spectrum-sliced elastic optical path network" or *SLICE* [13]. The major difference between RWA and RSA lies in the SLICE network architecture as it flexibly adjusts to the format of the modulation [14]. The main driver of the EON architecture is the ability to allocate bandwidth at the granularity of an OFDM subcarrier rather than at the coarse unit of a wavelength in a fixed-grid network, using bandwidth-variable and format-agile transponders that may be reconfigured dynamically via software [10]. Optical signals are routed along the path to the destination by multi-granular optical switches that adapt to the data rate and center frequency of incoming channels via software control [15,16]. Bandwidth-variable transponders and switches make it possible to support efficiently a range of traffic demands, from sub- to super-wavelength, by *slicing off* just a sufficient amount of spectral

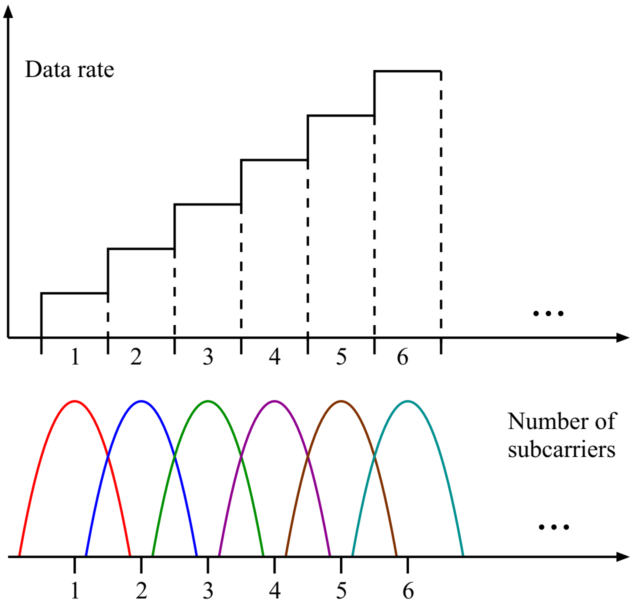

Fig. 1. Elastic data rates as a function of the number of allocated subcarriers (adapted from [19]).

resources along end-to-end paths to satisfy the client requirements. OFDM-based EONs have several advantages relative to existing WDM networks, including [7,9]

- *Resilience to physical impairments*: Since each subcarrier operates at a low symbol rate, inter-symbol interference is reduced and the effects of physical impairments are alleviated.
- *Elastic data rates*: The number of allocated subcarriers and the modulation format may be adjusted dynamically, on a per-connection basis, to account for (1) demand granularity, making it possible to support data rates from Gbps to Tbps, (2) path distance, so as to trade off spectrum utilization for reach, and (3) the time-varying nature of demands. Therefore, EONs may support multiple data rates, either by grouping together any number of subcarriers, as shown in Fig. 1, or by supporting a different data rate per subcarrier depending on network conditions. Importantly, EONs are highly scalable in that a transition to higher data rates would not require major changes in system design.
- *Spectral efficiency*: Two features of OFDM enable highly efficient use of spectral resources. On one hand, adjacent subcarriers may overlap in the spectrum due to their orthogonality, as shown in the bottom part of Fig. 1. This reuse of spectrum increases the overall system capacity. On the other hand, adapting the data rate to demand size, path length, and time variations achieves better use of existing spectrum.

While the finer granularity of bandwidth allocation is the key feature that makes OFDM attractive for future optical networks, it also introduces new and formidable challenges in the design and control of such networks. These challenges call for new spectrum management techniques that address effectively and efficiently issues related to

- *Scalability*: With an OFDM subcarrier as the unit of bandwidth allocation, the number of spectral resources to be managed network-wide is significantly larger than the number of wavelengths in existing WDM networks.
- *Spectrum contiguity*: If a demand requires t units of spectrum, then t *contiguous* subcarriers must be allocated to it. This constraint, not encountered in wavelength-based networks, may, unless appropriately accommodated, lead to severe fragmentation of the spectral resources that counters the inherent efficiency of fine-grain allocation.
- *Spectrum continuity*: The same t contiguous subcarriers must be allocated on each link along the end-to-end path of a demand. This constraint is analogous to the wavelength continuity constraint in WDM networks, and further contributes to potential spectrum fragmentation across the network links.
- *Variable data rates*: Support for elastic data rates, a core feature of EONs, requires precise tracking of the spectral width and center frequency of optical signals, and tight coordination of bandwidth-variable transponders and switches along end-to-end paths.

Due to the spectrum continuity constraint, there is a tight coupling between spectrum allocation and routing of a demand. Consequently, routing and spectrum assignment (RSA) [17,18] has emerged as the essential problem for spectrum management in EONs. Since the performance of a network depends not only on its physical resources (e.g., transponders, physical links, usable spectral width, optical switches, etc.) but also on how it is controlled, the objective of an RSA algorithm is to achieve the best possible performance within the limits of physical constraints. The RSA problem can be cast in numerous forms. The different variants of the problem, however, can be classified under one of the two broad versions: *offline RSA*, whereby the traffic demands are known in advance, and *online RSA*, in which a sequence of client requests arrive in some random fashion. The next two sections discuss spectrum management techniques for offline and online RSA, respectively.

3. The offline RSA problem

In the offline routing and spectrum assignment (RSA) problem, the input typically consists of a set of traffic demands, and the objective is to assign a physical path and contiguous spectrum to each demand so as to minimize the total amount of allocated spectrum (either over the whole network or on any link). Offline RSA arises whenever the traffic patterns in the network are reasonably well-known in advance and any traffic variations take place over long time scales. For instance, offline RSA is an effective technique for provisioning a set of semipermanent connections. Since these connections are assumed to remain in place for relatively long periods of time, it is worthwhile to attempt to optimize the way in which network resources (e.g., physical links and spectrum) are assigned to each connection, even though optimization may require a considerable computational effort.

We consider the following basic definition of the offline problem.

Definition 1 (*RSA*). Given a graph $G=(\mathcal{V},\mathcal{A})$ where $\mathcal{V}$ is the set of nodes and $\mathcal{A}$ the set of arcs (directed links), and a spectrum demand matrix $T=[t_{sd}]$ where t_{sd} is the amount of spectrum required to carry the traffic from source node s to destination node d, assign a physical path and spectrum to each demand so as to minimize the total amount of spectrum used on any link in the network, under three constraints: (1) each demand is assigned a contiguous spectrum (spectrum contiguity constraint), (2) each demand is assigned the same spectrum along all links of its path (spectrum continuity constraint), and (3) demands that share a link are assigned non-overlapping parts of the available spectrum (non-overlapping spectrum constraint).

RSA is a generalization of the well-known routing and wavelength assignment (RWA) problem [20]. If a single route for each source–destination pair is provided as part of the input, and each traffic demand is constrained to follow the given route, the RSA problem reduces to the spectrum assignment (SA) problem.

Definition 2 (*SA*). The RSA problem under the additional constraint that all traffic from source s to destination d must follow the given physical path r_{sd}.

3.1. Complexity results

A recent study [21] considered the complexity of the offline SA problem in chain (path) networks, in which no routing decision is involved. Using results from graph coloring theory, it was shown in [21] that the SA problem in paths is NP-hard, and that a $(2+\epsilon)-$approximation algorithm (where ϵ is an arbitrary real number that approaches zero) for computing the interval chromatic number of an interval graph may be used for solving the SA problem with the same performance bound. The study also extends this algorithm to solve the SA problem in ring networks with a performance bound of $(4+2\epsilon)$.

In [22], it was shown that the SA problem can be viewed as a problem of scheduling tasks on multiprocessor systems in which each task requires more than one processor simultaneously. Consider the following scheduling problem [23,24].

Definition 3 ($P|fix_j|C_{max}$). Given a set of n tasks and a set of identical processors, a processing time p_j and a prespecified set fix_j of processors for task $j, j=1,\dots,n$, schedule the tasks so as to minimize the makespan (or schedule length) $C_{max}=\max_j C_j$, where C_j denotes the completion time of task j, under the following constraints: (1) each processor can work on at most one task at a time, (2) each task must be processed simultaneously by all processors in fix_j, and (3) preemptions are not allowed.

It has been shown [24] that the three-processor problem $P3|fix_j|C_{max}$ is strongly NP-hard for general processing times, but that if the number of processors m is fixed and all tasks have unit times, i.e., $Pm|fix_j, p_j=1|C_{max}$, then the problem is solvable in polynomial time. Approximation algorithms and/or polynomial time approximation schemes (PTAS) have been developed for several versions of the problem [25].

Any instance of the SA problem can be transformed into an instance of the $Pm|fix_j|C_{max}$ problem, and hence, an algorithm for solving the latter problem also solves the former one [22]. However, the reverse is not true [22], i.e., there exist instances of $Pm|fix_j|C_{max}$ for which there is no corresponding instance of the SA problem.

By building upon this new perspective, it was shown in [22] that (1) $P3|fix_j|C_{max}$ transforms to the SA problem in a unidirectional ring with three links, hence the latter is NP-hard, and (2) the SA problem is solvable in polynomial time on paths with at most three links, but is NP-hard for paths with four or more links. The latter result confirms the conclusion in [21] that the SA problem is harder than the wavelength assignment problem which can be solved in polynomial time on paths.

3.2. ILP formulations and heuristics

Several variants of the RSA problem have been studied in the literature, and take into account various design aspects. Accordingly, a variety of integer linear program (ILP) formulations have been proposed, each tailored to a specific problem variant. Since the problem is intractable, these ILP formulations cannot be solved within a reasonable amount of time for problem instances corresponding to network topologies encountered in practice. Therefore, an array of heuristic algorithms have been put forward to obtain reasonably good solutions efficiently.

Link-based ILP formulations of RSA as a multicommodity flow problem have been studied in [18,26,27]. In addition to the spectrum contiguity, spectrum continuity, and non-overlapping spectrum constraints, the formulations also impose guard carrier constraints in allocating frequency slots (i.e., subcarriers) to the traffic demands. These studies consider three main objectives to minimize: the maximum number of subcarriers allocated on any fiber, the maximum subcarrier index allocated on any fiber, and the total number of subcarriers over all fibers. Upper and lower bounds on the optimal solution for ring and mesh networks were presented under both predetermined and non-predetermined routing. The lower bounds were obtained either using cut-set techniques or the even-load method, and the latter was shown to provide tighter bounds. It was also shown that, in ring networks with uniform demands, the lower and upper bounds are tight. Finally, two heuristic algorithms were developed to solve the RSA problem efficiently. The first algorithm, referred to as shortest path with maximum spectrum reuse (SPSR), uses shortest path routing and the first-fit spectrum allocation strategy to assign frequency slots to demands in decreasing order of their size. The second algorithm, balanced load spectrum allocation (BLSA), considers the k shortest paths as candidates for each demand, and selects the one that minimizes the maximum link load, so as to balance the use of spectrum across the network links.

A similar link-based formulation of an RSA variant referred to as the routing, wavelength assignment, and spectrum allocation (RWSA) problem was studied in [28,29]. In these studies, the network was assumed to support only a given set of line rates (e.g., 10/40/100/400/1000 Gbps), and each line rate requires a predetermined amount of spectrum that does not depend on the path length. The objective was to select a set of line rates for each demand and a corresponding set of lightpaths (one lightpath per line rate) so as to minimize the total spectrum width in the network. For each lightpath, the ILP formulation determines the center wavelength, the spectral width, and the physical route across the network. It was shown that RWSA is a generalization of the RWA problem in mixed line rate (MLR) optical networks. Due to the complexity of the problem, the ILP formulation was solved only for a six-node network in [28]. Consequently, the problem is decomposed in [29] into two subproblems: (1) the line rate selection problem that determines the line rates for each demand, and which is solved using dynamic programming, and (2) a variant of the routing and wavelength assignment (RWA) problem that seeks to establish lightpaths at the specific line rates. The latter subproblem is solved using three greedy heuristic methods.

A different, path-based ILP formulation of the RSA problem was presented in [17]. In a path-based formulation, k paths, where k is a small integer, are pre-computed for each demand, such that the demand may be routed only along one of these paths. The ILP is used to assign one of the predetermined paths to each demand, while also satisfying the spectrum contiguity and non-overlapping spectrum constraints; the spectrum continuity constraint is implicitly satisfied since the assignment of subcarriers is along a whole path. In [17], the objective is to minimize the number of subcarriers that are used in any link in the network. Although this path-based formulation is more compact than the link-based ILP of [18], the number of decision variables and constraints is substantially large that it cannot be solved directly. Accordingly, a heuristic algorithm called adaptive frequency assignment with collision avoidance (AFA-CA) was presented to select the path for each demand. The algorithm uses a link metric that captures the number of subcarriers that may potentially be allocated to a link, based on the paths that use this link. Then, it processes demands in an adaptive order that depends on previously allocated demands and the current usage of frequency slots in the network, and uses the link metric above to avoid selecting paths that will result in congested links.

Another path-based ILP formulation for the RSA problem was presented in [30,19], where the objective is to minimize the maximum subcarrier index assigned on any link in the network. In addition to the joint formulation, the study in [19] also presents a decomposition into two formulations that address the routing (R) and spectrum assignment (SA) aspects of the problem, respectively, and is referred to as R+SA. The first formulation takes as input a predetermined set of paths for each demand, and selects one path per demand so as to minimize the amount of traffic flow (i.e., spectrum use) on any link. In the second phase, the SA formulation assigns frequency slots to demands so as to minimize the maximum subcarrier index. While each ILP formulation in the decomposition is more compact and scalable than the joint ILP, such a sequential solution is not guaranteed to yield an optimal solution to the joint problem. A greedy heuristic algorithm is also presented that processes demands in decreasing order of either their size or their shortest path length; this order is fixed, unlike the heuristic in [17] that adapts the order as the algorithm progresses. A simulated annealing meta-heuristic that builds upon the greedy algorithm was also presented. Finally, the spectrum utilization in an elastic network was compared to that in a fixed-grid WDM network in [30].

The above approaches model the spectrum contiguity constraint by including a set of problem constraints that increase significantly the complexity of the ILP formulation. To overcome this difficulty, the concept of a “channel” as a set of contiguous subcarriers of a given width was introduced in [31]. For a given spectrum width t, all possible channels with t subcarriers (i.e., one starting at the first subcarrier, one starting at the second subcarrier, and so on) are defined on each link of the network. Then, the RSA problem is transformed to one of “routing and channel allocation,” in which channel assignment implies allocation of contiguous spectrum, and no explicit spectrum contiguity constraints are needed. The result is a more compact formulation that achieves a significant speed-up in running time compared to those that directly account for the spectrum continuity constraints [31].

4. The online RSA problem

Under a dynamic traffic scenario, clients submit to the network requests for optical paths to be set up as needed. Thus, connection requests are initiated in some random fashion. Depending on the state of the network at the time of a request, the available spectral resources may or may not be sufficient to establish a connection between the corresponding source–destination node pair. The network state consists of the physical path (route) and spectrum assignment for all active connections. The state evolves randomly in time as new connections are admitted and existing connections are released. Thus, each time a request is made, an algorithm must be executed in real time to determine whether it is feasible to accommodate the request, and, if so, to perform routing and spectrum assignment. If a request for a connection cannot be accepted because of lack of resources, it is blocked. Therefore, the blocking probability of connection requests arises as the key performance metric of interest in an online RSA scenario.

4.1. Heuristic algorithms

Because of the real-time nature of the problem, RSA algorithms in a dynamic traffic environment must be simple and fast. Since combined routing and spectrum assignment is a hard problem, most studies devise heuristic algorithms. Online RSA algorithms can be broadly classified into two categories depending on whether they tackle the routing and spectrum assignment aspects jointly (i.e., in one step) or separately (i.e., in two steps).

4.1.1. Two-step algorithms

Two-step online RSA heuristics decompose the problem into two subproblems, the routing problem and the spectrum assignment subproblem, which are then solved sequentially. Specifically, such an approach consists of the following steps:

- Compute a number of candidate physical paths for each source–destination node pair and arrange them in a path list.
- Starting with the path at the top of the corresponding list, use a spectrum allocation policy to assign a feasible path and set of contiguous slots, if they exist, for the requested connection.

The routing algorithm may be *static*, in which case the paths are computed and ordered offline, or *adaptive*, in that the paths computed and their order may vary according to the current state of the network. The number of path choices is another important parameter. Often, as in [32], multiple alternate paths are computed for each request. The spectrum allocation policy determines which set of available contiguous slots are assigned to a request, and is crucial to the performance of an online RSA algorithm. A *first-fit* policy [32,33] selects the lowest index set, a *random-fit* policy [34] randomly allocates one of the available sets, whereas a *best-fit* policy selects the smallest set that can satisfy the request. Note that an improvement in the operation of the first-fit policy has been proposed in [35]. This study proposed an evolutionary algorithm to search for the most feasible spectrum ordering for first-fit so as to minimize the blocking probability; the study showed that the algorithm performs far better than the conventional first-fit policy.

The study in [33] investigated the optimal slot width for EONs by measuring the blocking probability under dynamic traffic using Monte Carlo simulations. Each demand was routed on its shortest path, and the first-fit policy was used for spectrum allocation. The main finding of the study was that the best performance is achieved when the slot width is equal to the greatest common factor of the spectrum widths of the data rates supported in the network. On the other hand, the KSP-based RSA algorithm proposed in [32] first computes k shortest paths for a given request, and then searches each path to find the required number of contiguous subcarriers using the first-fit policy.

The two-step online RSA algorithm in [36] proposes an interesting adaptive routing algorithm to solve the first subproblem. Specifically, routing tables at each node associate a probability with each (next hop, destination) pair. A dynamic ant colony optimization (ACO) algorithm is run continuously that updates these probabilities based on spectrum usage information on each link collected by ants as they traverse the network. When a request arrives, a path is selected starting at the source and visiting next-hop nodes based on the current values of the probabilities stored at the routing tables until the destination has been reached. Once the path has been determined, the first-fit policy is used to assign frequency slots to the request. This ACO-based online RSA algorithm compares favorably to the KSP-based algorithm in [32].

4.1.2. One-step algorithms

One-step online RSA heuristics, on the other hand, solve the two subproblems (i.e., routing and spectrum assignment) simultaneously but sub-optimally, typically using a greedy approach. In [37], a dynamic version of the RWSA problem studied in [28,29] is considered, and is referred to as D-RWSA. As with RWSA, D-RWSA was decomposed into a rate selection problem (similar to the one we discussed earlier), and a dynamic routing and channel selection problem that is a variant of online RSA (in this work, the notion of a channel is similar to that defined in [31]). This latter problem was solved using an auxiliary graph that represents the state (i.e., established connections) of the network. With this representation, finding a path and contiguous spectrum for a new request reduces to finding a shortest path on the auxiliary graph.

Two different one-step online RSA algorithms were introduced in [32]. In the first, a modified Dijkstra shortest path algorithm is employed, in which, every time a link is considered to extend the shortest path to the destination node, the link is checked to ensure that it has available contiguous spectrum in common with links already in the path. If so, the link is added, otherwise the algorithm considers other links even if they have a higher cost. The second one-step RSA algorithm builds a path vector tree routed at the source node such that all links along each path have sufficient contiguous spectrum for the request. The algorithm then searches the path vectors to find one that has available spectrum and minimum cost; this algorithm is more computationally demanding but results in better blocking performance. A modified Dijkstra algorithm similar to the one in [32] was proposed in [38] to solve the online RSA problem in one step. The main difference is that the spectrum availability along the path is maintained by keeping track of the channels (as defined in [31]) that may fit the given connection request.

4.2. Performance models

Evaluating the performance (e.g., in terms of blocking probability, spectrum utilization, or fragmentation rate) of an EON operating under a given online RSA algorithm requires analytical models that are able to capture accurately the evolution of the system's state. However, developing such models is a challenging task due to the nature of the RSA problem. The spectrum continuity constraint introduces load correlation between subcarriers in adjacent links, similar to the wavelength continuity constraint in current WDM networks. In addition, the spectrum contiguity constraint introduces correlation between subcarriers in the same link. Hence, it is not possible to use multiclass $M/M/K/K$ models directly, as these do not account for the fact that a new connection must occupy a set of *contiguous* servers simultaneously.

An exact analytical model for a *single* link of an EON has been developed in [34]. The model captures three important features of such a system: the load correlation among subcarriers due to the contiguity requirement; the fact that subcarriers are not equivalent (e.g., the lowest and highest index subcarriers are not equivalent to other subcarriers since they have contiguous subcarriers on one side only); and the existence of multiple classes of calls, each requiring a different spectral width. Under the assumption of Poisson arrivals, two continuous-time Markov chain models were developed, one each for the first-fit and random-fit allocation policies. Since the state space grows exponentially with the number of subcarriers, two heuristic algorithms were also developed to compute the stationary distribution from which the blocking probability and other parameters may be obtained.

Exact and approximate Markov chain models for a special version of a spectrum allocation problem are developed in [39]. This study considers demands with time-varying requirements and proposes spectrum extraction and contraction policies to adjust dynamically the amount of spectrum assigned to each demand in response to these requirements. The analytical models are used to evaluate the blocking performance of the various policies. These blocking models are then used within an iterative RSA algorithm to minimize the average blocking probability.

5. Distance-adaptive RSA (DA-RSA)

The design of commercial WDM networks has traditionally focused on optimizing the transmission performance in the worst case [40]. Specifically, for a given data rate (e.g., 40 Gbps), the modulation format is determined so as to ensure that the signal is transmitted with sufficient quality along the worst path (typically, the longest path and/or the one with the most hops) in the network. Then, the same format is used for every, say, 40 Gbps demand. In other words, each such demand is assigned the exact same amount of optical spectrum irrespective of its path length, the number of nodes it traverses, or the level of impairments it encounters. Engineering for the worst-case scenario leads to low spectrum utilization given that the transmission quality along most paths is far better than that in the worst case.

In OFDM-based networks, on the other hand, it is possible to adjust the modulation format and/or number of bits per symbol to account for link impairments (e.g., available signal-to-noise ratio (SNR)) so that demands with the same data rate are allocated different amounts of spectral resources depending on the quality of their path [10,41–45]. Distance-adaptive (DA) spectrum allocation, a concept first introduced in [40], exploits the tradeoff between spectrum width and reach (for the same data rate) to improve utilization [46] by tailoring the modulation format to the level of impairments: a high-level modulation format with narrow spectrum and low SNR tolerance may be selected for a short path, whereas a low-level modulation with a wider spectrum and high SNR tolerance may be used for a longer path [47]. In fact, it has been argued [40] that the utilization of spectral resources depends not only on the network size and topology, node and link characteristics, or traffic pattern, but also on the specific RSA algorithm employed.

An offline version of the DA-RSA problem, referred to as the routing, modulation level, and spectrum allocation (RMLSA) problem, was studied in [48]. In this problem, each demand is mapped to a modulation level based on the requested data rate and the distance of the path over which it is routed, with the mapping function provided as input to the problem. A path-based ILP formulation for RMLSA was first provided, and then the problem was decomposed into two subproblems, routing and modulation level (RML) and spectrum assignment (SA) and solved sequentially (RML+SA) using ILPs. Finally, a greedy heuristic and a simulated annealing meta-heuristic, similar to the ones in [19], were also presented.

Most studies of DA-RSA consider the online version of the problem and develop heuristic algorithms to accommodate randomly arriving connection requests. Several algorithms follow a two-step approach similar to the one we discussed in Section 4.1.1 for the basic (i.e., non-distance-adaptive) online RSA problem. The main difference is that a spectrum allocation model is used to determine the number of subcarriers as a function of data rate and path length. Therefore, in the second step of the algorithm, the number of slots that the spectrum allocation policy must search for varies depending on the length of the path considered.

Two-step heuristics for the online DA-RSA problem were presented in [40,49]. These algorithms compute a number of fixed-alternate paths (i.e., they use static routing) for each source–destination pair, and order them in decreasing length. In the second step, they employ the first-fit spectrum allocation policy and sequentially consider each path until a number of contiguous frequency slots that can accommodate the requested data rate over the given path length is found; if no spectral resources are available on any of the paths in the list, the request is blocked. The main difference between the two algorithms is that the latter only considers paths for which the number of slots is the same as for the shortest path. A similar algorithm was used in [43] for routing of superwavelength demands.

A version of the online DA-RSA problem referred to as dynamic impairment aware routing and spectrum allocation (IARSA) was studied in [47]. The objective was to select a feasible path for each request and allocate subcarriers by using an appropriate modulation format with the transmission reach for the requested data rate. Two variants of IARSA were investigated, one in which regenerators may modify the modulation format of the incoming signal, and one in which the modulation format does not change in the network. The heuristic algorithm used to solve the IARSA problem works as follows. For each modulation format, each link in the network is assigned a weight equal to the ratio of the required spectrum over the number of free slots on this link. For each modulation format, a modified version of Dijkstra's algorithm is used to find a minimum cost path with sufficient contiguous spectrum for this request; this approach is similar to the one-step heuristics for the basic online RSA problem we discussed in Section 4.1.2. Finally, the path and modulation format with the smallest cost (if any is found) is assigned to the request.

A quality of transmission (QoT) aware online RSA technique was proposed in [14], consisting of three stages: path calculation, path selection, and spectrum assignment. The Dijkstra and *k*-shortest algorithms were adapted for computing paths, while fiber impairments and non-linearity effects at the physical layer were modeled using a closed-form expression that estimates the QoT along a given path. For each request, the most feasible route is chosen with respect to the optical signal-to-noise ratio (ONSR), as determined by the physical layer model.

6. Fragmentation-aware RSA (FA-RSA)

Whereas the ability to allocate variable data rates is an attractive feature of EONs in terms of supporting heterogeneous applications, two challenges arise in operation when connections arrive and depart dynamically [50]:

- *Fragmentation*: Fragmentation of spectral resources emerges as allocation and de-allocation of blocks of contiguous slots on demand may cause part of the spectrum to become unusable. There are two sources of stranded spectrum in EONs [50]. The spectrum continuity constraint causes *horizontal* fragmentation in that the same block of spectrum may not be available along successive links of a path despite the fact that each link may have sufficient bandwidth for a request; this issue is similar to wavelength fragmentation in WDM networks. Variable data rates along with the spectrum contiguity constraint, on the other hand, are the cause of *vertical* fragmentation, a situation whereby the spectral resources on a single link are fragmented into small non-contiguous blocks that cannot be allocated to a single large demand.
- *Fairness*: If spectral resources become fragmented across the network links, heterogeneous connection requests will experience blocking rates that depend strongly on their data rate and/or path length. Due to vertical fragmentation, large contiguous blocks of slots may become sparse, hence high-rate connections will be more likely to be rejected than the low-rate ones. On the other hand, for a given amount of spectrum, horizontal fragmentation makes it more difficult to find continuous blocks on long paths compared to the short ones. Therefore, if left unchecked, fragmentation may lead to starvation of high-rate and/or long-path connections.

Fragmentation-aware RSA algorithms attempt to improve the blocking performance, fairness, and spectrum utilization of EONs by minimizing the extent of spectrum fragmentation. The first step in such an approach is to develop metrics that quantify the degree of fragmentation in the network. The utilization entropy concept introduced in [51] assesses the spectrum fragmentation of a link by counting the number of neighboring pairs of slots that have different status (i.e., one slot is used but the other is free), and normalizing this number to get a value in $(0,1)$ such that low (respectively, high) values represent low (respectively, high) vertical fragmentation. A similar concept was defined to capture horizontal fragmentation along a path by considering the status of a given slot on pairs of successive links in the path. The fragmentation ratios proposed in [52,53] to measure the bandwidth fragmentation of a link or path are inspired by similar ratios proposed for storage systems, and take into account not only the status of neighboring slots or that of a slot on adjacent links, but also the size of free blocks of slots. In [50], each contiguous block of unused slots is assigned a value equal to the maximum data rate that it can support, and the fragmentation index of a path is computed by taking the ratio of the sum of the value of available blocks to the value of the sum of the slots in all blocks; the latter sum represents the value of these slots as if they were contiguous. The fragmentation index concept is extended to capture spectrum fragmentation over the whole network by taking the average of the fragmentation indices of the shortest paths between every source–destination pair.

Fragmentation-aware RSA algorithms may be classified as *proactive* or *reactive* [50], as we discuss next.

6.1. Proactive FA-RSA

Proactive FA-RSA techniques attempt to prevent or minimize spectrum fragmentation at the time a new request is admitted to the network. Since support for variable data rates is a main contributor to fragmentation, Wang and Mukherjee [50] identify four spectrum management techniques for allocating spectrum to connections of different data rates. With “complete sharing,” all connections share the whole spectrum using the first-fit policy. With “pseudo partition,” low-rate connections are allocated bandwidth from one end of the spectrum (using first-fit), while high-rate connections are allocated from the other end (again using first-fit). Under the “dedicated partition” scheme, spectrum is partitioned and each partition is dedicated to serving connections of a given data rate; hence, vertical fragmentation due to variable rates is eliminated and each partition reduces to a wavelength-routed network. Finally, “shared partition” is a generalization of the previous scheme in that spectrum is partitioned but higher data rate connections may access partitions assigned to lower data rate ones. It was shown in [50] that the dedicated and shared partition schemes improve both fairness and fragmentation compared to complete sharing or pseudo partition.

A similar concept of spectrum reservations is studied in [54]. Instead of partitioning spectrum such that each partition is shared only among connections with the same rate, Christodoulopoulos et al. [54] propose that a block of contiguous subcarriers be reserved for each source–destination pair. In addition, subcarriers that are not reserved may be shared on demand among all connections. This study assumes that demands vary with time, but as long as a demand stays within its reservation, it can always be accommodated. However, if a connection requires additional bandwidth, an FA-RSA algorithm is executed to allocate shared subcarriers along one of a set of candidate paths.

Two FA-RSA algorithms that make spectrum allocation decisions based on the current state of fragmentation were introduced in [53]. The first algorithm assigns each new request to a path that minimizes a network-wide fragmentation ratio defined in the same study. The second algorithm attempts to utilize slots that are already used the most in the network. The maximize common large segment (MCLS) algorithm in [55] generates a number of candidate paths for a request. For each candidate path, the algorithm considers all the links in the path (i.e., candidate links), as well as links that are adjacent to the candidate links at any node on the path. The algorithm then computes a metric that captures the availability of contiguous slots between candidate and adjacent links after spectrum for this request has been allocated. Finally, it selects the path and spectrum allocation that has the smallest value for this metric, so as to maximize the probability that future requests will find a sufficient number of contiguous slots to be accepted.

We also note that all the FA-RSA algorithms above [54,53,55] can be classified as two-step heuristics based on our discussion in Section 4.1.1.

6.2. *Reactive FA-RSA*

Recognizing that fragmentation may not be completely eliminated in a dynamic environment, reactive FA-RSA algorithms employ defragmentation techniques to restore the network's ability to accommodate high-rate and long-path connections. The objective of defragmentation is to rearrange the spectrum allocation of existing traffic demands so as to consolidate available slots into large contiguous and continuous blocks that may be used to establish future requests. Defragmentation strategies may be broadly classified as *periodic* or *path-triggered* [38]. Periodic defragmentation runs at long time scales and is initiated either at regular intervals or whenever a metric indicative of network-wide fragmentation exceeds a certain threshold. The process tackles the fragmentation of spectral resources across the whole network, and hence it is computationally expensive and may disrupt ongoing connections along large parts of the network. The scope of path-triggered defragmentation, on the other hand, is more narrow. It is invoked when the online RSA algorithm is unable to find adequate resources to satisfy a new traffic demand, with the objective of assembling continuous/contiguous blocks of spectrum sufficient for this demand. Therefore, path-triggered defragmentation is executed at shorter time scales but only disrupts connections sharing a link with the new request. With either strategy, a make-before-break rerouting (MBBR) [44] scheme may be used to minimize the disruption to existing connections as they are moved to a new spectrum block. Spectrum defragmentation has been demonstrated experimentally in [56].

The network-wide defragmentation problem was studied in [57] under two objectives. The problem was defined as one of the rearranging existing connections so as to minimize the total spectrum required for these connections (and, hence, maximize spectrum consolidation for future connections). However, this objective may be achieved with an offline RSA approach, and does not take into account the disruption to connections caused by the rearrangement. Therefore, the study also considered the secondary objective of minimizing the number of connections interrupted during the reconfiguration process. The problem was formulated as an ILP, and two heuristics were proposed to solve it.

A path-triggered defragmentation strategy was proposed and investigated in [38]. When a new request arrives, an online RSA algorithm is first run to accommodate it. If a path with sufficient spectrum cannot be found, then a defragmentation algorithm is triggered that consists of the following steps. First, the number of available (but not necessarily continuous or contiguous) slots along each of a set of candidate paths is computed. If a path with a sufficient number of free slots does not exist, then the request is blocked. Otherwise, an attempt is made to create a block of continuous and contiguous slots along the shortest path with sufficient available resources. To this end, connections that share links with this shortest path must be reallocated a different block of spectrum. The problem of identifying the connections to be rearranged and the new spectrum allocation (for the old connections and the new request) was formulated as an ILP. The objective was to minimize the number of connections to be rearranged, under the constraint that the routes of existing connections are not modified. In order to solve the problem for large networks, a greedy randomized adaptive search (GRASP) meta-heuristic was proposed.

The spectral defragmentation algorithm in [58] consists of two phases. When a new request arrives but cannot be accommodated, a greedy heuristic is executed in the first phase to identify the spectrum block for the new request that minimizes the number of ongoing connections that will be affected. In the second phase, the problem of assigning new spectrum to the affected connections is formulated as a bipartite matching problem. If a perfect matching (i.e., one that satisfies all connections) exists, then existing connections are moved to the new spectrum and the request is assigned the spectrum block selected in the first phase; otherwise, the request is blocked or the algorithm repeats from the first phase to determine a new spectrum block for the request.

A set of reactive defragmentation strategies that capitalize on hitless optical path shift (HOPS) were described in [59]. HOPS technology, assessed in [60] where it is referred to as “push-pull defragmentation,” allows a connection to shift to a new block of spectrum as long as the route of the connection does not change and the move to the new spectrum does not affect other established connections. A simple technique was proposed in [59] to consolidate the spectrum freed by a terminated connection with other blocks of spectrum available along the links of its path. Specifically, upon the departure of a connection, ongoing connections that share a link with the just terminated connection are shifted, whenever possible, to the lower end of the spectrum, thus creating larger available blocks at the higher end for future connections. A more complex algorithm was proposed to create continuous and contiguous blocks of spectrum for new requests whenever they cannot be accommodated by a basic online RSA algorithm that is applied first.

7. Traffic grooming with RSA (TG-RSA)

Traffic grooming [61] is an optimization problem that arises in the design and control of networks with multigranular traffic. In conventional WDM networks, the main objective is to aggregate sub-wavelength demands onto lightpaths of fixed capacity so as to improve the utilization of wavelengths and reduce the number of add/drop ports that represents a major cost for the network [62]. The aggregation of traffic takes place at specific grooming nodes that are analogous to hub airports where passengers are "groomed" onto fixed-capacity airplanes.

In EONs, lightpaths (or "LambdaFlex connections" [63]) may be set up with the exact amount of spectrum to meet the data rate and reach for a traffic demand. Nevertheless, traffic grooming in EONs has the potential to improve spectrum utilization as well as yield significant cost reduction in terms of transponders. Specifically, aggregating demands for transport on a single lightpath presents two opportunities for spectrum savings. First, depending on the modulation format, transporting a single high-rate (e.g., 400 Gbps) connection is generally more spectrum efficient than transporting a number of low-rate connections with the same total data rate (e.g., 4 × 100 Gbps) [64]. Second, a smaller number of independent lightpaths results in a smaller number of guard bands around these lightpaths, hence reducing the spectrum required to support a given set of demands (or, equivalently, increasing the amount of spectrum available to carry client traffic). Importantly, in EONs it is possible to perform traffic grooming at the optical layer thereby eliminating expensive O/E/O operations. Optical traffic grooming involves setting up an optical tunnel [63] that carries several LambdaFlex connections in a contiguous block of spectrum; the bandwidth-variable transceiver and switch requirements to implement such optical tunnels are discussed in [65].

Once grooming considerations are included in the RSA problem (whether online or offline), the new TG-RSA problem becomes significantly more difficult to tackle optimally. Mixed integer linear programming (MILP) formulations of offline TG-RSA variants have been presented in [65–67]. The objective in [65,66] was to minimize the total amount of spectrum used on any link, whereas the objective function of [67] is an average spectrum utilization rate weighted by fiber length. The formulations in [65,66] are path-based, and, in fact, the formulation presented in [65] leverages the formulation of the basic offline RSA in [30] and adds grooming-specific constraints. Two TG-RSA heuristics are presented in [65] to solve the problem efficiently on realistic networks. The least spectrum grooming (LSG) heuristic attempts to minimize the spectral resources by grooming demands with the same source and paths that share the most links. The minimum transmitter grooming (MTG) algorithm is similar to LSG but its goal is to minimize the number of transponders.

The online TG-RSA problem has been addressed in [68], and an auxiliary graph was used to model the problem. Specifically, each physical node is represented using several layers in the auxiliary graph: the transponder layer captures the O/E/O conversion (i.e., grooming) ability and the spectrum layers represent the available spectrum resources. The spectrum continuity constraint is enforced by having edges between the same spectrum layers of adjacent nodes. The spectrum contiguity constraint is enforced by deleting edges between spectrum layers if the amount of contiguous spectrum is not sufficient to carry a demand. A path for a demand is found by running a shortest path algorithm on the auxiliary graph representing the state of the network at the time a request arrives, as well as the data rate requirement of the request. This auxiliary graph approach makes it possible to apply a range of grooming policies (e.g., minimize the number of transponders, the number of physical or logical links, etc.) by setting the weights of the edges appropriately.

The auxiliary graph model of [68] was used in [69] to find both a working and a backup lightpath in a dynamic traffic grooming scenario under shared protection. The working path is selected among the k shortest paths in the auxiliary graph, while the backup path is link-joint with the working path. This study proposes a protection scheme referred to as elastic separate-protection-at-connection (ESPAC) that allows the backup lightpaths of two connections to share spectrum if the respective working lightpaths are link-disjoint. ESPAC maximizes the opportunities for spectrum sharing by using the first-fit policy to assign spectral resources, starting with one end of the spectrum for working lightpaths and on the other end for backup lightpaths. Protection and restoration schemes for EONs are discussed in more detail in the following section.

8. Survivability and RSA

OFDM lays the foundation for EONs to support individual data rates of 400–1000 Gbps [1] and, hence, aggregate throughput per fiber of potentially tens to hundreds Tbps. Since the failure of even a single network element (e.g., a fiber link) is likely to affect numerous connections and result in immense data loss, network survivability assumes critical importance. Survivability refers to the ability of the network to reconfigure itself so as to restore the connections affected by a failure. In an optical network, three types of failures are generally considered: link failures (e.g., caused by cable cuts), node failures (e.g., due to equipment malfunction at a switch or router), and channel failures (e.g., caused by the failure of transmitting or receiving equipment specific to that channel).

Several survivability mechanisms have been explored in WDM networks [70] and can be classified in one of the two broad categories: *protection* or *restoration*. Protection schemes are carried out at the network design or planning phase in *anticipation* of a network failure. Backup resources reserved by these schemes may be *dedicated* to a single connection or *shared* among multiple connections. During normal network operation, reserved resources remain idle. Upon occurrence of a failure, affected connections are redirected to reserved resources according to the plan determined at the planning phase (typically, at the time a connection was set up). A restoration scheme, on the other hand, does not reserve resources for a connection at the time of establishment. Instead, resources for recovering a connection are discovered dynamically immediately *after* the occurrence of a failure affecting the connection. With either protection or restoration, a recovery scheme may be either *failure-independent* (also referred to as path-based) or *failure-dependent* (also referred to as link-based). In failure-independent recovery, the source and destination nodes of each affected connection take action to switch to a backup path that is disjoint with the corresponding failed working path, regardless of the location or type of failure. In contrast, with failure-dependent recovery, it is the nodes at either side of the failure (e.g., the endpoints of a failed link) that redirect all affected connections around the failure without intervention of the various source and destination nodes. As was remarked in [71], similar survivability mechanisms may be applied to EONs.

8.1. Protection

The problem of ensuring dedicated or shared protection of established connections in an EON may be viewed as a variant of the offline and online RSA problems with additional constraints to account for the backup paths and the sharing (if any) of backup spectrum. The following studies all consider failure-independent mechanisms. Note that the ESPAC algorithm [69] we discussed in Section 7 also implements a failure-independent shared-protection scheme in a traffic grooming context.

Two-fiber ring networks with 1+1 dedicated and 1:1 shared protection were considered in [72], and DA-RSA algorithms were proposed to assign spectrum to the working and backup paths. In both cases, the shortest (respectively, alternate) path around the ring was assigned as working (respectively, backup). With 1+1 dedicated protection, the same set of frequency slots are assigned to the working and backup paths of a connection independently of other connections. With 1:1 shared protection, on the other hand, the algorithm assigns the same set of slots to the working paths of several connections as long as they are pair-wise disjoint, and similarly for the backup paths of the same set of connections.

The offline RSA problem was considered in [73] under the additional constraint that all demands be assigned two disjoint paths, one working and one backup, and that spectral resources on backup paths be shared among connections with disjoint working paths. A MILP formulation of the problem was developed with the objective of minimizing the amount of spectrum allocated on the most congested link. Furthermore, a heuristic algorithm was proposed. The algorithm finds the k shortest *cycles* for each demand, and processes demands in decreasing order of shortest cycle length. It uses the first-fit policy to assign spectrum for the demand on each cycle (hence, on the working and backup paths), and selects the one that yields the lowest congestion on any link of the cycle.

The online RSA problem with dedicated protection was studied in [74]. The objective in this case is to find a working and backup path for each arriving connection request that are link-disjoint; the two paths are assumed to share transponders, hence they must be assigned the exact same spectrum block. The problem is solved by first constructing a set of auxiliary graphs that represent the network state, and then running Suurballe's algorithm on each graph to find the shortest pair of disjoint paths for the request. The i-th auxiliary graph includes an edge between two nodes if the corresponding fiber link has a number of contiguous slots starting at slot i that may satisfy the demand. The algorithm selects the shortest pair of paths among the ones constructed.

The study in [75] considered the online RSA problem with shared protection, and introduced two policies for spectrum sharing among backup paths. Under the conservative policy, sharing is allowed only if the backup paths have the same bandwidth, whereas under the aggressive policy, sharing is allowed even if the backup paths have different bandwidths. Of course, with the aggressive policy, the number of resources allocated to the backup path equals the maximum bandwidth on the working paths that share this backup path. It is noted that, while the aggressive policy leads to more sharing opportunities, it may fragment the spectrum along backup paths. A heuristic algorithm was proposed that first computes the working path and then the backup path (each selected from respective k shortest paths), and allocates spectrum using the first-fit policy.

A recent study [76] considered three types of networks (single- and multi-rate WDM networks, and EONs) and compared them in terms of cost and energy efficiency under dedicated (1+1 or 1:1) and shared protection. The respective RWA and RSA problems were solved using heuristics. The heuristics are similar to the ones discussed earlier in this section in that candidate pairs of working/backup paths for each connection were computed using a k-shortest path algorithm. The main difference is that each candidate pair was evaluated using a metric that accounts for power consumption on both paths.

8.2. Restoration

A new survivability scheme called bandwidth squeezed restoration (BSR) was first proposed in [77] and was further refined in [78]. BSR takes advantage of the variable, fine-granularity spectrum allocation possible in EONs, and may be adapted for operation under either a pre-planned protection or a dynamic restoration mechanism. In its original form [77], BSR assumes that the data rate assigned to a connection is the sum of a committed rate and an excess rate. Consequently, recovery does not start until a failure causes the data rate of a connection to fall below the committed rate. At that time, a new backup path is selected (either from an existing protection plan or discovered dynamically) that can support the committed rate of the connection. As part of the recovery process, the bandwidth allocated to some connections that are not affected by the failure (e.g., connections with a working path that overlaps with the backup path of an affected connection) may be *squeezed* to the corresponding committed rate, so as to provide maximum survivability for the given failure.

The BSR scheme was extended in [78] to allow for three recovery scenarios: (i) full bandwidth guaranteed recovery (FBGR), (ii) partial bandwidth guaranteed recovery (PBGR), or (iii) best-effort recovery (BER). FBGR is a conventional recovery scheme, while PBGR is akin to restoring the committed rate as in [77]. The BER scheme is unique to EONs and makes it possible to allocate any amount of spectral resources available at the time of failure so that a connection may proceed even at a low rate. Since the amount of available resources must be discovered after the failure, even if the backup path has been pre-planned, BSR may be classified as a dynamic restoration scheme.

The dynamic load balancing shared-path protection (DLBSPP) algorithm presented in [71,79] is a hybrid protection/restoration algorithm. Similar to shared protection schemes we discussed in the previous subsection, DLBSPP protects from single failures by allocating a working and backup path to each connection at setup time. An interesting feature of the algorithm is that, for the purpose of path computation, link weights are set proportional to the number of available frequency slots, in an attempt to balance the load across the network links. On the other hand, DLBSPP adopts a restoration strategy to recover from the simultaneous failure of multiple links. In the event of a multi-link failure, the algorithm attempts to restore all connections for which both the working and backup paths become unavailable. To this end, affected connections are prioritized in decreasing order of their spectrum demands, and the algorithm computes, on the fly, a new path for each connection, starting with the connection with the highest priority (i.e., the highest data rate). The motivation for this approach, referred to as traffic adaptive restoration (TAR), is based on the observation that it is more difficult to find paths for high-rate connections, hence these should be accommodated before allocating resources to low-rate connections.

A pure restoration scheme for EONs was presented in [46] as part of a study to evaluate the benefits of elastic bandwidth allocation relative to fixed-grid optical networks. Initially, only a working path is provisioned for each connection; however, if a single working path that can support the requested data rate is not available, at most two working paths of lower data rates may be assigned to a connection. Upon occurrence of a link failure, new routes are dynamically computed for all connections affected by the failure. Furthermore, if there are multiple connections with the same source and destination with failed working paths, they are all aggregated into a single connection so as to minimize the amount of resources required for restoration.

Recall that restoration schemes do not reserve backup resources for each connection. A key challenge, therefore, is to provision backup resources to ensure that all connections affected by a single failure can be fully recovered. This problem is referred to as the spare capacity allocation (SCA) problem, and it has been addressed in the context of EONs in [80]. This study considered span restoration, a failure-dependent restoration scheme whereby all connections affected by a span failure are redirected over a different path that connects the two endpoints of the failed span. A MILP formulation of the SCA problem was presented in [80] under the assumption that each node supports *spectrum conversion*; hence, the formulation did not include the continuity constraints that were included in variants of the RSA problem that we discussed earlier. The MILP uses a path-based formulation with the dual objective of minimizing spare capacity and maximizing restorable traffic capacity. The study showed that both working and spare capacity increase linearly with the traffic load, but that spare capacity redundancy improves (i.e., drops) due to more opportunities for spare capacity sharing.

9. Multi-path RSA

Multi-path routing is a technique widely used in SONET networks to improve both the network-wide utilization of link capacity and the survivability of individual connections [81–83]. Multi-path routing introduces flexibility in the design and operation of the network by allowing operators to split traffic demands into multiple streams that are routed independently of each other to the destination. It also affords more options in recovering from network failures. These features can be useful in addressing and overcoming the fragmentation of spectral resources in an elastic optical network.

An algorithm that uses a hybrid single-/multi-path routing (HSMR) scheme to provision a request was presented in [52]. The algorithm considers a set of candidate paths in order of increasing weight, and takes into account fragmentation as follows. For each path, it allocates bandwidth to the request based on the number of contiguous slots (if any) available on the path and the modulation level (which depends on path distance). If the allocated bandwidth is not sufficient for the requested data rate, the algorithm proceeds to allocate bandwidth on the next path, and so on, until a sufficient number of slots has been assigned; the request is blocked if the sum of the available slots on the candidate paths is not sufficient to support its data rate. This technique overcomes the limitations of fragmentation by using multi-path routing whenever the request cannot be provisioned along a single path. Two versions of HSMR were studied, a static one in which paths are computed in advance and an adaptive one that computes paths at the time a request arrives. In the latter case, the link weights were calculated using a metric that captures their fragmentation status. Similar single-/multi-path routing schemes are presented in [84]. The two algorithms in [84] differ in the way they compute paths. The first algorithm computes the paths at the time the request arrives, and in doing so, it takes into account the availability of sufficient bandwidth. The second algorithm uses pre-computed paths, and considers them in the order of increasing delay when assigning bandwidth to a request. A main difference between the algorithms in [84] and the ones in [52] is that the former take the differential delay into account when allocating paths.

The dynamic multi-path service scheme presented in [85] selects the paths on which to provision a request such that the maximum differential delay along these paths does not exceed a given threshold. The algorithm also imposes a minimum bandwidth allocation granularity on each path, so as to avoid splitting the service across a large number of routing paths.

An offline version of the multi-path RSA problem was studied in [86] in the context of supporting high-speed

Ethernet in EONs. The problem of establishing a set of connections, each potentially using more than one path, is formulated as a path-based ILP; in addition to the usual spectrum continuity and contiguity constraints, this ILP also imposes differential delay constraints on the multiple paths for a given connection. An efficient two-step heuristic is also presented. The heuristic first computes a set of fiber-level paths for each source–destination pair. The second phase takes these paths as input, and first attempts to find a single path for each request. If a single path is not found, then the algorithm proceeds to compute a parallel transmission solution across multiple paths.

Multi-path solutions in EONs have also been investigated in the context of protection and restoration. The single-path provisioning, multi-path recovery (SPP-MPR) scheme in [87] provisions a demand on a single primary path. However, it provides for a multi-path squeezed recovery mechanism that utilizes two backup paths and allows for bandwidth squeezing similar to [77,78]. The problem is formulated as an MILP, but a genetic algorithm-based meta-heuristic is also presented.

An offline multi-path provisioning scheme for OFDM-based EONs was first introduced in [88]. The problem, referred to as the “static survivable multi-path routing and spectrum allocation” (SM-RSA), was first formulated as a path-based ILP. A three-step heuristic was also developed. In the first step, Bhandari's algorithm [89] is used to compute the largest number of disjoint paths between each source–destination pair. In the second step, requests are considered either in order of decreasing demands or decreasing path length, and are assigned paths and spectrum blocks. In the third step, requests on paths that use the maximum number of spectral slots are rerouted in an attempt to reduce the use of spectral resources. An online version of the same problem is studied in [90] and again an ILP formulation and a heuristic algorithm are presented. Both the ILP and heuristic limit the number of paths allocated to each connection to at most three; they also allow for connections that require partial (not full) protection.

10. Concluding remarks

The elastic network paradigm is a compelling research area, and the future is likely to see it becoming more essential. We have presented a comprehensive review of recent research in spectrum management techniques for EONs. Our focus has been on techniques to solve the offline and online RSA problem, as well as variants that take into account features including distance-adaptive modulation, fragmentation, traffic grooming, and survivability. An exciting variety of directions in attacking the problem more ambitiously and in more general contexts remain, and we are certain that many significant results are forthcoming.

References

[1] P.J. Winzer, Beyond 100G ethernet, IEEE Commun. Mag. 48 (July (7)) (2010) 26–30.

[2] A. Nag, M. Tornatore, Optical network design with mixed line rates, Opt. Switch. Netw. 6 (3) (2009) 227–237.

[3] ITU-T G.694.1, Spectral Grids for WDM Applications: DWDM Frequency Grid, February 2002.

[4] G. Shen, M. Zukerman, Spectrum-efficient and agile CO-OFDM optical transport networks: architecture, design, and operation, IEEE Commun. Mag. 50 (5) (2012) 82–89.

[5] M. Jinno, T. Ohara, Y. Sone, A. Hirano, O. Ishida, M. Tomizawa, Elastic and adaptive optical networks: possible adoption scenarios and future standardization aspects, IEEE Commun. Mag. 49 (10) (2011) 164–172.

[6] M. Jinno, H. Takara, B. Kozicki, Dynamic optical mesh networks: drivers, challenges and solutions for the future, in: Proceedings of 35th European Conference on Optical Communication (ECOC), September 2009, p. 7.7.4.

[7] 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) (2012) s12–s20.

[8] 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: Proceedings of 34th European Conference on Optical Communication (ECOC), September 2008, p. Th.3.F.6.

[9] G. Zhang, M. De Leenheer, A. Morea, B. Mukherjee, A survey on OFDM-based elastic core optical networking, IEEE Commun. Surv. Tutor. 15 (First Quarter (1)) (2013) 65–87.

[10] W. Shieh, OFDM for flexible high-speed optical networks, J. Lightwave Technol. 29 (May 15 (10)) (2011) 1560–1577.

[11] A.J. Lowery, J. Armstrong, Orthogonal-frequency-division multiplexing for dispersion compensation of long-haul optical systems, Opt. Express 14 (March (6)) (2006) 2079–2084.

[12] A.J. Lowery, L.B. Du, J. Armstrong, Performance of optical OFDM in ultralong-haul WDM lightwave systems, J. Lightwave Technol. 25 (January (1)) (2007) 131–138.

[13] M. Jinno, H. Takara, B. Kozicki, Yukio Tsukishima, Y. Sone, S. Matsuoka, Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies, IEEE Commun. Mag. 47 (11) (2009) 66–73.

[14] H. Beyranvand, J. Salehi, A quality-of-transmission aware dynamic routing and spectrum assignment scheme for future elastic optical networks, J. Lightwave Technol. 31 (September 15 (18)) (2013) 3043–3054.

[15] S. Frisken, G. Baxter, D. Abakoumov, H. Zhou, I. Clarke, S. Poole, Flexible and grid-less wavelength selective switch using LCOS technology, in: Proceedings of OFC/NFOEC, 2011, Paper OTuM3.

[16] 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 (January (1)) (2005) 54–61.

[17] M. Klinkowski, K. Walkowiak, Routing and spectrum assignment in spectrum sliced elastic optical path network, IEEE Commun. Lett. 15 (8) (2011) 884–886.

[18] Y. Wang, X. Cao, Y. Pan, A study of the routing and spectrum allocation in spectrum-sliced elastic optical path networks, in: Proceedings of IEEE INFOCOM, 2011, pp. 1503–1511.

[19] K. Christodoulopoulos, I. Tomkos, E. Varvarigos, Routing and spectrum allocation in OFDM-based optical networks with elastic bandwidth allocation, in: Proceedings of IEEE Globecom, 2010.

[20] G.N. Rouskas, Routing and wavelength assignment in optical WDM networks, in: J. Proakis (Ed.), Wiley Encyclopedia of Telecommunications, John Wiley & Sons, New Jersey, 2001.

[21] S. Shirazipourazad, Ch. Zhou, Z. Derakhshandeh, A. Sen, On routing and spectrum allocation in spectrum-sliced optical networks, in: Proceedings of IEEE INFOCOM, April 2013, pp. 385–389.

[22] S. Talebi, E. Bampis, I. Katib, G.N. Rouskas, The spectrum assignment (SA) problem in optical networks: a multiprocessor scheduling perspective, 2014, submitted for publication.

[23] E. Bampis, M. Caramia, J. Fiala, A. Fishkin, A. Iovanella, Scheduling of independent dedicated multiprocessor tasks, in: Proceedings of the 13th Annual International Symposium on Algorithms and Computation, Lecture Notes in Computer Science, vol. 2518, 2002, pp. 391–402.

[24] J.A. Hoogeveen, S.L. Van de Velde, B. Veltman, Complexity of scheduling multiprocessor tasks with prespecified processor allocations, Discrete Appl. Math. 55 (1994) 259–272.

[25] E. Bampis, A. Kononov, On the approximability of scheduling multiprocessor tasks with time dependent processing and processor requirements, in: Proceedings of the 15th International Parallel and Distributed Processing Symposium, San Francisco, 2001.

[26] Y. Wang, X. Cao, Q. Hu, Routing and spectrum allocation in spectrum-sliced elastic optical path networks, in: Proceedings of IEEE International Conference on Communications (ICC), June 2011.

[27] 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) (2012) 906–917.

[28] A.N. Patel, Ph.N. Ji, J.P. Jue, T. Wang, Routing, wavelength assignment, and spectrum allocation in transparent flexible optical WDM (FWDM) networks, in: Proceedings of Integrated Photonics Research, Silicon and Nanophotonics and Photonics in Switching (IPR/PS), July 2010, p. PDPWG1.

[29] A.N. Patel, Ph.N. Ji, J.P. Jue, T. Wang, Routing, wavelength assignment, and spectrum allocation algorithms in transparent flexible optical WDM networks, Opt. Switch. Netw. 9 (3) (2012) 191–204.

[30] K. Christodoulopoulos, I. Tomkos, E. Varvarigos, Spectrally/bitrate flexible optical network planning, in: Proceedings of 36th European Conference on Optical Communication (ECOC), 2010, p. WE.8.D.3.

[31] L. Velasco, M. Klinkowski, M. Ruiz, J. Comellas, Modeling the routing and spectrum allocation problem for flexgrid optical networks, Photonic Netw. Commun. 24 (2012) 177–186.

[32] X. Wan, L. Wang, N. Hua, H. Zhang, X. Zheng, Dynamic routing and spectrum assignment in flexible optical path networks, in: Proceedings of Optical Fiber Communication Conference and the National Fiber Optic Engineers Conference (OFC/NFOEC), March 2011, p. JWA55.

[33] X. Wang, Q. Zhang, I. Kim, P. Palacharla, M. Sekiya, Blocking performance in dynamic flexible grid optical networks—what is the ideal spectrum granularity?, in: Proceedings of 37th European Conference on Optical Communication (ECOC), September 2011, p. Mo.2.K.6.

[34] Y. Yu, J. Zhang, Y. Zhao, X. Cao, X. Lin, W. Gu, The first single-link exact model for performance analysis of flexible grid WDM networks, in: Proceedings of OFC/NFOEC, 2013, Paper JW2A.

[35] R.C. Almeida Jr., R.A. Delgado, C.J.A. Bastos-Filho, D.A.R. Chaves, H.A. Pereira, J.F. Martins-Filho, An evolutionary spectrum assignment algorithm for elastic optical networks, in: Proceedings of 15th International Conference on Transparent Optical Networks (ICTON), June 2013.

[36] Y. Wang, J. Zhang, Y. Zhao, J. Wang, W. Gu, Routing and spectrum assignment by means of ant colony optimization in flexible bandwidth networks, in: Proceedings of Optical Fiber Communication Conference and the National Fiber Optic Engineers Conference (OFC/NFOEC), March 2012, Paper NTu2J.3.

[37] A.N. Patel, Ph.N. Ji, J.P. Jue, T. Wang, Dynamic routing, wavelength assignment, and spectrum allocation in transparent flexible optical WDM networks, in: Proceedings of SPIE, vol. 7959, January 2011, pp. 79590M-1–79590M-8.

[38] A. Castro, L. Velasco, M. Ruiz, M. Klinkowski, Juan Pedro Fernndez-Palacios, Davide Careglio, Dynamic routing and spectrum (re)allocation in future flexgrid optical networks, Comput. Netw. 56 (12) (2012) 2869–2883.

[39] K. Christodoulopoulos, I. Tomkos, E. Varvarigos, Time-varying spectrum allocation policies and blocking analysis in flexible optical networks, IEEE J. Sel. Areas Commun. 31 (January (1)) (2013) 13–25.

[40] 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) (2010) 138–145.

[41] H. Takara, B. Kozicki, Y. Sone, M. Jinno, Spectrally-efficient elastic optical path networks, in: Proceedings of 15th Optoelectronics and Communications Conference (OECC), July 2010, pp. 116–117.

[42] 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: Proceedings of Optical Fiber Communication and National Fiber Optic Engineers Conference (OFC/NFOEC), March 2010, p. OMU3.

[43] H. Takara, B. Kozicki, Y. Sone, T. Tanaka, A. Watanabe, A. Hirano, K. Yonenaga, M. Jinno, Distance-adaptive super-wavelength routing in elastic optical path network (SLICE) with optical OFDM, in: Proceedings of 36th European Conference and Exhibition on Optical Communication (ECOC), September 2010, p. We.8.D.2.

[44] 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: Proceedings of 37th European Conference on Optical Communication (ECOC), September 2011, p. Mo.2.K.3.

[45] H. Takara, K. Yonenaga, M. Jinno, Spectrally-efficient elastic optical path networks toward 1 Tbps era, in: Optical Fiber Communication Conference, March 2012, p. OTh3B.3.

[46] A. Morea, O. Rival, Advantages of elasticity versus fixed data-rate schemes for restorable optical networks, in: Proceedings of 36th European Conference on Optical Communication (ECOC), September 2010, p. Th.10.F.5.

[47] S. Yang, F. Kuipers, Impairment-aware routing in translucent spectrum-sliced elastic optical path networks, in: Proceedings of 17th European Conference on Networks and Optical Communications (NOC), June 2012.

[48] K. Christodoulopoulos, I. Tomkos, E.A. Varvarigos, Elastic bandwidth allocation in flexible OFDM-based optical networks, J. Lightwave Technol. 29 (9) (2011) 1354–1366.

[49] 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: Proceedings of Optical Fiber Communication Conference and the National Fiber Optic Engineers Conference (OFC/NFOEC), March 2011, p. OTuI7.

[50] R. Wang, B. Mukherjee, Spectrum management in heterogeneous bandwidth optical networks, Opt. Switch. Netw. 11 (2014).

[51] X. Wang, Q. Zhang, I. Kim, P. Palacharla, M. Sekiya, Utilization entropy for assessing resource fragmentation in optical networks, in: Proceedings of OFC/NFOEC, 2012, Paper OTh1A.2.

[52] 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 (January (1)) (2013) 15–22.

[53] 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: Proceedings of Asia Communications and Photonics Conference (ACP), vol. ATh2D.3, November 2012.

[54] K. Christodoulopoulos, I. Tomkos, E. Varvarigos, Dynamic bandwidth allocation in flexible OFDM-based networks, in: Proceedings of Optical Fiber Communication Conference and the National Fiber Optic Engineers Conference (OFC/NFOEC), March 2011, p. OTuI5.

[55] Y. Sone, A. Hirano, A. Kadohata, M. Jinno, O. Ishida, Routing and spectrum assignment algorithm maximizes spectrum utilization in optical networks, in: Proceedings of 37th European Conference on Optical Communication (ECOC), September 2011, p. Mo.1.K.3.

[56] N. Amaya, M. Irfan, G. Zervas, K. Banias, M. Garrich, I. Henning, D. Simeonidou, Y.R. Zhou, A. Lord, K. Smith, V.J.F. Rancano, S. Liu, P. Petropoulos, D.J. Richardson, Gridless optical networking field trial: flexible spectrum switching, defragmentation and transport of 10G/40G/100G/555G over 620-km field fiber, in: Proceedings of 37th European Conference on Optical Communication (ECOC), September 2011, p. Th.13.K.1.

[57] A.N. Patel, P.N. Ji, J.P. Jue, T. Wang, Defragmentation of transparent flexible optical WDM (FWDM) networks, in: Proceedings of Optical Fiber Communication Conference and the National Fiber Optic Engineers Conference (OFC/NFOEC), March 2011, p. OTuI8.

[58] 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: Proceedings of 37th European Conference on Optical Communication (ECOC), September 2011, p. Mo.2.K.4.

[59] R. Wang, B. Mukherjee, Provisioning in elastic optical networks with non-disruptive defragmentation, J. Lightwave Technol. 31 (August (15)) (2013) 2491–2500.

[60] F. Cugini, F. Paolucci, G. Meloni, G. Berrettini, M. Secondini, F. Fresi, N. Sambo, L. Poti, P. Castoldi, Push–pull defragmentation without traffic disruption in flexible grid optical networks, J. Lightwave Technol. 31 (January (1)) (2013) 125–133.

[61] R. Dutta, A.E. Kamal, G.N. Rouskas (Eds.), Traffic Grooming for Optical Networks: Foundations, Techniques and Frontiers, Springer, New York, 2008.

[62] R. Dutta, G.N. Rouskas, Traffic grooming in WDM networks: past and future, IEEE Netw. 16 (November/December (6)) (2002) 46–56.

[63] O. Gerstel, Flexible use of spectrum and photonic grooming, in: Proceedings of Integrated Photonics Research, Silicon and Nanophotonics and Photonics in Switching (IPR/PS), vol. PMD3, July 2010.

[64] B. Kozicki, H. Takara, Y. Tsukishima, T. Yoshimatsu, T. Kobayashi, K. Yonenaga, M. Jinno, Optical path aggregation for 1-Tb/s transmission in spectrum-sliced elastic optical path network, IEEE Photonics Technol. Lett. 22 (17) (2010) 1315–1317.

[65] G. Zhang, M. De Leenheer, B. Mukherjee, Optical traffic grooming in OFDM-based elastic optical networks [Invited], J. Opt. Commun. Netw. 4 (November (11)) (2012) B17–B25.

[66] W. Zheng, Y. Jin, W. Sun, W. Guo, W. Hu, On the spectrum-efficiency of bandwidth-variable optical OFDM transport networks, in: Proceedings of Optical Fiber Communication and National Fiber Optic Engineers Conference (OFC/NFOEC), March 2010, p. OWR5.

[67] Y. Zhang, X. Zheng, Q. Li, N. Hua, Y. Li, H. Zhang, Traffic grooming in spectrum-elastic optical path networks, in: Proceedings of Optical

Fiber Communication Conference and the National Fiber Optic Engineers Conference (OFC/NFOEC), March 2011, p. OTuI1.

[68] S. Zhang, C. Martel, B. Mukherjee, Dynamic traffic grooming in elastic optical networks, IEEE J. Sel. Areas Commun. 31 (January (1)) (2013) 4–12.

[69] M. Liu, M. Tornatore, B. Mukherjee, Survivable traffic grooming in elastic optical networks—shared protection, J. Lightwave Technol. 31 (March (6)) (2013) 903–909.

[70] D. Zhou, S. Subramaniam, Survivability in optical networks, IEEE Netw. 14 (November/December (6)) (2000) 16–23.

[71] B. Chen, J. Zhang, Y. Zhao, Ch. Lv, W. Zhang, Sh. Huang, X. Zhang, W. Gu, Multi-link failure restoration with dynamic load balancing in spectrum-elastic optical path networks, Opt. Fiber Technol. 18 (1) (2012) 21–28.

[72] T. Takagi, H. Hasegawa, K. Sato, T. Tanaka, B. Kozicki, Y. Sone, M. Jinno, Algorithms for maximizing spectrum efficiency in elastic optical path networks that adopt distance adaptive modulation, in: Proceedings of 36th European Conference on Optical Communication (ECOC), September 2010, p. We.8.D.5.

[73] A. Eira, J. Pedro, J. Pires, Optimized design of shared restoration in flexible-grid transparent optical networks, in: Proceedings of OFC/NFOEC, 2012, Paper JTh2A37.

[74] A.N. Patil, P.N. Ji, J.P. Jue, T. Wang, Survivable transparent flexible optical WDM (FWDM) networks, in: Proceedings of OFC/NFOEC, 2011, Paper OTuI2.

[75] X. Shao, Y.-K. Yeo, Z. Xu, X. Cheng, L. Zhou, Shared path protection in OFDM-based optical networks with elastic bandwidth allocation, in: Proceedings of OFC/NFOEC, 2012, Paper OTh4B.4.

[76] J.L. Vizcaino, Y. Ye, V. Lopez, F. Jimenez, F. Musumeci, M. Tornatore, A. Pattavina, P.M. Krummrich, Protection in optical transport networks with fixed and flexible grid: cost and energy efficiency evaluation, Opt. Switch. Netw. 11 (2014).

[77] 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: Proceedings of Optical Fiber Communication (OFC), March 2009, p. OThO2.

[78] Y. Sone, A. Watanabe, W. Imajuku, Y. Tsukishima, B. Kozicki, H. Takara, M. Jinno, Bandwidth squeezed restoration in spectrum-sliced elastic optical path networks (SLICE), J. Opt. Commun. Netw. 3 (3) (2011) 223–233.

[79] B. Chen, J. Zhang, Y. Zhao, Ch. Lv, W. Zhang, Y. Gu, Sh. Huang, W. Gu, A novel recovery algorithm for multi-link failures in spectrum-elastic optical path networks, in: Proceedings of Asia Communications and Photonics Conference (ACP), vol. 8310, November 2011, pp. 83101V-1–83101V-6.

[80] Y. Wei, G. Shen, Sh. You, Span restoration for CO-OFDM-based elastic optical networks under spectrum conversion, in: Proceedings of Asia Communications and Photonics Conference (ACP), November 2012, p. AF3E.7.

[81] D. Cavendish, K. Murakami, S.-H. Yun, O. Matsuda, M. Nishihara, New transport services for next-generation SONET/SDH systems, IEEE Commun. Mag. 40 (May (5)) (2002) 80–87.

[82] E. Hernandez-Valencia, Hybrid transport solutions for TDM/data networking services, IEEE Commun. Mag. 40 (May (5)) (2002) 104–112.

[83] International Telecommunication Union (ITU), Link capacity adjustment scheme (LCAS) for virtually concatenated signals, in: ITU-T G.7042, 2004.

[84] X. Chen, Y. Zhong, A. Jukan, Multi-path routing in elastic optical networks with distance-adaptive modulation formats, in: Proceedings of the IEEE International Conference on Communications (ICC), June 2013.

[85] 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 (January (1)) (2013) 158–160.

[86] X. Chen, A. Jukan, A. Gumaste, Optimized parallel transmission in elastic optical networks to support high-speed Ethernet, J. Lightwave Technol. 32 (January 15 (2)) (2014) 228–238.

[87] A. Castro, L. Velasco, M. Ruiz, J. Comellas, Single-path provisioning with multi-path recovery in flexgrid optical networks, in: Proceedings of the 4th International Workshop on Reliable Networks Design and Modeling (RNDM), 2012.

[88] L. Ruan, N. Xiao, Survivable multipath routing and spectrum allocation in OFDM-based flexible optical networks, J. Opt. Commun. Netw. 5 (March (3)) (2013) 172–182.

[89] R. Bhandari, Survivable Networks: Algorithms for Diverse Routing, Kluwer, USA, 1999.

[90] L. Ruan, Y. Zheng, Dynamic survivable multipath routing and spectrum allocation in OFDM-based flexible optical networks, J. Opt. Commun. Netw. 6 (January (1)) (2014) 77–85.