

Offline Distance-Adaptive Routing and Spectrum Assignment in Mesh Elastic Optical Networks

Sahar Talebi[†], George N. Rouskas^{†*}, Iyad Katib^{*}

[†]Operations Research and Department of Computer Science, North Carolina State University, Raleigh, NC 27695-8206 USA

^{*}King Abdulaziz University, Jeddah, Saudi Arabia

Abstract—The routing and spectrum assignment (RSA) problem has emerged as the key design and control problem in elastic optical networks. Distance adaptive spectrum allocation exploits the tradeoff between spectrum width and reach to improve resource utilization by tailoring the modulation format to the level of impairments along the path. In this paper, we consider the distance-adaptive RSA (DA-RSA) problem with fixed alternate routing. We first show that the DA-RSA problem in networks of general topology is a special case of a well-studied multiprocessor scheduling problem. We then leverage insights from scheduling theory to (1) present new results regarding the complexity of the DA-RSA problem, and (2) build upon list scheduling concepts to develop a computationally efficient solution approach that is effective in utilizing the available spectrum resources.

I. INTRODUCTION

Optical networking technologies underlie the delivery and availability of reliable and survivable Internet services. As optical transmission speeds approach 1 Tbps, new technologies, including flexible spectrum switches and bandwidth variable transceivers [1] continue to drive novel capabilities at the optical layer. Elastic optical networks [2], [3] offer the promise to harness the properties of *individual* optical devices to deliver novel features and capabilities at the *network* scale. Furthermore, using finer spectrum granularity than conventional fixed-grid WDM technology, elastic networks have the potential to accommodate efficiently the ever-increasing traffic demands by tailoring both modulation format and spectrum resources to the data rate and path impairments [2], [4].

Elastic optical network technology is in the early stages of development and/or deployment, yet relevant network design techniques have been the subject of considerable research and development activities in recent years. The routing and spectrum assignment (RSA) problem [5] addresses the network-wide allocation and management of spectral resources and is fundamental to the design and control of elastic optical networks. The objective of RSA is to assign spectrum and a physical path to each demand, so as to optimize spectrum utilization. Several aspects of the problem have been studied in the literature, including offline RSA [6], [7], online RSA [8], [9], distance adaptive RSA (DA-RSA) [10], [11], fragmentation-aware RSA (FA-RSA) [12], RSA and traffic grooming [13], and RSA with restoration [14]; for a recent survey of the literature, we refer the reader to [5]. Most

existing studies approach the problem using classical network design techniques. Network design problems are notoriously hard, and optimal methods (e.g., integer programming formulations) do not scale to topologies encountered in practice. This issue is even more pronounced in elastic optical networks as the network designer has to take into account additional dimensions including variable bandwidth demands (rather than single-wavelength ones as in fixed-grid WDM) and trade-offs in reach versus spectral efficiency.

In this paper, we provide new insight into the structure of the offline distance-adaptive RSA (DA-RSA) problem by relating it to a well-known problem of scheduling multiprocessor tasks on dedicated processors. We also present a computationally efficient solution approach for mesh (i.e., general-topology) networks, based on list scheduling, that is effective in utilizing the available spectrum resources. Specifically, the remainder of the paper is organized as follows. In Section II, we define DA-RSA with fixed alternate routing and show that this problem is a special case of a multiprocessor scheduling problem in which a task may be executed by alternate sets of processors. Accordingly, we leverage scheduling theory to investigate the complexity of the DA-RSA problem (in Section III), and to develop a list scheduling algorithm to solve it (in Section IV). In Section V, we present the results of an experimental study to evaluate the list scheduling algorithm on various network topologies and traffic distributions, and we conclude the paper in Section VI.

II. DA-RSA IN GENERAL GRAPHS AS A SPECIAL CASE OF MULTIPROCESSOR SCHEDULING

The concept of distance-adaptive (DA) spectrum allocation was introduced in [15] to exploit the tradeoff between reach and spectrum width, by tailoring the modulation format to the level of impairments along the path so as to improve spectrum utilization. Specifically, for the same data rate, a high-level modulation format with low SNR tolerance and narrow spectrum may be selected for a short path, whereas a low-level modulation with high SNR tolerance and a wider spectrum may be used for a longer path [11].

We consider the following general definition of the distance-adaptive routing and spectrum assignment (DA-RSA) problem with fixed alternate routing in elastic optical networks.

- *DA-RSA Inputs:* (1) a directed graph $G = (\mathcal{V}, \mathcal{A})$, where \mathcal{V} is the set of nodes and \mathcal{A} is the set of arcs (directed links); (2) k alternate routes, $r_{sd}^1, \dots, r_{sd}^k$, from node s

This work was supported 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.

to node d , where $k \geq 1$ is a small integer, (3) a spectrum demand matrix $T = [t_{sd}^l]$, such that (i) t_{sd}^l is the number of spectrum slots required to carry the traffic from source s to destination d along the l -th route between the two nodes, $l = 1, \dots, k$, and (ii) spectrum demands may increase (but not decrease) with the path length, i.e.,

$$|r_{sd}^l| \leq |r_{sd}^h| \Rightarrow t_{sd}^l \leq t_{sd}^h. \quad (1)$$

- *DA-RSA Objective*: select one of the k possible routes for each spectrum demand and assign spectrum slots along all the arcs of this route such that the total amount of spectrum used on any arc in the network is minimized.
- *DA-RSA Constraints*: (1) spectrum contiguity: each demand is assigned contiguous spectrum slots; (2) spectrum continuity: each demand uses the same spectrum slots along all arcs of its route; and (3) non-overlapping spectrum: demands that share an arc are assigned non-overlapping parts of the available spectrum.

Now, consider the following multiprocessor scheduling problem $P|set_j|C_{max}$, defined as [16]:

- $P|set_j|C_{max}$ *Inputs*: a set of m identical processors, a set of n tasks, a set $set_j = \{S_j^1, \dots, S_j^k\}$ of k alternative processor sets that may execute each task j , where k is an integer, and the processing time p_j^l of task j when it is to be executed on processor set $S_j^l, l = 1, \dots, k$.
- $P|set_j|C_{max}$ *Objective*: assign one of the k sets of processors to execute each task, and schedule the tasks so as to minimize the makespan $C_{max} = \max_j C_j$ of the schedule, where C_j denotes the finish time of task j .
- $P|set_j|C_{max}$ *Constraints*: (1) no preemption is allowed; (2) all the processors in the selected set must work on task j simultaneously, and (3) each processor may execute at most one task at any given time.

The next two lemmas show that the DA-RSA problem with fixed-alternate routing in networks of general topology is a special case of the $P|set_j|C_{max}$ scheduling problem. Lemma 2.1 first shows that DA-RSA transforms to $P|set_j|C_{max}$, and hence, any algorithm for the latter problem also solves the former. Lemma 2.2 shows by counter-example that the reverse result is not true, i.e., that there exist instances of $P|set_j|C_{max}$ for which there is no corresponding instance of DA-RSA.

Lemma 2.1: DA-RSA with fixed-alternate routing in mesh networks transforms to $P|set_j|C_{max}$.

Proof. Consider an instance of the DA-RSA problem with fixed-alternate routing on a general directed graph $G = (\mathcal{V}, \mathcal{A})$, a set of k routes $\{r_{sd}^1, \dots, r_{sd}^k\}$ for each source-destination pair (s, d) , and demand matrix $T = [t_{sd}^l], l = 1, \dots, k$. It is possible to construct an instance of $P|set_j|C_{max}$ such that: (1) there is a processor i for every arc in $a_i \in \mathcal{A}$, (2) there is a task j for each source-destination pair (s, d) , (3) there is a $set_j = \{S_j^1, \dots, S_j^k\}$ for each task j with $S_j^l = \{i : a_i \in \{r_{sd}^l\}\}$ where (s, d) is the source-destination pair corresponding to task j , and (4) the processing time of task j on processor set S_j^l is $p_j^l = t_{sd}^l, l = 1, \dots, k$. In this

transformation, each arc in the DA-RSA problem maps to a processor in the scheduling problem, each spectrum demand to a task, each alternate route of a demand to one of the alternate processor sets of the corresponding task, and the number of spectrum slots along a route of a demand to the processing time of the task on the corresponding set of processors. Note that, because of (1), the processing times of each task j in the $P|set_j|C_{max}$ instance will obey this relationship:

$$|S_j^l| \leq |S_j^h| \Rightarrow p_j^l \leq p_j^h. \quad (2)$$

With this transformation, the spectrum contiguity constraint in allocating slots to a demand implies that processing of the corresponding task in the constructed scheduling problem will continue with no preemption. The spectrum continuity constraint along the arcs of the route taken by a demand guarantees that all the processors within the set assigned to a task will execute this task simultaneously. Also, the non-overlapping spectrum constraint assures that a processor works at most on one task at a time.

Finally, the total amount of required spectrum required for all the demands using an arc of graph G in the DA-RSA problem, is equivalent to the completion time of the last task executed on the corresponding processor. Accordingly, minimizing the spectrum use on any arc of the DA-RSA problem is equivalent to minimizing the makespan of the schedule in the corresponding problem $P|set_j|C_{max}$. ■

Lemma 2.2: There exist instances of $P|set_j|C_{max}$ for which there is no corresponding instance of the DA-RSA problem with fixed-alternate routing.

Proof. By counter-example. Consider an instance of $P|set_j|C_{max}$ with $m = 3$ processors labeled P_1, P_2, P_3 , and $n = 3$ tasks τ_1, τ_2 , and τ_3 . Each task τ_j must be executed by a single set S_j^1 of processors (i.e., $k = 1$), as shown in the following table; the processing time of each task can be arbitrary:

task	S_j^1
τ_1	$\{P_1, P_2\}$
τ_2	$\{P_2, P_3\}$
τ_3	$\{P_1, P_3\}$

The graph of the corresponding DA-RSA instance would have to consist of three directed links, L_1, L_2 , and L_3 , corresponding to processors P_1, P_2, P_3 , respectively. Consider task τ_1 . Since this task must be executed simultaneously on processors P_1 and P_2 , then the corresponding demand in the DA-RSA instance must be routed along the two arcs L_1 and L_2 . Suppose that the path of this demand is the directed pair of arcs $\langle L_1, L_2 \rangle$; if the path is the directed pair of arcs $\langle L_2, L_1 \rangle$, then similar arguments may be used to reach the same conclusion. Now consider task τ_2 . Similarly, the path of the corresponding demand can be either (a) $\langle L_2, L_3 \rangle$, or (b) $\langle L_3, L_2 \rangle$. In case (a), the graph of the DA-RSA instance would have to be the directed three-link chain network $\langle L_1, L_2, L_3 \rangle$. In case (b), the graph of the

DA-RSA instance would have to be a a three-link network in which links L_1 and L_3 feed into link L_2 . Consequently, there is no feasible path for the spectrum demand corresponding to the third task τ_3 in either graph, as it is not possible for traffic on link L_1 to continue onto link L_3 , or vice versa. Therefore, an instance of DA-RSA does not exist. ■

III. COMPLEXITY RESULTS

The problem $P2|set_j|C_{max}$ in which the number of processors is fixed to $m = 2$, is NP-hard [17]. Moreover, it has been shown that, unless $P = NP$, no constant-ratio polynomial time approximation algorithm exists for the general problem $P|set_j|C_{max}$ [18]. However, since the DA-RSA problem is a special case of $P|set_j|C_{max}$, it is possible that polynomial or approximation algorithms exist for special topologies or spectrum demand matrices. In this section, we present theoretical results on the complexity of the DA-RSA problem.

Before we proceed, we introduce two definitions. First, we let K_N denote a *complete digraph* with N nodes. Since every pair of distinct nodes in K_N is connected by a pair of distinct arcs, one in each direction, the total number of arcs in the graph is equal to $N(N - 1)$. Second, in the context of multiprocessor scheduling, we refer to tasks as *compatible* if they can be executed simultaneously, i.e., if the processor sets assigned to the tasks are pairwise disjoint. We now have the following lemma.

Lemma 3.1: The DA-RSA problem on complete digraphs K_N , $N \geq 2$, is solvable in polynomial time.

Proof. From Lemma 2.1 the multiprocessor scheduling problem instance corresponding to a DA-RSA instance on K_N contains $N(N - 1)$ processors, one for each arc of K_N . Let us select the shortest path (i.e., direct arc) for each spectrum demand in the DA-RSA instance. Then, each task in the scheduling instance is to be executed on its own distinct processor. Therefore, the problem reduces to that of scheduling a set of single-processor tasks that are pairwise compatible. Since all tasks may be executed in parallel, the makespan of the schedule is equal to the processing time of the longest task. Recall that all instances of $P|set_j|C_{max}$ constructed from an instance of the DA-RSA problem are such that processing times of tasks satisfy expression (2). Therefore, this makespan is optimal. ■

Although DA-RSA may be solved optimally on a complete digraph using shortest path routing, as the above lemma implies, the next three results show that DA-RSA on general topologies derived by deleting arcs from a complete digraph, is NP-Complete.

Theorem 3.1: DA-RSA on a digraph G obtained by deleting the two arcs¹ between any pair of nodes of K_4 , is NP-Complete.

¹In typical telecommunication networks, two nodes are directly connected using two links, one in each direction. Hence, in this and the next theorem, we only consider the case of removing both arcs between a pair of nodes. Although, it is possible to extend the results to the case of deleting one arc at a time, we will consider this task in future work.

The proof is omitted due to page constraints, but it is available in the first author's dissertation [19]. Theorem 3.1 shows that removing the two arcs between any pair of nodes of K_4 renders the DA-RSA problem on the resulting graph K'_4 NP-Complete. The following theorem shows that removing two pairs of arcs from K_5 yields a problem that is also NP-Complete.

Theorem 3.2: DA-RSA on a digraph K'_5 obtained by deleting the two arcs between any two pairs of nodes of K_5 , is NP-Complete.

Again, the proof is omitted but can be found in [19].

We now provide the following complexity result for the DA-RSA problem on general graphs.

Lemma 3.2: Let G be a digraph. If either K'_4 or K'_5 is a vertex-induced subgraph of G , then the DA-RSA problem on G is NP-Complete.

Proof. Let K'_4 be a vertex-induced subgraph of G ; identical arguments apply when K'_5 is a vertex-induced subgraph of G . Consider an instance of DA-RSA on G with the following spectrum demands: (1) arbitrary, between nodes of the K'_4 subgraph, (2) equal to a large number M , between adjacent nodes not in the K'_4 subgraph, and (3) equal to zero, between non-adjacent nodes not in the K'_4 subgraph. Similar to the observations in the previous theorem, in the optimal solution, each arc of G that is not part of the subgraph K'_4 only carries the traffic between the directly connected nodes. Hence, this instance reduces to a DA-RSA sub-problem on digraph K'_4 , which, according to Theorem 3.1, is NP-Complete. ■

IV. A LIST SCHEDULING ALGORITHM FOR $Pm|set_j|C_{max}$

In this section, we propose a list scheduling (LS) algorithm for the $Pm|set_j|C_{max}$ problem. Since DA-RSA is a special case of $Pm|set_j|C_{max}$, this algorithm can be used to solve the DA-RSA problem in networks of general topology. This is accomplished in three steps: (1) the DA-RSA instance at hand is first transformed to an instance of $Pm|set_j|C_{max}$ following the process described in Lemma 2.1, (2) the LS algorithm is applied to construct a schedule that solves the scheduling instance, and (3) the schedule is transformed back to a solution of the DA-RSA instance.

The input to the LS algorithm is a list of tasks L , along with their corresponding k alternate sets of processors. Tasks in the list are sorted in decreasing order of the processing time on their smallest processor set; ties are broken by the size (i.e., the number of processors) of their smallest processor set, and further ties are broken arbitrarily. For each task, its alternate processor sets are sorted in increasing order of their size.

At each scheduling instant t , the algorithm scans the list L to find the first task j and processor set S_j^t that is compatible with the tasks already executing at this time t . This set S_j^t of processors is selected to execute task j starting at time t , and the algorithm removes the task from L . The algorithm updates the set of free processors at time t , and continues scanning list L , repeating the above process until no other compatible task is found. Then, the algorithm advances t to the earliest time

List Scheduling Algorithm for $Pm|set_j|C_{max}$

Input: A list L of n tasks on m processors, each task j defined by the set $set_j = \{S_j^1, \dots, S_j^k\}$ of k alternative processor sets on which it may be executed, and the corresponding processing times $\{p_j^1, \dots, p_j^k\}$

Output: A schedule of tasks, i.e., the time T_j when task j starts execution, along with the set S_j of processors assigned to it and the corresponding processing time p_j

begin

1. $t \leftarrow 0$ //Scheduling instant
2. $F \leftarrow \{1, \dots, m\}$ //Set of currently idle processors
3. **while** list $L \neq \emptyset$ **do**
4. $j \leftarrow$ first task in list L
5. $S_j \leftarrow \emptyset$ //Set of processors to execute task j
6. $p_j \leftarrow 0$ //Processing time of task j
7. **for** $z \leftarrow 1$ to k
8. **if** $S_j^z \subseteq F$ **then**
9. $S_j \leftarrow S_j^z$
10. $p_j \leftarrow p_j^z$
11. $T_j \leftarrow t$
12. $F \leftarrow F \setminus S_j$
13. Remove the task j from list L
14. **break**
15. **while** not at the end of list L **or** $F \neq \emptyset$ **do**
16. $i \leftarrow$ first task in list
17. **for** $w \leftarrow 1$ to k
18. **if** $S_i^w \subseteq F$ **then**
19. $S_i \leftarrow S_i^w$
20. $p_i \leftarrow p_i^w$
21. $T_i \leftarrow t$
22. $F \leftarrow F \setminus S_i$
23. Remove the task i from list L
24. **break**
25. **end while** // no more tasks may start at time t
26. $j \leftarrow$ the first task executing at time t to complete
27. $t \leftarrow T_j + p_j$
28. $F \leftarrow F \cup S_j$
19. **end while**

end

Fig. 1. A list scheduling (LS) algorithm to select one set S_j and its corresponding processing time p_j to execute each task j of the $Pm|set_j|C_{max}$ problem.

$t' > t$ at which one of the currently executing tasks will be completed, releases the set of processors assigned to the just completed task, and repeats the above actions for time t' . The algorithm continues in this manner until all tasks in list L have been scheduled.

A pseudocode description of the LS algorithm is provided in Figure 1. Both the outer and inner **while** loops of the algorithm take at most $O(n)$ time, in the worst case, where n is the number of tasks in the scheduling problem. Both **for** loops take time $O(k)$ in the worst case, where k is the number of alternate processor sets. Therefore, the running time complexity of the LS algorithm is $O(kn^2)$. Since the number of tasks corresponds to the number of spectrum demands, the complexity of the algorithm when applied to the DA-RSA problem is $O(kN^4)$, where N is the number of nodes and k the number of alternate paths.




Fig. 2. The NSFNet topology

V. NUMERICAL RESULTS

We have evaluated the performance of the LS algorithm by carrying out simulation experiments with a large number of DA-RSA problem instances. Each problem instance is characterized by three parameters: (1) the network topology, (2) the number k of shortest paths for each source-destination pair, and (3) a randomly generated spectrum demand matrix.

Due to space constraints, in this paper we only show results for the 14-node, 42-arc (directed link) NSFNet shown in Figure 2; for results on other general topology networks of varying size and average nodal degree, the reader is referred to [19].

We used Yen's algorithm [20] to compute the k loop-less shortest paths, $k = 1, \dots, 7$, between each pair of nodes in each topology. Yen's algorithm takes time $O(N^3)$, where N is the number of nodes. For the experiments we present in this section, we assumed that all links have unit weight for purposes of computing shortest paths.

A. Spectrum Demand Matrix

For each DA-RSA problem instance we randomly generate a spectrum demand matrix in two steps: traffic demand generation and distance-adaptive spectrum allocation.

1) *Traffic Demand Generation:* We assume that the elastic optical network supports the following data rates (in Gbps): 10, 40, 100, 400, and 1000. Therefore, in the first step, traffic rates between every pair of nodes are drawn from one of three probability distributions:

- *Distance-independent:* each value in the set $\{10, 40, 100, 400, 1000\}$ is selected with equal probability.
- *Distance-increasing:* the probability assigned to each value in the set $\{10, 40, 100, 400, 1000\}$ depends on the length of the shortest path between the source and destination nodes, such that the probability of higher values in the set *increases* with the length of the shortest path.
- *Distance-decreasing:* the probability assigned to higher values in the set $\{10, 40, 100, 400, 1000\}$ *decreases* with the length of the shortest path between the source and destination nodes.

2) *Distance-Adaptive Spectrum Allocation:* In the second step, we determine the number t_{sd}^l of spectrum slots required for the traffic demand to be carried on the l -th alternate path,

$l = 1, \dots, k$, from source s to destination d . In distance-adaptive spectrum allocation, the number of slots depends on both the data rate and the length of the path [2], [15]. We adopt the parameters of the study in [15], and assume a slot width of 12.5 GHz and three modulation formats:

- *Paths with up to 4 links*: the 64-QAM modulation format is used such that data rates of 10, 40, 100, 400, and 1000 Gbps require 1, 1, 2, 6, and 14 spectrum slots, respectively.
- *Paths with 5-9 links*: the 16-QAM modulation format applies, such that rates of 10, 40, 100, 400, and 1000 Gbps are assigned 1, 1, 2, 8, and 20 slots, respectively.
- *Paths with 10 or more links*: the QPSK modulation format is utilized, and data rates of 10, 40, 100, 400, and 1000 Gbps are allocated 1, 2, 4, 16, and 40 spectrum slots, respectively.

B. Evaluation Metrics

The first metric we consider is the maximum number of spectrum slots on any link in the network required by the solution to a DA-RSA problem instance obtained by the LS algorithm. We denote this value as $MaxSlots_{LS}$; as the reader may recall, this value is equivalent to the length of the schedule constructed by the LS algorithm for the corresponding scheduling problem instance. This metric can provide insight into the impact of the number k of alternate paths or the traffic rate distribution on the use of spectrum resources in the network.

In order to evaluate the quality of the LS algorithm, and since the optimal solution cannot be obtained in polynomial time, it is important to compute a lower bound (LB). Let D_q^{in} and D_q^{out} denote the in- and out-degrees of node q . A simple lower bound for the DA-RSA problem can be calculated as follows:

$$LB = \max\left\{\max_s \sum_d t_{sd}/D_s^{out}, \max_d \sum_s t_{sd}/D_d^{in}\right\} \quad (3)$$

where t_{sd} in the above expression is the spectrum demand for the traffic from s to d along the shortest path between the two nodes. The metric we use to characterize of the LS algorithm is the ratio

$$R = MaxSlots_{LS}/LB. \quad (4)$$

Clearly, $R \geq 1.0$; the closer R is to 1.0, the better the performance of the algorithm. We note, however, that the lower bound in (3) only considers spectrum demands in and out of each node, and does not account for the interaction of these demands along the links of the network; therefore, we expect the bound to be loose.

The figures we present in the next section report average values for either $MaxSlots_{LS}$ or R . Specifically, each data point on these figures is the average of 10 replications of a random experiment; in turn, each replication is the average of 30 random instances generated for the stated parameters (i.e., topology, number k of paths, and traffic rate distribution). The figures also report 95% confidence intervals which can be seen to be narrow.




Fig. 3. Spectrum slots vs. number k of paths in NSF

C. Results and Discussion

Figure 3 plots the maximum number of spectrum slots, $MaxSlots_{LS}$, as a function of the number k of alternate paths, for the NSFNet topology. The figure includes three curves, each representing results for problem instances with spectrum demand matrices generated by the distance-independent, distance-increasing, and distance-decreasing distributions, respectively.

We first observe that the amount of spectrum increases with the size of the network, reflecting the corresponding increase in traffic demands due to the larger number of source-destination pairs. Nevertheless, the overall behavior of the curves is consistent across the three traffic distributions. Specifically, the amount of spectrum resources is high for shortest path routing ($k = 1$), but drops sharply (between 20-50%, depending on the distribution) when demands may be routed along one of $k = 2$ alternate paths. As the number k of alternate paths increases further, the number of spectrum slots decreases more slowly and eventually levels off, indicating the diminishing returns of employing each additional path. Very similar results have been observed for other networks of general topology, and are reported in [19].

A final observation from the figure is that the solution to the DA-RSA problem is highly sensitive to the traffic demand distribution. Specifically, everything else being equal, the distance-increasing distribution requires more spectrum than the distance-independent distribution, which in turn is more resource-intensive than the distance-decreasing distribution. This result can be explained by the fact that demands between nodes that are far away from each other consume more spectral resources in the network than the same demands between two nearby nodes due to (1) the larger number of links in the paths they travel, and (2) the wider spectrum that is required to carry the demand if the length of its path crosses the threshold into a lower-level modulation with high SNR tolerance.

Let us now turn our attention to Figure 4 which plots the average ratio R in expression (4) against the number k of paths for NSFNet topology; again, the figure includes three plots, one per demand distribution. Note that the lower bound in (3)




Fig. 4. Average ratio vs. number k of paths in NSF

is independent of the number k of alternate paths for each demand. Since the number of required slots, $MaxSlots_{LS}$ decreases with k , as seen in the previous figure, we expect R to decrease as well, and this is exactly what we observe in Figures 4.

Nevertheless, there is an important difference between the figure that plots the absolute value of spectrum slots required and the one that shows the average ratio. Specifically, we observe that there are significant gaps between the various curves in Figure 3, which, as we explained above, are due to the combined effects of the demand distribution and distance-adaptive spectrum allocation. On the other hand, the curves of the various distributions in Figure 4 are closer to each other and the average ratios of the three distributions converge to similar values. Recall that the lower bound in (3) depends on the demands in and out of each node in the network, and hence it depends on the traffic distribution. Therefore, the behavior of the curves in Figure 4 is a strong indication that, for the distributions we considered in this study, the LS algorithm is capable of exploiting alternate paths to construct solutions that move towards the lower bound, regardless of the absolute value of spectrum slots required in each problem instance.

Overall, the results in this section (and in [19]) indicate that the LS algorithm is effective in using a small number of alternate paths (i.e., $k = 5, 6$) to utilize spectrum resources efficiently, by balancing the traffic demands across the network links.

VI. CONCLUDING REMARKS

We have shown that the distance-adaptive routing and spectrum assignment (DA-RSA) problem with fixed alternate routing in mesh networks transforms to a well-known processor scheduling problem. We have also developed a computationally efficient algorithm that builds upon list scheduling concepts to jointly tackle the routing and spectrum assignment aspects of DA-RSA. Our work explores the tradeoffs involved in DA-RSA algorithm design, and opens up new research directions in leveraging the vast literature in scheduling theory to address important and practical problems in network design.

REFERENCES

- [1] G. Zhang, M. De Leenheer, A. Morea, and B. Mukherjee. A survey on OFDM-based elastic core optical networking. *IEEE Communications Surveys & Tutorials*, 15(1):65–87, First Quarter 2013.
- [2] O. Gerstel, M. Jinno, A. Lord, and S. J B Yoo. Elastic optical networking: a new dawn for the optical layer? *IEEE Communications Magazine*, 50(2):s12–s20, 2012.
- [3] M. Jinno, H. Takara, and B. Kozicki. Dynamic optical mesh networks: Drivers, challenges and solutions for the future. In *Proceedings of 35th European Conference on Optical Communication (ECOC)*, page 7.7.4, September 2009.
- [4] G. Shen and M. Zukerman. Spectrum-efficient and agile CO-OFDM optical transport networks: architecture, design, and operation. *IEEE Communications Magazine*, 50(5):82–89, 2012.
- [5] S. Talebi, F. Alam, I. Katib, M. Khamis, R. Khalifah, and G. N. Rouskas. Spectrum management techniques for elastic optical networks: A survey. *Optical Switching and Networking*, 13:34–48, July 2014.
- [6] Y. Wang, X. Cao, and Y. Pan. A study of the routing and spectrum allocation in spectrum-sliced elastic optical path networks. In *Proceedings of IEEE INFOCOM*, pages 1503–1511, 2011.
- [7] A. N. Patel, Ph. N. Ji, J. P. Jue, and T. Wang. Routing, wavelength assignment, and spectrum allocation algorithms in transparent flexible optical WDM networks. *Optical Switching and Networking*, 9(3):191–204, 2012.
- [8] X. Wan, L. Wang, N. Hua, H. Zhang, and 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)*, page JWA55, March 2011.
- [9] X. Wang, Q. Zhang, I. Kim, P. Palacharla, and 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)*, page Mo.2.K.6, September 2011.
- [10] B. Kozicki, H. Takara, Y. Sone, A. Watanabe, and 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)*, page OMU3, March 2010.
- [11] S. Yang and 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.
- [12] K. Wen, Y. Yin, D. J. Geisler, Sh. Chang, and 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)*, page Mo.2.K.4, September 2011.
- [13] Y. Zhang, X. Zheng, Q. Li, N. Hua, Y. Li, and 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)*, page OTu11, March 2011.
- [14] A. Eira, J. Pedro, and J. Pires. Optimized design of shared restoration in flexible-grid transparent optical networks. In *Proceedings of OFC/NFOEC 2012*, page Paper JTh2A37, 2012.
- [15] M. Jinno, B. Kozicki, H. Takara, A. Watanabe, Y. Sone, T. Tanaka, and A. Hirano. Distance-adaptive spectrum resource allocation in spectrum-sliced elastic optical path network. *IEEE Communications Magazine*, 48(8):138–145, 2010.
- [16] L. Bianco, J. Blazewicz, P. Dell’Olmo, and M. Drozdowski. Scheduling multiprocessor tasks on a dynamic configuration of dedicated processors. *Annals of Operations Research*, 58(7):493–517, 1995.
- [17] M. Kubal. The complexity of scheduling independent two-processor tasks on dedicated processors. *Information Processing Letters*, 24(3):141–147, 1987.
- [18] L. Torres A. Miranda and J. Chen. On the approximability of multiprocessor task scheduling problems. In *Proceedings of the 13th Annual International Symposium on Algorithms and Computation*, volume LNCS 2518, pages 403–415, 2002.
- [19] Sahar Talebi. *On Routing and Spectrum Assignment in Elastic Optical Networks*. PhD thesis, North Carolina State University, Raleigh, NC, December 2015.
- [20] J. Y. Yen. Finding the k shortest loopless paths in a network. *Management Science*, 17(11):712–716, 1971.