

First-Fit: A Universal Algorithm for Spectrum Assignment

George N. Rouskas
North Carolina State University

Abstract—First-Fit (FF) is a well-known and widely deployed algorithm for spectrum assignment (SA), but until our recent study [1], investigations of the algorithm had been experimental in nature and no formal properties of the algorithm with respect to SA were known. In this work, we show that FF is a *universal* algorithm for the SA problem in the sense that 1) it can be used to construct solutions equivalent to, or better than, any solution obtained by any other algorithm, and 2) it can construct an optimal solution. This universality property applies to both the min-max and min-frag objectives, and to variants of the SA problem with or without guard band constraints. Consequently, the spectrum symmetry-free model of [1] extends to all known SA variants, which therefore reduce to permutation problems. Accordingly, all variants may be solved by similar, intuitive, effective and highly parallelizable algorithms. Our results unlock new algorithmic approaches for optical network design problems that encompass SA as an integral subproblem.

I. INTRODUCTION

Spectrum allocation (SA) underlies much of optical network design [2] and has been studied extensively in a wide range of optimization problems, usually coupled to objectives including routing [3]–[5], virtual topology design [6], [7], traffic grooming [8], and network survivability [9]. The SA problem is intractable even when considered in isolation, i.e., separately from other aspects of network design [10]. Consequently, a number of heuristic algorithms, including first-fit, best-fit, most-used, and least loaded [11] have been developed and studied experimentally.

The first-fit (FF) algorithm is a simple heuristic for the SA problem that operates without global knowledge and has been shown to be effective across a wide range of network topologies and traffic demands [11]–[13]. Accordingly, it is commonly employed in practice. Recently, we proved an optimality property of FF that forms the basis for a symmetry-free model for spectrum assignment in networks of general topology. This model eliminates from consideration symmetric solutions, i.e., equivalent solutions derived from spectrum slot permutations, and opens up new algorithmic directions for the SA problem specifically, and optical network design, more generally.

Before its application to spectrum assignment, the FF algorithm had been investigated theoretically since the early days of computing in the context of bin packing [14] and memory

This work was supported by the National Science Foundation under Grant CNS-1907142.

allocation [15], among other fields. For instance, FF has been shown to be an approximation algorithm for the bin packing problem and a series of studies over four decades gradually improved the approximation ratio [14], [16], [17]. To the best of our knowledge, however, before we derived the FF optimality property in [1], investigations of the FF algorithm within the context of spectrum assignment, on its own or as part of optical network optimization problems, were purely experimental and no formal properties of the algorithm were known.

With this work we aim to close this gap. Specifically, we carry out a theoretical investigation of FF and show that it is a *universal* algorithm for all known variants of the SA problem. In this paper, we use the term *universal* to indicate that, for any instance of any SA problem variant, the pure FF algorithm or an appropriately modified version is capable of 1) constructing a solution that is equivalent to, or better than, any solution obtained by any other algorithm, and 2) constructing an optimal solution.

The paper is organized as follows. In Section II we review our recent results regarding the min-max SA problem, the variant that has been most extensively studied in the literature. We discuss the optimality property of FF with respect to min-max SA that allowed us to develop the first symmetry-free model for networks of general topology; in turn, this model led to intuitive and effective algorithms. In the following three sections, we extend these results to several SA variants, and specifically, min-max DSA (Section III), min-frag SA (Section IV), and min-frag DSA (Section V). We conclude the paper in Section VI.

II. MIN-MAX SA AS A PERMUTATION PROBLEM

Let $G = (V, A)$ represent the topology graph of an optical network with a set V of nodes and a set A of directed fiber links, where $N = |V|$ and $L = |A|$ denote the number of nodes and links, respectively. The traffic offered to the network consists of C connections $T_i, i = 1, \dots, C$, where each connection is a tuple $T_i = (s_i, d_i, p_i, t_i)$ such that: s_i and d_i are the endpoints of the connection, p_i is the path between nodes s_i and d_i that the connection must follow, and t_i is the number of spectrum slots that are required to carry the traffic of the connection along path p_i .

We assume that the spectrum slots on each link are indexed $1, 2, 3, \dots$, and we define the min-max spectrum assignment (SA) problem as:

Definition 2.1 (min-max SA):

Input: Graph G and C connections $\{T_i = (s_i, d_i, p_i, t_i)\}$.

Output: An assignment of t_i spectrum slots to each connection T_i along its path p_i .

Objective: Minimize the index of the highest spectrum slot used on any link in the network.

Constraints:

- 1) *Contiguity:* each connection T_i is assigned a block of t_i contiguous spectrum slots starting at index f_i , i.e., block $[f_i, f_i + 1, \dots, f_i + t_i - 1]$.
- 2) *Continuity:* each connection is assigned the same block of slots on each link $\ell \in p_i$ along its path p_i ; and
- 3) *Nonoverlap:* connections whose paths share a link are assigned non-overlapping blocks of spectrum slots.

The first-fit (FF) algorithm [11] considers connections in a fixed order and assigns to each one a contiguous block of spectrum slots that starts at the lowest-indexed slot available along the links of the connection's path. Let P denote a permutation of the C connections, and let $FF(P)$ denote the solution constructed by applying the FF algorithm to the connections in the order implied by P . In recent work [1] we have shown that there exists a permutation P^* of the C connections such that $FF(P^*)$ is an optimal solution to the min-max SA problem. This result has several implications:

- 1) **min-max SA as a permutation problem:** to find an optimal spectrum assignment it is sufficient to examine the connection permutations, hence min-max SA is transformed into a permutation problem; accordingly, we have developed an optimal, recursive, branch-and-bound algorithm, recursive first-fit (RFF), to search the permutation space efficiently [1], [18].
- 2) **Symmetry-free spectrum assignment:** in selecting among the various connection permutations there is no need to consider a spectrum assignment other than the one produced by the FF algorithm; hence, symmetric solutions are automatically eliminated from consideration and the size of the solution space is reduced by orders of magnitude [1].
- 3) **Inherent parallelism:** the connection permutation space is naturally represented as a tree which may be decomposed into non-overlapping subtrees [1], [18]; accordingly, multi-threaded implementations of RFF may explore the subtrees in parallel.
- 4) **FF subsumes any other min-max SA algorithm:** the results of [1] imply that for any solution S constructed by *any other* SA algorithm (e.g., best-fit, most-used, etc. [11]), there exists a connection permutation P such that $FF(P)$ has an objective value equal to or better than that of S .

In the following, we extend these results to other SA problem variants.

III. THE MIN-MAX DSA PROBLEM

The min-max distance SA (min-max DSA) problem is a variant introduced in [19] that includes the min-max SA

problem of the previous section as a special case. Specifically, min-max DSA arises when it is required to allow for a guard band between blocks of spectrum slots assigned to different connections on the same link, e.g., to prevent crosstalk or reduce security threats at the optical layer. Min-max DSA constitutes a fairly general model in that the size of the guard band is not fixed but depends on each connection pair, e.g., on the number of spectrum slots, path, or other attribute of the two adjacent connections [19].

Formally, the min-max DSA problem can be defined as follows:

Definition 3.1 (min-max DSA):

The min-max SA problem of Definition 2.1 with:

Additional Input: Guard band size $g_{ij} \geq 0$ for each connection pair (T_i, T_j) , $i \neq j$.

Additional Constraint:

- 4) *Guard band:* if connections T_i and T_j are assigned adjacent slots on a link ℓ shared by their paths, their spectrum blocks must be separated by a block of unallocated (empty) slots of size at least g_{ij} .

Clearly, if $g_{ij} = 0 \forall i, j$, min-max DSA reduces to min-max SA.

A. The DSA-FF Algorithm and Its Universality Property

Consider a permutation P of the C connections. Without loss of generality, assume that the connections are relabeled so that $P = \langle T_1, T_2, \dots, T_C \rangle$. Let DSA-FF be a modified version of the FF algorithm for the min-max DSA problem. Similar to FF, DSA-FF takes as input some permutation P and assigns spectrum to each connection at a time in the order implied by P , i.e., T_1, T_2, \dots, T_C . After considering connections T_1, \dots, T_{i-1} , DSA-FF assigns to connection T_i a contiguous spectrum block of t_i slots starting at slot f_i such that f_i is the slot with the lowest index satisfying two conditions:

- 1) slots $[f_i, f_i + 1, \dots, f_i + t_i - 1]$ are free along all links of the connection's path p_i , and
- 2) for each link $\ell \in p_i$, if some connection T_j , $j < i$, has been assigned a block of slots immediately below slot f_i on link ℓ , then there is a guard band of size at least g_{ij} between the two blocks; i.e., $f_i \geq (f_j + t_j - 1) + g_{ij}$.

We now prove the following DSA-FF universality property which implies that, for any feasible solution to the min-max DSA problem produced by any algorithm, the DSA-FF algorithm may construct an equivalent or better solution, i.e., one with an equal or lower objective value.

Lemma 3.1 (Universality Property for min-max DSA): Let S be any feasible solution to the min-max DSA problem. There exists a permutation P of the C connections such that applying the DSA-FF algorithm in the order implied by P yields S or another feasible solution S' with equal or lower objective value.

Proof. By construction.

Consider a feasible solution S to the min-max DSA problem with objective value equal to SOL , and label the slots on

each link as $1, 2, \dots, SOL$. By definition, feasible solution S satisfies all four constraints of the min-max DSA problem. Let f_i denote the slot with the lowest index within the block of t_i slots allocated to each request $T_i, i = 1, \dots, C$, under solution S .

Let P_S be the complete permutation in which the requests T_i are listed in increasing order of f_i in solution S , with ties broken arbitrarily. Consider the block of t_j contiguous spectrum slots, starting at slot f_j , allocated to some connection T_j . Let us remove this block of t_j slots from solution S . In the remaining partial solution, it is possible that there exists a block of t_j slots starting at a lower indexed slot $f'_j < f_j$ such that 1) they are available on all links of path p_j , and 2) there are free slots above and below this block such that allocating this block to T_j will not violate the guard band constraints. If so, we can allocate the lower-indexed t_j slots starting with slot f'_j to connection T_j without 1) affecting the feasibility of the solution, or 2) increasing the objective value beyond SOL .

Based on this observation, we modify solution S by considering the connections one by one, in increasing order of f_i as listed in permutation P_S . For each connection T_i , we remove its block of spectrum slots that starts at slot f_i from the solution, and we allocate to it an equal block of slots starting at the lowest possible slot index f'_i in the partial solution such that none of the problem constraints are violated, keeping in mind that f'_i may be equal to f_i . This modified solution S' does not use more than SOL slots on any link, since any modifications involve the allocation of lower-indexed spectrum slots to connections. At the same time, since modifications are applied to the starting solution S only if no constraints are violated, the modified solution S' is also feasible. Hence, the modified solution S' is either 1) identical to S , if it was not possible to move any spectrum blocks, or 2) a different feasible solution with an objective function equal to or lower than SOL , if the spectrum slots allocated to some connection(s) in S were moved to a lower-indexed block. Importantly, by construction the modified solution S' is such that no connection may be allocated to a spectrum block that starts at a lower-indexed slot.

Let P be the permutation in which the connections are reordered so that they are listed in increasing order of f'_i in the modified solution S' , and let us apply the DSA-FF algorithm to this permutation. The algorithm allocates to each connection T_i a block of t_i contiguous slots starting at the lowest-indexed slot for which all problem constraints are satisfied. Therefore, the DSA-FF algorithm will construct the modified solution S' above that is feasible and has an objective value no larger than SOL . ■

Figure 1 illustrates the proof of Lemma 3.1 using a simple example. Figure 1(a) shows a feasible solution to the min-max DSA problem on a four-link chain network and $C = 5$ connections labeled A, B, C, D, E . Each connection is represented by a rectangle of a different color. The length of each rectangle spans all the links in the corresponding connection's path,

whereas the width represents the number of slots allocated to the connection. For instance, the bottommost connection A in the figure spans all four links of the network and has been allocated the contiguous block consisting of slots 1 and 2 along these links. For the problem instance shown in the figure, we have $g_{AB} = 3, g_{AC} = g_{AD} = 2$ and $g_{AE} = g_{BE} = 1$. Therefore, the solution in Figure 1 is feasible as there is a guard band (i.e., block of empty slots) of appropriate size between the blocks of adjacent connections.

Figure 1(b) is the modified solution constructed by the proof of Lemma 3.1. In this solution, connection E has been moved from slot 11 to slot 4, as the three free slots between the blocks of connections A and B in Figure 1 are sufficient to allocate one slot to connection E and form two guard bands of size 1 between A and E and between E and B . Similarly, connection D has been assigned slots 5 and 6, instead of slots 6 and 7 in the original solution. The new solution is also feasible and has an objective value of 9, lower than the objective value 11 of the original solution. The reader may also verify that 1) the solution of Figure 1(b) is the one constructed by the DSA-FF algorithm on permutation $\langle A, B, C, D, E \rangle$, and 2) the solution of Figure 1(a) would not have been produced by DSA-FF on any permutation of the five connections. We also note that the solution in Figure 1(b) is optimal since the objective value is equal to the number of slots required to accommodate the amount of traffic and corresponding guard bands associated with the connections whose path includes the bottleneck link 4.

B. Symmetry-Free Model for min-max DSA

Lemma 3.1 applies to all feasible solutions, including optimal ones, yielding this corollary:

Corollary 3.1: There exists a permutation P^* of the C connections such that $DSA\text{-FF}(P^*)$ is an optimal solution to the min-max DSA problem.

Consequently, the symmetry-free model we introduced in [1] for the min-max SA problem extends to this general min-max DSA variant. Specifically, min-max DSA reduces to a permutation problem in that, to find an optimal solution, one only needs to examine the solutions that the DSA-FF algorithm produces on the various connection permutations. As a result, symmetric solutions, i.e., solutions derived from a DSA-FF solution by permuting blocks of spectrum slots, are eliminated from consideration. Clearly, the min-max DSA problem remains NP-hard [19]. However, as we explain in [1], the number of symmetric solutions is exponential. Therefore, the size of the symmetry-free solution (i.e., permutation) space is orders of magnitude smaller compared to that explored by conventional integer linear programming (ILP) formulations [20], [21].

Furthermore, with the elimination of symmetric solutions, the permutation space has a well-defined structure that is amenable to recursive and multi-threaded exploration. In particular, the recursive branch-and-bound RFF algorithm we developed in [1], [18] for min-max SA can be readily extended to tackle the min-max DSA problem. This DSA-RFF

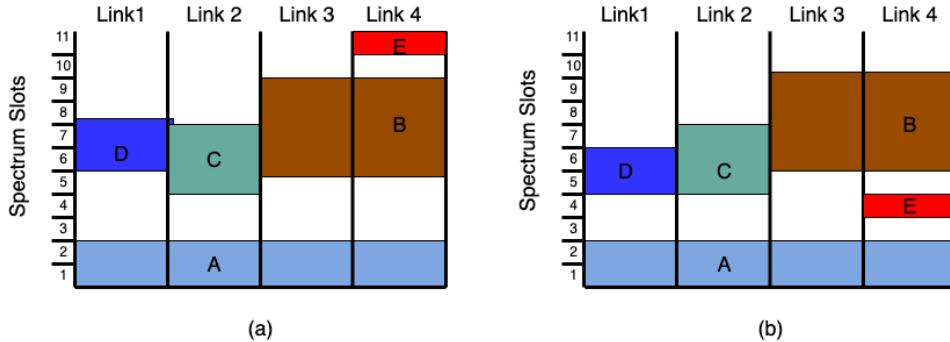


Fig. 1. (a) A feasible solution to a min-max DSA problem instance on a 4-link chain with $C = 5$ connections A, B, C, D, E . (b) The feasible (and optimal) solution constructed by the proof of Lemma 3.1.

algorithm operates identically to RFF; the only difference is that it applies DSA-FF, rather than the pure FF algorithm, as it recursively and incrementally builds and evaluates the connection permutations. Rather than repeating the details of RFF, we refer the reader to [1] for the operation of the algorithm and to [18] for two alternative multi-threaded implementations. We also note that min-max DSA includes more constraints than min-max SA, hence its solution space is smaller (since permutations that might yield feasible solutions for min-max SA may lead to infeasible solutions for min-max DSA). Therefore, we expect DSA-RFF to be more time-efficient than pure RFF. Nevertheless, an experimental evaluation of DSA-RFF is outside the scope of this paper and will be the subject of future research.

IV. THE MIN-FRAG SA PROBLEM

A spectrum *fragment* on a link is a block of one or more contiguous unused (free) slots located *between* two assigned spectrum blocks. The min-max objective we have discussed so far attempts to pack the assigned spectrum blocks tightly within lower-index slots so as to minimize spectrum fragmentation and allow for growth in demand. A different way for achieving the same goal would be to construct solutions that minimize the number of unused slots contained within spectrum fragments. Therefore, we define the min-frag SA problem as:

Definition 4.1 (min-frag SA):

The min-max SA problem of Definition 2.1 with this new objective:

Objective: Minimize the sum of spectrum fragment sizes over all links of the network.

We now show that the optimal solutions to the min-max SA and min-frag SA problems can be very different.

Lemma 4.1: An optimal solution to the min-max SA problem is not necessarily optimal for the min-frag SA problem, and vice versa.

Proof. By counter-example.

Figure 2(a) shows an optimal solution to a min-max SA problem instance on a 4-link chain with $C = 9$ connections; each connection is represented similarly to Figure 1. The

solution is optimal since the index of the highest slot used is 13, which is equal to the lower bound, i.e., the number of slots necessary to carry the traffic on connections using link 3. This is also a feasible solution to the min-frag SA problem with an objective value of 2 representing the sum of the two one-slot spectrum fragments on links 1 and 4.

Figure 2(b), on the other hand, shows an optimal solution to the min-frag SA problem instance where there is a single spectrum fragment of one slot. This solution is also a feasible solution for the min-max objective but the index of the highest slot used is 14, higher than that of the optimal min-max solution in Figure 2(a).

Hence, the min-max optimal solution of Figure 2(a) is suboptimal for the min-frag objective, whereas the min-frag optimal solution of Figure 2(b) is suboptimal under the min-max objective. ■

As we can see from Figure 2, there is a tradeoff between the min-max and min-frag objectives. Specifically, to minimize the index of the highest slot used it may be necessary to introduce additional spectrum fragments as in Figure 2(a); and conversely, minimizing the spectrum fragments may require the use of higher-indexed slots as in Figure 2(b).

Nevertheless, the min-max SA and min-frag SA problems have something in common, namely, that they are both solved by the FF algorithm. We showed that FF solves the min-max SA problem in [1]; to show that FF also solves the min-frag SA problem we first prove the following more general result.

Lemma 4.2 (FF Universality Property for min-frag SA):

Let S be a feasible solution to the min-frag SA problem. There exists a permutation P of the C connections such that applying the FF algorithm in the order implied by P yields S or another feasible solution S' with equal or lower objective value.

Proof. By construction.

The construction is very similar to that of the proof of Lemma 3.1 in that we consider the connections in solution S one at a time and in the same order, remove their assigned spectrum block from the solution, and attempt to place the block at a lower-indexed starting slot. The main difference

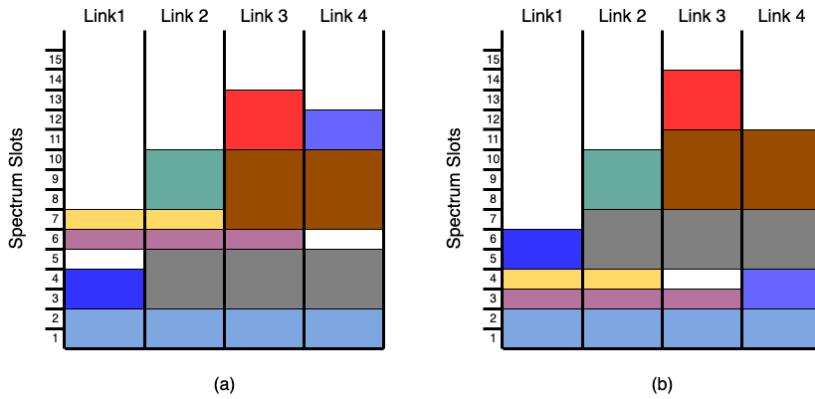


Fig. 2. (a) An optimal solution to a min-max SA problem instance on a 4-link chain with $C = 9$ connections; also a feasible (but suboptimal) solution to the corresponding min-frag SA problem instance. (b) An optimal solution to the min-frag SA problem instance; also a feasible (but suboptimal) solution to the corresponding min-max SA problem instance.

is that there are no guard bands to consider when placing a spectrum block.

Note also that if a spectrum block is moved to a lower-indexed starting slot, the objective function cannot increase (but may decrease). To see this, consider a spectrum block of size t that is moved to lower-indexed slots. There are two cases to consider: (a) links in which another spectrum block is assigned slots of higher index than this block, and (b) links for which this spectrum block is the one with the highest-index slots. In case (a), removing the block creates a spectrum fragment of size t , but placing the block at a lower-indexed slot removes a fragment of the same size; therefore, the net change in the objective function is zero. Case (b) has two sub-cases that correspond to the spectrum blocks of connections D and E in Figure 1(a), ignoring the guard bands implicit in that figure. In both sub-cases, when the spectrum blocks of the two connections are moved to a lower-indexed starting slot, the objective function (i.e., the sum of fragment sizes) decreases. ■

Similar to Corollary 3.1, we have this result:

Corollary 4.1: There exists a permutation P^* of the C connections such that $FF(P^*)$ is an optimal solution to the min-frag SA problem.

Therefore, the symmetry-free model, along with all the implications we discussed earlier, also extends to the min-frag SA problem. In particular, the RFF algorithm [1] may be modified in a straightforward manner to explore the permutation space for an optimal solution under the min-frag, rather than the min-max, objective; multi-threaded implementations of RFF [18] are also applicable in this case.

However, there is a crucial difference between the min-max and min-frag objectives. Specifically, as the FF algorithm operates on a certain permutation to build a solution under the min-max objective, the objective value (i.e., the index of the highest assigned slot) is monotonically non-decreasing as a function of the number of connections considered. This property allows RFF to determine if a permutation that has

been partially considered will not lead to a solution that is better than a baseline one; if so, RFF eliminates the current subtree of the permutation space and backtracks to explore a different subtree. The min-frag objective, on the other hand, does not have the same property: a connection further down in a permutation may fit within an existing spectrum fragments, thus *reducing* the objective value. Consequently, RFF will have to operate on each and every connection of a permutation without the possibility of backtracking; hence, it will take more time to explore the same fraction of the permutation space than for the min-max objective. As we mentioned in the previous section, an experimental investigation of this modified version of RFF is outside the scope of this paper, but we are considering other options for speeding up the exploration of the min-frag solution space.

V. THE MIN-FRAG DSA PROBLEM

Recall that the min-frag objective attempts to minimize the number of unused slots stranded in fragments between assigned spectrum blocks. In the presence of guard band constraints, however, certain slots within a spectrum fragment may represent a required guard band and cannot be considered as “unused.” Consider, for instance, the 3-slot fragment in link 1 of Figure 1(a). Since $g_{AD} = 2$, two slots of that fragment represent the required guard band between connections A and D , hence only one slot is unused. Therefore, once connection D is moved to slots 5 and 6 in the modified solution of Figure 1(b), the two-fragment slot on link 1 represents the guard band and does not include any unused slots.

With these observations, the min-frag DSA problem is derived from min-max DSA with the new objective of minimizing the size of the spectrum fragments that do not represent guard bands. More formally:

Definition 5.1 (min-frag DSA):

The min-max DSA problem of Definition 3.1 with this new objective:

Objective: Minimize, over all links of the network, the sum of spectrum fragment sizes *minus* the sum of all guard bands.

Using a construction proof similar to the ones we presented in the previous two sections, we can show that DSA-FF is also a universal algorithm for the min-frag DSA problem; specifically, we have the following lemma:

Lemma 5.1 (Universality Property for min-frag DSA): Let S be any feasible solution to the min-frag DSA problem. There exists a permutation P of the C connections such that applying the DSA-FF algorithm in the order implied by P yields S or another feasible solution S' with equal or lower objective value.

Similar to our earlier observations, the implication of Lemma 5.1 is that the recursive DSA-RFF algorithm we discussed in Section III-B may be used to search the symmetry-free solution (permutation) space for a permutation on which the DSA-FF algorithm produces an optimal solution to min-frag DSA. As we discussed in the previous section, the min-frag objective is such that the objective value is not monotonically non-decreasing as a function of the number of connections in the permutation to which the DSA-FF algorithm has assigned spectrum. Therefore, DSA-RFF cannot operate in a branch-and-bound mode, and has to completely evaluate each permutation, i.e., perform an exhaustive search of the permutation space. A quantitative evaluation of DSA-RFF for this problem is the subject of ongoing research in our group.

VI. CONCLUDING REMARKS

We have studied two variants of the spectrum assignment problem, one without guard band constraints (i.e., SA) and one with such constraints (i.e., DSA). For each variant, we considered two objectives: min-max, which minimizes the highest spectrum slot index assigned to any connection, and min-frag, which minimizes the sum of the number of unused slots in spectrum fragments across all links. We have shown that the FF algorithm may be used to solve optimally the min-max SA and min-frag SA problems. Likewise, the DSA-FF algorithm which operates similarly to FF but takes into consideration the guard band constraints, may be used to solve optimally the min-max DSA and min-frag DSA problems. We also showed that while the two objectives, min-max and min-frag, attempt to minimize fragmentation, they may lead to different solutions.

Our results transform all four spectrum assignment problems into permutation problems with a well-structured and symmetry-free solution space. With this new insight into the nature of spectrum assignment, we were able to develop intuitive recursive algorithms, RFF and DSA-RFF, that differ radically from conventional approaches and may be used to explore the corresponding solution (permutation) spaces effectively and in parallel. Our work shows that seemingly disparate variants of spectrum assignment have a common underlying structure and may be tackled using similar approaches. Finally, since spectrum assignment underlies most optical network optimization problems, our results represent

the first step towards new algorithmic approaches to optical network design.

REFERENCES

- [1] G. N. Rouskas and C. Bandikatla, "Recursive first fit: A highly parallel optimal solution to spectrum allocation," *IEEE/Optica Journal of Optical Communications and Networking*, vol. 14, pp. 165–176, April 2022.
- [2] J. M. Simmons, *Optical Network Design and Planning*. Springer, 2008.
- [3] B. Jaumard, C. Meyer, and B. Thiongane, "Comparison of ILP formulations for the RWA problem," *Optical Switching and Networking*, vol. 3-4, pp. 157–172, 2007.
- [4] M. Klinkowski, P. Lechowicz, and K. Walkowiak, "Survey of resource allocation schemes and algorithms in spectrally-spatially flexible optical networking," *Optical Switching and Networking*, vol. 27, no. C, pp. 58–78, 2018.
- [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*, vol. 13, pp. 34–48, July 2014.
- [6] R. Dutta and G. N. Rouskas, "A survey of virtual topology design algorithms for wavelength routed optical networks," *Optical Networks*, vol. 1, pp. 73–89, January 2000.
- [7] H. Wang and G. N. Rouskas, "Hierarchical traffic grooming: A tutorial," *Computer Networks*, vol. 69, pp. 147–156, August 2014.
- [8] B. Chen, G. N. Rouskas, and R. Dutta, "Clustering methods for hierarchical traffic grooming in large-scale mesh WDM networks," *Journal of Optical Communications and Networking*, vol. 2, pp. 502–514, August 2010.
- [9] D. Zhou and S. Subramaniam, "Survivability in optical networks," *IEEE Network*, vol. 14, pp. 16–23, November/December 2000.
- [10] S. Talebi, E. Bampis, G. Lucarelli, I. Katib, and G. N. Rouskas, "Spectrum assignment in optical networks: A multiprocessor scheduling perspective," *Journal of Optical Communications and Networking*, vol. 6, pp. 754–763, August 2014.
- [11] J. Simmons and G. N. Rouskas, "Routing and wavelength (spectrum) allocation," in *B. Mukherjee, I. Tomkos, M. Tornatore, P. Winzer, and Y. Zhao (Editors), Springer Handbook of Optical Networks*, Springer, 2020.
- [12] H. Zang, J. P. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *Optical Networks*, vol. 1, pp. 47–60, January 2000.
- [13] Y. Zhu, G. N. Rouskas, and H. G. Perros, "A comparison of allocation policies in wavelength routing networks," *Photonic Network Communications*, vol. 2, pp. 265–293, August 2000.
- [14] D. S. Johnson, A. J. Demers, J. D. Ullman, M. R. Garey, and R. L. Graham, "Worst-case performance bounds for simple one-dimensional packing algorithm," *SICOMP*, vol. 3, no. 4, 1974.
- [15] M. R. Garey, R. L. Graham, and J. D. Ullman, "Worst-case analysis of memory allocation algorithms," in *Proceedings of the Fourth Annual ACM Symposium on Theory of Computing*, pp. 143–150, 1972.
- [16] M. R. Garey, R. L. Graham, D. S. Yao, and A. Chi-Chih, "Resource constrained scheduling as generalized bin packing," *Journal of Combinatorial Theory, Series A*, vol. 21, no. 3, pp. 257–298, 1976.
- [17] G. Dosa and J. Sgall, "First fit bin packing: A tight analysis," in *30th International Symposium on Theoretical Aspects of Computer Science (STACS 2013)*, pp. 538–549, 2013.
- [18] G. N. Rouskas, S. Gupta, and P. Sharma, "Experimental evaluation of a symmetry-free parallel algorithm for spectrum allocation," in *Proceedings of IEEE GLOBECOM 2022*, December 2022.
- [19] H. Wu, F. Zhou, Z. Zhu, and Y. Chen, "On the Distance Spectrum Assignment in Elastic Optical Networks," *IEEE/ACM Transactions on Networking*, vol. 25, pp. 2391–2404, August 2017.
- [20] B. Jaumard, C. Meyer, and B. Thiongane, "ILP formulations for the routing and wavelength assignment problem: Symmetric systems," in *Handbook of Optimization in Telecommunications*, pp. 637–677, Springer US, 2006.
- [21] R. Ramaswami and K. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Transactions on Networking*, vol. 3, pp. 489–500, October 1995.