Routing and Spectrum Allocation in OFDM-based Optical Networks with Elastic Bandwidth Allocation

[1]K. Christodoulopoulos, [2]I. Tomkos, [1]E. A. Varvarigos

[1] Computer Engineering and Informatics Department, University of Patras, Greece, and
Research Academic Computer Technology Institute, Patras, Greece
{kchristodou, manos}@ceid.upatras.gr

[2] Athens Information Technology Center, Paiania, Greece
itom@ait.edu.gr

Abstract— **Orthogonal Frequency Division Multiplexing (OFDM) has been recently proposed as a modulation technique for optical networks, due to its good spectral efficiency and impairment tolerance. Optical OFDM is much more flexible compared to traditional WDM systems, enabling elastic bandwidth transmissions. We consider the planning problem of an OFDM-based optical network where we are given a traffic matrix that includes the requested transmission rates of the connections to be served. Connections are provisioned for their requested rate by elastically allocating spectrum using a variable number of OFDM subcarriers. We introduce the Routing and Spectrum Allocation (RSA) problem, as opposed to the typical Routing and Wavelength Assignment (RWA) problem of traditional WDM networks, and present various algorithms to solve the RSA. We start by presenting an optimal ILP RSA algorithm that minimizes the spectrum used to serve the traffic matrix, and also present a decomposition method that breaks RSA into two substituent subproblems, namely, (i) routing and (ii) spectrum allocation (R+SA) and solves them sequentially. We also propose a heuristic algorithm that serves connections one-by-one and use it to solve the planning problem by sequentially serving all traffic matrix connections. To feed the sequential algorithm, two ordering policies are proposed; a simulated annealing meta-heuristic is also proposed to obtain even better orderings. Our results indicate that the proposed sequential heuristic with appropriate ordering yields close to optimal solutions in low running times.**

Keywords- Optical OFDM, Elastic spectrum optical paths, Planning (offline) problem, Routing and spectrum allocation, Spectrum continuity constraint.

I. INTRODUCTION

The continuous growth of consumers IP traffic in combination with emerging high-rate applications, such as video on demand, high definition TV, cloud computing and grid applications, require a cost-efficient and scalable networking infrastructure. To meet the increasing capacity requirements, recent innovations in optical communication systems including advanced modulation formats together with digital equalization in electrical domain have enabled per-channel bandwidths of 40 and 100 Gb/s with improved transmission distance features. The high channel capacity and the extended optical reach enable high rate transmission over multiple WDM links and wavelength cross-connects (WXCs) without optical-electrical-optical (OEO) regeneration. Thus, wavelength routed transparent mesh networks seem to offer a cost-effective solution for high capacity transport networks.

Although wavelength routed transparent networks offer obvious advantages, they still have a drawback due to their rigid and coarse granularity. Currently, wavelength-routed networks require full allocation of a wavelength to a connection even when the traffic between the end nodes is not sufficient to fill the entire capacity. Wavelength level granularity leads to inefficient capacity utilization, a problem expected to become even more significant with the deployment of higher capacity WDM networks (i.e., systems of 40 and 100 Gbps per channel).

The need for flexibility and efficiency requires the use of resources with subwavelength granularity. In addition, high-end applications requiring super-wavelength capacity would benefit from a more agile network infrastructure. Ideally, an adaptive network would have a fine granularity so as to elastically provide the required capacity to sub- or super-wavelength demands. Approaches such as optical packet switching (OPS) and optical burst switching (OBS) that meet these requirements have been proposed in the literature. However, these approaches can be viewed as long-term solution since their enabling technologies are not yet mature.

Recently, Orthogonal Frequency-Division Multiplexing (OFDM) has been proposed as a modulation technique in optical networks [1]-[1]. Optical OFDM distributes the data on a high number of low data rate subcarriers and, thus, can provide fine-granularity capacity to connections by the elastic allocation of subcarriers according to the connection demands. Enabling technologies, such as bandwidth-variable transponders and bandwidth-variable WXCs, have been designed and demonstrated in Spectrum-sLICed Elastic optical path network ("SLICE") [4][5]. To achieve high spectral efficiency the bandwidth-variable transponder generates an optical signal using just enough spectral resources, in terms of subcarriers, to serve the client demand. Every WXC on the route allocates a cross-connection with the corresponding spectrum to create an appropriate-sized end-to-end optical path. To route and receive the data with acceptable signal performance, adjacent optical paths require appropriate spectrum separation, implemented by spectrum guardbands [6].

The use of optical OFDM as a bandwidth-variable and highly spectrum-efficient modulation format in SLICE can provide scalable and flexible sub- and super-wavelength granularity, compared to the conventional, fixed-bandwidth fixed-grid WDM network. However, this new concept poses new challenges on the networking level, since the routing and wavelength assignment (RWA) algorithms of traditional WDM networks are no longer directly applicable. The wavelength continuity constraint of traditional WDM networks is transformed to a spectrum continuity constraint. Moreover, a connection requiring capacity larger than one OFDM subcarrier has to be assigned a number of contiguous subcarrier slots. To address these issues, new routing and spectrum allocation (RSA) algorithms, as well as appropriate extensions to network control and management protocols have to be developed.

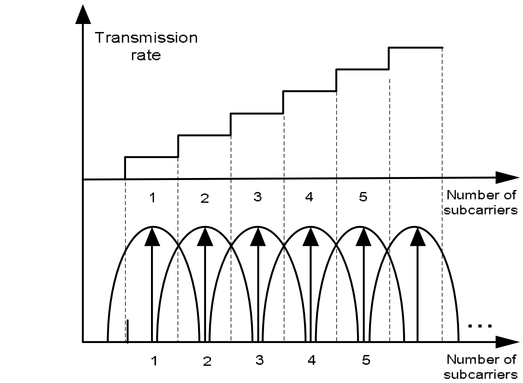

Figure 1. Variable bandwidth transmission by elastically controling the number of OFDM subcarriers.

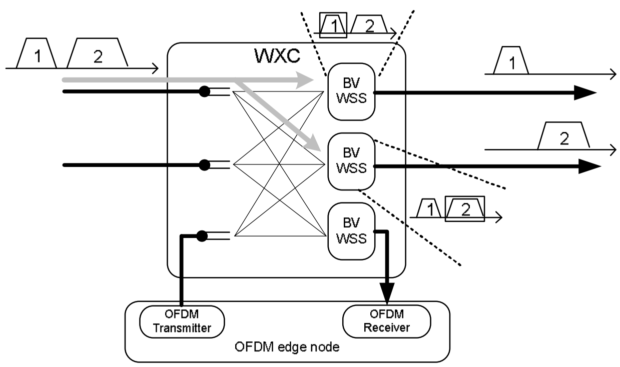

Figure 2. Spectrum flexible WXC

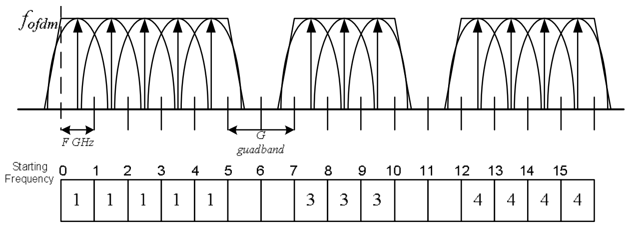

Figure 3. Spectrum allocation as a table of subcarier slots assigned to connections. Spectrum guardaband of *G* subcariers is used to separate the path flows so as to route and receive them with acceptable signal performance.

In this paper we focus on the routing and spectrum allocation (RSA) problem in OFDM-based elastic optical networks. We consider the planning phase (offline problem) of such network where we are given a traffic matrix with the requested transmission rates of all connections. Our objective is to serve the connections through adequate spectrum allocation, with the constraint that no spectrum overlapping is allowed among these connections, and minimize the utilized spectrum. We initially present an optimal RSA integer linear programming (ILP) formulation. To reduce the size of the RSA problem we decompose it into its substituent sub-problems, namely (i) routing and (ii) spectrum allocation, which are solved sequentially (R+SA), without, however, being guaranteed to find an optimal solution for the joint RSA problem. Since these formulation cannot be solved efficiently for large networks, we present a heuristic algorithm that solves the planning RSA problem by sequentially serving one-by-one the connections. The ordering in which the connections are served plays an important role in the performance of such a heuristic algorithm. We propose and evaluate two ordering policies and also use a simulated annealing meta-heuristic to find good orderings that yield near-optimal performance.

II. OFDM-Based Optical Networks

In this section we shortly present the architecture of an OFDM-based optical network.

In OFDM the data is transmitted over multiple orthogonal subcarriers. This technology has been widely implemented in various systems, such as wireless local area network (LAN) and asymmetric digital subscriber line (ADSL). Recently, extensive research efforts have focused on an optical version of OFDM as a means to overcome transmission impairments [1][3]. Besides the advantages of low symbol rate of each subcarrier and coherent detection that mitigate the effects of physical impairments, OFDM also brings unique benefits in terms of spectral efficiency. Moreover, OFDM enables elastic bandwidth transmission realized by allocating a variable number of low-rate subcarriers for a transmission (Figure 1).

The signal transmitted over the optical path (using the spectrum determined by the volume of client traffic) is routed through bandwidth variable (BV) wavelength cross-connects (WXCs) towards the receiver. Every BV WXC on the route allocates a cross-connection with the corresponding spectrum to create an appropriate-sized end-to-end optical path. To do so, the WXC has to configure its switching window in a contiguous manner according to the spectral width of the incoming optical signal. MEMS- or liquid crystal-based wavelength-selective switches (WSSs) can be employed as BV WXC switching elements. Figure 2 shows the switching operation of a BV WXC. Figure 3 presents an example of the utilization of a link in an OFDM-based optical network. Signals of different optical paths are multiplexed in the frequency domain. Each optical path can utilize a different number of OFDM subcarriers that are mapped to subcarrier slots. The use of optical OFDM increases the overall spectral efficiency and improves the granularity and flexibility of the network when compared to a fixed-grid WDM network.

III. Networking Issues

A. *Transmission rate service guarantee*

Although the transmission rate of a connection can fluctuate with time, from the operators' perspective the network has to be planned to guarantee the service of a connection for a requested rate. This would translate to non-overlapping spectrum allocation to all connections for their requested transmission rates. Although planning a network in this way may result in a waste of spectrum resources, when the connections under-utilize their provisioned bandwidth, there are still major gains compared to traditional WDM networks. The gains are (i) the high spectrum efficiency due to OFDM format, (ii) the fine granularity of low-rate subcarrier level, (iii) impairment tolerance due to OFDM features, and (iv) a possible reduction in power consumption by partially deactivating the transmitters adjusting them to the rate at a specific time. Moreover, at a specific time, unused spectrum can be shared and allocated to connections that surpass their requested transmission rates or to best-effort traffic, but this spectrum will be de-allocated when the initially provisioned connection requires it.

Additional gains in spectrum efficiency can be obtained by network planning based on time scheduling using time-varying traffic models, or by allowing overlapping spectrum allocation, based on stochastic traffic models. For example, connections that have complementary transmission rates in time, in the sense that when the rate of a connection increases, the opposite tends to happen to that of another one, could be served by shared spectrum slots. Another approach would be to have for each connection a requested transmission rate, the service of which is guaranteed, and a probabilistic model for exceeding this rate, where spectrum overlapping between connections could be performed. In this paper we focus on planning an OFDM-based optical network so as to guarantee the requested rates of the connections, assuming no spectrum overlapping between them. In the future we plan to examine planning methods for OFDM networks where spectrum overlapping is allowed, based on time or probabilistic traffic models, as a way to further improve spectrum utilization efficiency.

B. Routing and Spectrum Assignment (RSA) requirements

We now argue that typical routing and wavelength assignment (RWA) algorithms devised for standard grid WDM systems are not applicable to OFDM-based optical networks.

To transform the OFDM routing and spectrum allocation problem to a typical RWA formulation we have to map a subcarrier to a wavelength of the same capacity. Thus, a connection that requires a number of subcarriers is utilized by the same number of wavelengths. Although typical RWA algorithms are able to serve such a connection, the wavelengths that are going to be provisioned are not going to be contiguous in the spectrum domain. To select contiguous wavelengths the RWA algorithms have to be modified accordingly, and it is unclear how this can be done with most algorithms. Moreover, the majority of RWA algorithms proposed in the literature utilize variables and constraints that depend on the number of wavelengths, which in a typical WDM network seldom exceeds 80, beyond which the operators have to resort to a parallel network, installing additional fibers per link. The high number of OFDM subcarrier (of the order of several hundreds) poses limitations to the applicability of the traditional RWA algorithms.

From the above discussion it is clear that we have to develop new algorithms that would (i) serve a connection utilizing a contiguous spectrum and (ii) use variables and constraints that do not depend on the number of subcarriers.

IV. ROUTING AND SPECTRUM ALLOCATION ALGORITHMS

We assume an OFDM optical network as presented in Section II and [4]. The spectral granularity of the transmitters and WXCs is one subcarrier corresponding to F GHz of spectrum. The capacity of a subcarrier is equal to C Gbps. Although C can adapt depending on the used OFDM level, i.e., BPSK, QPSK, 8-QAM, or higher, we will assume a given constant C. To route the paths through the WXC a guardband of G subcarriers has to separate adjacent spectrum paths. We assume that elastic OFDM transmitters can be tuned to utilize a given number of subcarriers forming a continuous spectrum with a step of F GHz. In this context, the spectrum starting from frequency f_{OFDM} is divided in subcarrier slots of F GHz (Figure 3). Serving a connection i that requires T_i subcarriers is translated to finding a starting subcarrier frequency f_i after which it can use T_i contiguous subcarriers (in addition to the guardbands). For example, with respect to Figure 3, connection '*2*' that requires 3 subcarriers is assigned the starting frequency f_2=7, assuming that f_{OFDM} corresponds to zero frequency.

A network topology is represented by a connected graph $G=(V,E)$. V denotes the set of nodes, which we assume to be equipped with bandwidth variable WXCs. E denotes the set of (point-to-point) single-fiber links. Let $N=|V|$ and $L=|E|$ denote the number of nodes and the number of links of the network. The planning version of the RSA problem assumes an a-priori known traffic scenario. We assume that there is a function f that connects the transmission rate of a connection and the allocated spectrum, so that a bandwidth demand of B_i can be mapped to a demand of T_i subcarriers (e.g., $T_i=\left\lceil B_i/C \right\rceil$, for a given C). Thus, the traffic scenario is given in the form of a matrix of non-negative integers T, called the spectrum traffic matrix. Then T_{sd} denotes the number of subcarriers required for the communication between source s and destination d. We assume that for connection (s,d) we utilize a continuous spectrum (a continuous set of subcarriers), so that T_{sd} subcarriers are allocated over a single path that connects (s,d).

In the following, we propose algorithms for Routing and Spectrum Allocation (RSA) in OFDM optical networks. We assume that physical layer impairments (PLI) are not significant (due to the low symbol rate of OFDM subcarriers and coherent detection) [3], and thus are not accounted for in the proposed algorithms. The reader is referred to [7] for the problem of designing a typical WDM network under PLIs.

A. Combinatorial RSA Algorithms

1) Joint RSA Algorithm

We initially present an optimal integer linear programming (ILP) formulation [8] that minimizes the utilized spectrum.

For each commodity s-d we pre-calculate k paths. Let P_{sd} be a set of candidate paths for s-d and $P=\bigcup_{(s,d)} P_{sd}$ be the total set of candidate paths. Note that we can formulate the problem without using any set of predefined paths, but allow routing over any feasible path. Such an algorithm will give at least as good solutions as the algorithm that uses pre-calculated paths, but will use a much higher number of variables and constraints and thus would be less scalable. In any case, the optimal solution can be also found with an algorithm that uses pre-calculate paths, given a large enough set of paths.

Variables:
x_p: Boolean variable that denotes the utilization of path $p\in P$ (x_p equals to 0 if path p is not utilized, and 1 if p is utilized)
f_{sd}: Integer variable that denotes the starting frequency for connection (s,d). Frequency f_{sd} is relative to f_{OFDM}. Assuming $T_{total}=\sum_{(s,d)} T_{sd}$, we have $0 \le f_{sd} < T_{total}$.
$\delta_{sd,s'd'}$: Boolean variable that equals 0 if the starting frequency of connection (s',d') is smaller than the starting frequency of connection (s,d) (i.e., $f_{s'd'}<f_{sd}$), and 1 otherwise (i.e., $f_{sd}<f_{s'd'}$).
c: maximum utilized spectrum slot

ILP Routing and Spectrum Allocation (RSA) formulation:

minimize c

subject to the following constraints:

- Cost function

$$c \geq f_{sd} + T_{sd}\text{, for all } (s,d) \text{ pairs} \quad (1)$$

- Single path routing constraints

$$\sum_{p \in P_{sd}} x_p = 1\text{, for all } (s,d) \text{ pairs} \quad (2)$$

- Starting frequencies ordering constraints

For all commodities (s,d) and (s',d') that have $p_i \in P_{sd}$ and $p_j \in P_{s'd'}$, with p_i and p_j sharing at least one common link l ($\forall (s,d),(s',d') : \exists p_i \in P_{sd} \cap \exists p_j \in P_{s'd'} \cap (l \in p_i \cap l \in p_j)$), the following constraints are employed:

$$\delta_{sd,s'd'} + \delta_{s'd',sd} = 1, \quad (3)$$

$$f_{s'd'} - f_{sd} < T_{total} \cdot \delta_{sd,s'd'}, \quad (4)$$

$$f_{sd} - f_{s'd'} < T_{total} \cdot \delta_{s'd',sd}, \quad (5)$$

Constraints (3)-(5) ensure that either $\delta_{sd,s'd'}=1$, meaning that the starting frequency f_{sd} of connection (s,d) is smaller than the starting frequency $f_{s'd'}$ of (s',d') (i.e., $f_{sd} < f_{s'd'}$), or $\delta_{s'd',sd}=1$ (i.e., $f_{sd} > f_{sd}$). Note that f_{sd} and $f_{s'd'}$ are bounded by constant T_{total}, so their difference is always less than T_{total}.

- Spectrum continuity and non-overlapping spectrum allocation

For all commodities (s,d) and (s',d') that have $p_i \in P_{sd}$ and $p_j \in P_{s'd'}$, with p_i and p_j sharing at least one common link l, the following constraints are employed:

$$f_{sd} + T_{sd} + G - f_{s'd'} \leq (T_{total} + G) \cdot (1 - \delta_{sd,s'd'} + 2 - x_{pi} - x_{pj}) \quad (6)$$

$$f_{s'd'} + T_{s'd'} + G - f_{sd} \leq (T_{total} + G) \cdot (1 - \delta_{s'd',sd} + 2 - x_{pi} - x_{pj}) \quad (7)$$

When one (or both) of the paths p_i and p_j is not utilized ($x_{pi} \neq 1$ or $x_{pj} \neq 1$), then we do not have to consider the overlapping of their spectrum. In this case, constraints (6) and (7) are deactivated (hold always, irrespectively of f_{sd} and $f_{s'd'}$), since the right hand side of the constraints take a value larger than T_{total}, which is always higher than the left hand side.

Now, assume that both paths p_i and p_j are utilized ($x_{pi}=1$ and $x_{pj}=1$). Then one of the constraints (6) or (7) are activated according to the values of $\delta_{sd,s'd'}$ and $\delta_{s'd',sd}$. In particular, constraint (6) is activated when $\delta_{sd,s'd'}=1$ (that is when $f_{sd} < f_{s'd'}$), in which case (6) becomes:

$$f_{sd} + T_{sd} + G \leq f_{s'd'},$$

ensuring that the spectrum used by the two connections (s,d) and (s',d') do not overlap. When $\delta_{sd,s'd'}=1$, then $\delta_{s'd',sd}=0$, and constraint (7) is deactivated, since (7) becomes

$$f_{s'd'} + T_{s'd'} - f_{sd} \leq T_{total},$$

which holds always irrespectively of f_{sd} and $f_{s'd'}$.

In a similar manner, constraint (7) is activated when $\delta_{s'd',sd}=1$ (i.e., when $f_{sd} > f_{s'd'}$) and constraint (6) is deactivated.

Thus, constraints (6) and (7) ensure that the spectrums allocated to connections that utilize paths that have a common link do not overlap.

The above ILP algorithm finds the paths p (corresponding to $x_p=1$) and the starting frequencies f_{sd} of the connections over those paths so as to minimize the total used spectrum c. Spectrum continuity constraint is translated to non-overlapping spectrum allocation. Thus, the starting frequencies of the connections that utilize a common link are ordered so that their allocated spectrums do not overlap (accounting also for the required guardbands G in-between).

The number of variables and constraints used by the above ILP formulation depends on the overlapping of links between the paths considered (and thus depends on the network topology and the chosen k). In particular, Boolean variables $\delta_{sd,s'd'}$ and constraints (3)-(7) need to be employed only for commodities (s,d) and (s',d') that have at least one path that use a common link. In the worst case (where all commodities have paths that share links with all the other commodities), the formulation would require N^4 boolean $\delta_{sd,s'd'}$ variables, N^4 equality constraints for (3) and $4 \cdot N^4$ inequality constraints for (4)-(7). Practically, however, the number of variables and constraints are much lower than N^4. The rest of the formulation uses: $k \cdot N^2$ boolean and N^2 integer variables for x_p and f_{sd}, respectively, and N^2 equality constraints for (1) and N^2 inequality constraints for (2).

Note that if we did not use pre-calculated paths in our formulation, but instead used a multicommodity flow formulation that would allow the routing of a commodity over all links, we would have to utilize the complete set of N^4 boolean $\delta_{sd,s'd'}$ variables and all the corresponding (3)-(7) constraints, in addition to flow constraints that would ensure the construction of paths among the links. This is the main reason we chose to use pre-calculated paths in our formulation. From the above description, we observe that the number of variables and constraints in the ILP formulation does not depend on the number of subcarriers, which was one of the key desired properties we had in mind when designing this algorithm.

2) *Decomposing the problem (R+SA)*

The algorithm presented in this section breaks the problem to (i) the routing (R) and (ii) the spectrum allocation (SA) subproblems and addresses each problem separately and sequentially. Note that by decomposing the problem, the joint optimum of the RSA problem might not be found.

a) *Routing (Multicommodity) Phase*

As in the joint RSA algorithm described above, we start by calculating for each commodity *s-d* a set of candidate paths P_{sd}. The routing problem is formulated using the boolean x_p variables previously introduced in the joint RSA algorithm.

ILP Routing formulation:

minimize: c_R

subject to the following constraints:

- Cost function

$$c_R \geq F_l\text{, for all } (s,d) \text{ pairs}$$

- Flow cost per link

$$F_l = \sum_{sd} \sum_{\{p \in P_{sd} \mid l \in p\}} T_{sd} \cdot x_p\text{, for all links } l \in E$$

- Single path routing constraints

$$\sum_{p \in P_{sd}} x_p = 1\text{, for all } (s,d) \text{ pairs}$$

The solution forms a set P^*, containing one path p_{sd} per connection (s,d). Set P^* is passed to the second phase.

b) *Spectrum Allocation Phase*

The spectrum allocation is similar to the joint RSA formulation presented in Section IV.A.1, but instead of using the set P of pre-calculated paths, it uses the set P^* of paths that was calculated in the routing phase (previous subsection). Thus, for each connection (s,d) one path is included in P^* and the corresponding variables x_p of the RSA formulation are reduced and set equal to one. The ILP spectrum allocation formulation is omitted for brevity purposes.

B. Sequential establishment of Demands

Since the above ILP formulations (present in either the RSA or R+SA algorithms) cannot be solved efficiently for large networks, we now present a solution to the planning problem that can scale to networks of large size. The proposed approach uses a pre-ordering phase and then a single demand heuristic RSA algorithm to sequentially serve the demands one-by-one.

1) Single Demand RSA Heuristic Algorithm

We assume that each link $l \in E$ is characterized by a subcarrier availability boolean vector $\overline{U_l} = [u_{li}] = (u_{l1}, u_{2}, \ldots, u_{ld})$ of length $d = T_{total}$, which is equal to the maximum number of subcarriers required in the worst case to serve all demands. The i^{th} element of $\overline{U_l}$, i.e. u_{li}, records the availability of the i^{th} subcarrier, and equals 1 if the i^{th} subcarrier is available, and 0 if it has already been used by a path flow. We can calculate the subcarrier availability boolean vector of a path p by the availability vectors of the links $l \in p$ that comprise it as follows:

$$\overline{U_p} = [u_{pi}] = [\underset{l \in p}{\&} u_{li}], \qquad (8)$$

where "&" denotes the boolean AND operation.

The single demand RSA algorithm works as follows. As previously, we pre-calculate in a pre-processing phase a set P_{sd} of k paths for each source destination pair (s,d). To serve a new connection (s,d) that requires T_{sd} subcarriers, the algorithm takes as input the spectrum availability $\overline{U_l}$ of all links $l \in E$. We first use Eq. (8) to calculate the spectrum availability $\overline{U_p}$ of all the candidate paths $p \in P_{sd}$. We then search each spectrum availability vector $\overline{U_p}$ for the first possible placement of T_{sd} subcarriers (along with the required $2 \cdot G$ guardbands). In other words, we search for $T_{sd}+2 \cdot G$ continuous 1's in the spectrum availability vector (or $T_{sd}+G$ at the starting and ending limits of the spectrum vector). The algorithm selects the path with the lowest indexed starting subcarrier. Void filling is performed, in the sense that voids of size greater than $T_{sd}+2 \cdot G$ can be utilized. After selecting the path and the starting frequency, we update the spectrum availability of the links that comprise the selected path by setting 0's to the corresponding spectrum slots.

The above described algorithm is a quick and efficient greedy algorithm that finds for each new connection demand the lowest feasible starting subcarrier among the set of pre-calculated candidate paths. Pre-calculation of paths is used for speeding the procedure, especially in the simulated annealing case to be described shortly. Note that, depending on the network topology, the number of different paths between s and d can be high. Note also that instead of selecting the path with the minimum starting subcarrier we can use other approaches. For example, as in other void filing algorithms, we can select the void that leaves free the smallest remaining number of subcarriers (so that the subcarrier space is not fragmented), etc.

2) Ordering the demands and Simulated Annealing

The above described single demand heuristic algorithm serves the demands of the traffic matrix, one-by-one, in some particular order. The ordering is quite important in this process, and it is expected that different orderings will result in different spectrum utilization. For the scope of this paper we have evaluated the following ordering policies:

- *Most Subcarriers First* (*MSF*) ordering: We order the connection demands in decreasing order of the number of their requested subcarriers, and serve first the demand that requires the highest number of subcarriers.
- *Longest Path First* (*LPF*) ordering: We order the connection demands in decreasing order of the number of links their shortest paths utilize, and serve first the demand whose shortest path utilizes the highest number of links.

We also use a simulated annealing (SA) meta-heuristic to find good orderings that would provide good spectrum allocation solutions. Assuming an ordering of M connection demands, $\{(s_1,d_1),(s_2,d_2),\ldots,(s_M,d_M)\}$, we define the neighbor of this ordering as an ordering in which (s_i,d_i) is interchanged with (s_j,d_j). To generate a random neighbor we choose pivots (s_i,d_i) and (s_j,d_j) uniformly among the M demands.

The simulated annealing meta-heuristic works as follows. We start with MSF ordering and calculate its cost (viewed as "energy" in the simulated annealing setting) by serving the connection demands one-by-one, using the single demand heuristic algorithm described in the previous subsection IV.B.1 (this is the "fitness function"). We use the neighbor creation procedure described above to create some neighbors and follow the standard simulated annealing iterative procedure.

C. Extensions

The algorithms presented in the previous sections can be extended so as to incorporate additional features. An interesting extension is to examine the spectrum gains that can be obtained by allocating non-contiguous spectrum to the connections. Another important extension would be to enable the RSA algorithms to choose the OFDM level (BPSK, QPSK, 8-QAM, or higher) depending e.g. on the length of the path to be used. This would imply that the capacity C of subcarriers could differ for different paths. A possible solution to this problem would require the definition of path-related subcarrier demands (e.g., $T_{sd,p}$). Also, it would be interesting to examine the cost benefits that can be obtained by using a single transmitter to serve more than one connections, so as to reduce the total number of transmitters in the network. Finally, our future plans includes the development of algorithms that allow overlapping spectrum allocation based on time or probabilistic traffic models.

V. Performance Results

We evaluated the performance of the proposed algorithms through simulation experiments. We used Matlab to implement the algorithms, LINDO API [8] for ILP solving, and Matlab built-in simulated annealing meta-heuristic. Subcarrier traffic matrices were created as follows. The number of subcarriers T_{sd} required for connection (s,d) was chosen uniformly between 0 and D. In our experiments we used two different values of D (D=4 and D=40), to represent low and high load cases. The results were averaged over 100 randomly created matrices.

A. Small network experiments

Table I presents the results for the topology of Figure 4. We observe that for both values of D the optimal RSA ILP algorithm was able to find solutions. Decomposing the problem (R+SA ILP) reduces the average running time, while optimality is lost only in two traffic instances of high load (D=40). Regarding the sequential heuristic algorithm, MSF ordering performs slightly better than LPF ordering. LPF ordering takes into account the length of the paths that do not play a significant role in this small network. Simulated annealing (SA) improved vastly the performance of the sequential heuristic algorithm. Especially for low load (D=4) the ordering found by simulated annealing utilized the optimal spectrum in almost all traffic instances. Note that in Table I we present results for simulated annealing with 100 and 1000 iterations. We do not present results for a higher number of iterations, since the performance was only slightly improved, while the running time exceeded that of the optimal RSA ILP.

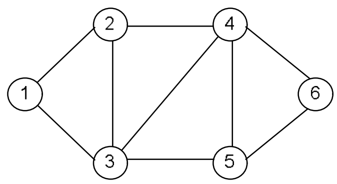

Figure 4. Small network topology used for simulation experiments

B. *Realistic network experiments*

Table II presents the result for the generic DT network topology consisting of 14 nodes and 46 directed links [10]. The optimal RSA ILP algorithm was unable to produce results for this network in reasonable time. Moreover, the results presented for the decomposed algorithm (R+SA ILP) were taken by stopping the ILP after 2 hours for each traffic instance. In particular, the routing (R) problem was solved optimally quite easily, while the spectrum allocation (SA) problem was quite hard, and thus we were not able to obtain optimal SA solutions in 2 hours time. From table II we observe that the decomposed R+SA ILP algorithm found the best solutions (although not the optimal ones within 2 hours). With respect to the sequential heuristic algorithm, LPF ordering slightly outperforms MSF ordering, especially for light load, in contrast to the results presented in the previous section. However, in this topology the path lengths differ among the connections. Connections traversing longer paths utilize more spectrum resources and thus performance benefits can be obtained by early serving these connections in an empty network, as LPF does. Again simulated annealing improved the performance of the sequential heuristic algorithm, and was able to find good orderings that produced results near to the R+SA ILP algorithm. Note that for this network we present results for simulated annealing with 1000 and 10000 iterations (as opposed to 100 and 1000 in the previous section).

VI. CONCLUSIONS

Recently, research has focused on optical OFDM as a spectrum-efficient modulation format that can provide elastic bandwidth transmission. We consider the problem of planning an OFDM-based optical network in which connections are provisioned based on their requested transmission rate (given in a traffic matrix) assuming no spectrum overlapping between them. We introduced the Routing and Spectrum Allocation (RSA) problem and presented various algorithms to solve it. We initially presented an optimal ILP algorithm that minimizes the spectrum used to serve the traffic matrix and then a decomposition method that breaks RSA into (i) routing and (ii) spectrum allocation subproblems and solves them sequentially. We also proposed a heuristic algorithm that serves connections one-by-one and used it to solve the planning problem by sequentially serving all connections. Two ordering policies and a simulated annealing meta-heuristic were used to feed the sequential algorithm. Our results indicate that the proposed sequential heuristic combined with an appropriate ordering can give close to optimal solutions in low running times.

TABLE I. PERFORMANCE RESULTS FOR THE SMALL NETWORK

Load	Algorithm	Average spectrum utilization (#subcarriers)	Average running time (sec)
D=4	RSA ILP	7.78	16.8
	R+SA ILP	7.78	8.67
	MSF+heuristic	8.24	0.29
	LPF+heuristic	8.32	0.29
	SA (100 iterations)	7.91	2.6
	SA (1000 iterations)	7.82	17.6
D=40	RSA ILP	70.7	95.2
	R+SA ILP	70.9	12.9
	MSF+heuristic	76.3	0.35
	LPF+heuristic	77.1	0.35
	SA (100 iterations)	73.1	6.6
	SA (1000 iterations)	71.8	53.9

TABLE II. PERFORMANCE RESULTS FOR THE REALISTIC NETWORK

Load	Algorithm	Average spectrum utilization (#subcarriers)	Average running time (sec)
D=4	RSA ILP	-	-
	R+SA ILP	27.2	7200 (*)
	MSF+heuristic	31.7	0.5
	LPF+heuristic	31.1	0.5
	SA (1000 iterations)	29.2	40.8
	SA (10000 iterations)	28.7	424.5
D=40	RSA ILP	-	-
	R+SA ILP	266	7200 (*)
	MSF+heuristic	302	0.6
	LPF+heuristic	298	0.6
	SA (1000 iterations)	285	51.4
	SA (10000 iterations)	273	546

ACKNOWLEDGMENTS

This work has been partially supported by the European Commission through the DICONET project [9].

REFERENCES

[1] A. J. Lowery, J. Armstrong, "Orthogonal frequency division multiplexing for dispersion compensation of long-haul optical systems", Opt. Express, vol. 14, no. 6, pp. 2079–2084, Mar. 2006.

[2] W. Shieh, X. Yi, Y. Tang, "Transmission experiment of multi-gigabit coherent optical OFDMsystems over 1000 km SSMF fiber", Electron. Lett., vol. 43, pp. 183–185, 2007.

[3] A. Lowery, L. B. Du, J. Armstrong, "Performance of Optical OFDM in Ultralong-Haul WDM Lightwave Systems", IEEE/OSA Journal of Lightwave Technology, Vol 25, No. 1, pp. 131-138, 2007.

[4] M. Jinno, H. Takara, B. Kozicki, Y. Tsukishima, Y. Sone, S. Matsuoka, "Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies" IEEE Communications Magazine, Vol. 47 , Is. 11, 2009.

[5] 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", European Conference on Optical Communication (ECOC), 2008.

[6] B. Kozicki, H. Takara, T. Yoshimatsu, K. Yonenaga, M. Jinno, "Filtering Characteristics of Highly-Spectrum Efficient Spectrum-Sliced Elastic Optical Path (SLICE) Network", Optical Fiber Communication, 2009.

[7] K. Christodoulopoulos, K. Manousakis, E. Varvarigos, "Offline Routing and Wavelength Assignment in Transparent WDM Networks", accepted for publication in IEEE/ACM Transactions on Networking.

[8] C. Papadimitriou, K. Steiglitz, "Combinatorial Optimization: Algorithms and Complexity", Dover publications, 1998.

[9] LINDO API: http://www.lindo.com/

[10] Dynamic Impairment Constraint Network for Transparent Mesh Optical Networks (DICONET), FP7 - GA 216338. http://www.diconet.eu/