

Elastic Bandwidth Allocation in Flexible OFDM-based Optical Networks

K. Christodoulopoulos, I. Tomkos, E. A. Varvarigos

Abstract—Orthogonal Frequency Division Multiplexing (OFDM) has recently been proposed as a modulation technique for optical networks, because of its good spectral efficiency, flexibility, and tolerance to impairments. We consider the planning problem of an OFDM 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 and choosing an appropriate modulation level, taking into account the transmission distance. We introduce the Routing, Modulation Level and Spectrum Allocation (RMLSA) problem, as opposed to the typical Routing and Wavelength Assignment (RWA) problem of traditional WDM networks, prove that is also NP-complete and present various algorithms to solve it. We start by presenting an optimal ILP RMLSA algorithm that minimizes the spectrum used to serve the traffic matrix, and also present a decomposition method that breaks RMLSA into its two subproblem subproblems, namely, (i) routing and modulation level, and (ii) spectrum allocation (RML+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 the connections in the traffic matrix. In the sequential algorithm, we investigate two policies for defining the order in which connections are considered. We also use a simulated annealing meta-heuristic to obtain even better orderings. We examine the performance of the proposed algorithms through simulation experiments and evaluate the spectrum utilization benefits that can be obtained by utilizing OFDM elastic bandwidth allocation, when compared to a traditional WDM network.

Index Terms— *Optical OFDM, elastic / flexible spectrum optical paths, planning (offline) problem, Routing Modulation Level and Spectrum Allocation, spectrum continuity constraints, non-overlapping spectrum constraints.*

I. INTRODUCTION

The continuous growth of 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-effective and scalable networking infrastructure. To meet the increasing capacity requirements, recent innovations in optical communication systems, including advanced modulation formats and digital equalization in the electronic domain, have enabled per-channel bandwidths of 40 and 100 Gbps with improved transmission distance in traditional fixed-grid single carrier WDM networks. The high channel capacity and the extended optical reach enable high rate transmission over multiple WDM links and optical cross-connects (OXC) without optical-electrical-optical (OEO) regeneration. Thus, wavelength

routed transparent networks seem to offer a cost-effective solution for high capacity transport networks [1].

Although wavelength routed WDM networks offer well-known advantages, they still exhibit a major 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 cost and energy efficiency and scalability requires a flexible network that would have a fine granularity so as to adaptively provide the required capacity to sub- or super-wavelength demands. Approaches such as optical burst switching (OBS) and optical packet switching (OPS) that meet these requirements can only be viewed as long-term solutions since their enabling technologies are not yet mature [2][3].

Recently, Orthogonal Frequency-Division Multiplexing (OFDM) has been proposed as a modulation technique in optical networks [4]-[7]. Optical OFDM distributes the data on several low data rate subcarriers (multi-carrier system). The spectrum of adjacent subcarriers can overlap, since they are orthogonally modulated, increasing the transmission spectral efficiency. Moreover, optical OFDM can provide fine-granularity capacity to connections by the elastic allocation of low rate subcarriers. A bandwidth-variable (BV) OFDM transponder generates an optical signal using just enough spectral resources, in terms of subcarriers with appropriate modulation level, to serve the client demand. Since, typically, the OFDM signal is generated at the RF domain, many transmission properties can be determined, enabling the choice of the number of modulated bits per symbol of the subcarriers. To establish a connection, every bandwidth variable (BV) OXC on the route allocates a cross-connection with sufficient spectrum to create an appropriately sized end-to-end optical path. Enabling technologies and sub-systems, such as bandwidth-variable transponders and bandwidth-variable OXCs, have been demonstrated in the Spectrum-sLICed Elastic optical path network (“SLICE”) [8]-[11].

The use of optical OFDM as a bandwidth-variable and highly spectrum-efficient modulation format can provide scalable and flexible sub- and super-wavelength granularity, in contrast to the conventional, fixed-grid, rigid-bandwidth WDM network. However, this new concept poses additional challenges on the networking level, since the routing and wavelength assignment (RWA) algorithms of traditional WDM networks are no longer directly applicable. A connection requiring capacity larger than that of an OFDM subcarrier has to be assigned a number of contiguous subcarrier slots for increased spectral efficiency (remember that OFDM uses overlapping orthogonally modulated adjacent subcarriers). In this context, the wavelength

K. Christodoulopoulos and E. Varvarigos are with Computer Engineering and Informatics Department, University of Patras, and with Research Academic Computer Technology Institute, Patra, Greece, (e-mail: kchristodou, manos@ceid.upatras.gr).

I. Tomkos is with Athens Information Technology (AIT) Center, Peania, Greece. (e-mail: itom@ait.edu.gr)

continuity constraint of traditional WDM networks is transformed to a spectrum continuity constraint, and the single wavelength assignment constraint is transformed to a non-overlapping spectrum constraint. Also, note that in OFDM many properties of the transmitted signal are determined in the electrical domain and can be managed by software. A feature that is particularly important for further increasing the flexibility and efficiency of an OFDM network is the choice of the number of modulated bits per symbol for each subcarrier (or for the set of subcarriers corresponding to a connection). To address these decision issues, new Routing, Modulation Level and Spectrum Allocation (RMLSA) algorithms as well as appropriate extensions to network control and management protocols have to be developed.

The problem of planning a flexible OFDM-based optical network has only recently received some attention. The spectrum allocation problem in an OFDM-based core network, in a slightly different setting than the one considered here, has been examined in [13]. In particular, the authors in [13] use shortest path routing and do not account for the requirement of contiguous spectrum allocation for the OFDM subcarriers. An OFDM-based access network (OFDMA) and OFDMA sub-wavelength spectrum assignment to form fixed-grid WDM wavelengths are presented in [14]. The planning of an opaque point-to-point OFDM-based network with adaptive modulation levels based on transmission distance restrictions has been examined in [15]. A comparison of the number of transponders required to design a bandwidth flexible OFDM network using fixed 50 GHz spaced grid to that of a traditional 50 GHz fixed-grid and rigid bandwidth WDM network is presented in [16]. Although not yet studied in depth, the flexibility of OFDM is also expected to offer significant benefits in a dynamic network environment with time-varying connection rates as well as in dynamic restoration scenarios.

In this paper we focus on the routing, modulation level and spectrum allocation (RMLSA) problem in OFDM-based elastic optical networks. We consider the planning phase (offline problem) of such a network, where we are given a traffic matrix with the requested transmission rates of all connections. Our objective is to serve the connections and minimize the utilized spectrum under the constraint that no spectrum overlapping is allowed among these connections. To serve the connections efficiently, the proposed algorithms exploit the two degrees of flexibility provided by OFDM, namely, the elastic spectrum allocation and the modulation level adaptation. Note that the proposed algorithms can function in a restricted flexibility network setting, where only one of the two flexibility degrees is provided. For example, the proposed algorithms can be used to plan a fixed spectrum OFDM network using standard fixed-grid OXCs, without employing the more exotic BV OXCs, similar to the study presented in [16].

We formulate the RMLSA problem and prove that it is NP-complete. We then present various algorithms to solve it. We initially present an optimal RMLSA integer linear programming (ILP) formulation. To reduce the size of the RMLSA problem we also present a formulation that decomposes it into its substituent sub-problems, namely (i) routing and modulation level, and (ii) spectrum allocation, which are solved sequentially (RML+SA), without, however, being able to guarantee an optimal solution for the joint RMLSA problem. Since these formulations cannot be solved efficiently for large networks, we

present a heuristic algorithm that solves the planning RMLSA problem by sequentially serving one-by-one the connections. The ordering plays an important role in the performance of the 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.

We compare the performance of the proposed RMLSA algorithms through simulations, using as performance metrics the utilized spectrum and the algorithms' running times. We also evaluate the spectrum utilization benefits that can be obtained by the flexible utilization of bandwidth enabled by OFDM, when compared to a typical, fixed-grid WDM network. Initially, we consider two restricted OFDM network settings, by constraining one at a time the two degrees of flexibility available. In particular, we examine the cases where the OFDM network (i) uses a specific modulation format but can elastically allocate the spectrum, or (ii) uses a fixed-grid spectrum but can adapt the modulation level of the connections. Finally, we compare the fully flexible OFDM network to a mixed-line-rate (MLR) fixed-grid WDM network.

In our previous works [20][21] we introduced the problem of routing and spectrum allocation (RSA) in an elastic OFDM optical network. In this paper we extend our previous work by proposing algorithms that can, apart from allocating routes and spectrum, also choose the modulation level of the connections. We provide a more complete analysis of the problem and sketch its NP-complete proof. We also present a more extensive study for the newly and previously proposed features of the algorithms, under realistic network and traffic parameters.

The rest of the paper is organized as follows. In Section II we provide an overview of an OFDM optical network that provides elastic bandwidth allocation. In Section III we introduce the routing, modulation level, and spectrum allocation problem and propose various algorithms to solve it, ranging from optimum ILP formulations to heuristic algorithms. Our performance results follow in Section IV, where we initially examine the performance of the proposed RMLSA algorithms and then compare the OFDM network performance to that of a typical WDM network. Finally, Section V concludes the paper.

II. OFDM-BASED OPTICAL NETWORK

In this section we shortly present the basic elements of the elastic OFDM optical network envisaged in our study.

In OFDM, 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, research efforts have focused on an optical version of OFDM as a means to overcome transmission impairments [4]-[7]. In addition to the advantages that stem from the low symbol rate of each subcarrier and the coherent detection that both help to mitigate the effects of physical impairments, OFDM also brings unique benefits in terms of spectral efficiency, by allowing the spectrum of adjacent subcarriers to overlap thanks to their orthogonal modulation. Moreover, OFDM enables elastic bandwidth transmission by allocating a variable number of low-rate subcarriers to a transmission. Fig. 1 presents an example of elastic bandwidth provisioning in the spectrum domain by controlling the number of used subcarriers [8] [9].

OFDM provides an additional degree of flexibility as described next. Optical OFDM signal can be generated either

electronically or optically. In the electrical approach, data are mapped onto individual subcarrier symbols and converted to the baseband OFDM signal through inverse fast Fourier transformation (IFFT) and digital-to-analog (DAG) conversion. Then the baseband OFDM signal is upconverted to the optical domain using a wavelength-tunable laser and an optical I/Q modulator [4]-[7]. In the optical approach, the OFDM signal is directly generated through modulating individual optical subcarriers that are then coupled together [9]. Both cases, and especially the electrical approach that uses Digital Signal Processing (DSP), provide the capability to adapt many properties of the transmitted signal, in contrast to hardware implementations of 10/40/100 Gbps transponders used in traditional WDM networks. In particular, each OFDM subcarrier can be modulated individually using, for example, single bit per symbol binary phase-shift keying (BPSK), QPSK (2 bits per symbol), 8QAM (3 bits per symbol), etc., using the same I-Q modulator for the transmission. In contrast, in traditional WDM networks, changing the modulation level/format requires the use of a different transponder [17].

The choice of the modulation level has to take into account the required Quality of Transmission (QoT) of the connection. A common assumption in optical OFDM, and the one adopted in this study, is that the transmission distance (or equivalently the number of traversed spans for constant spaced spans) of the optical path is the sole QoT factor [11][15][16] of interest. Thus, given the length of an optical path we can find the higher modulation level that can be used over the path (Fig. 2). Therefore, transmissions over shortest optical paths are able to utilize higher modulation levels. Moreover, a connection over a given path can use the same modulation level for all its subcarriers, irrespectively of their number. Note that other parameters related to transmitter/receiver characteristics, interference and non-linear physical layer impairments, can also affect the QoT and thus the modulation level choice. Such issues complicate significantly the problem and are not addressed in this work. The reader is referred to [18] for the problem of planning a typical WDM network under interference-related physical layer impairments.

A related issue is the choice of the spectral and capacity characteristics of the subcarriers. For example, consider a 50 Gbps connection that is served by an OFDM optical path consisting of 10 subcarriers, 5 GHz spaced, using QPSK to transmit 5 Gbps per subcarrier. Assume also that we can use a shorter path that supports 16 QAM modulation format with acceptable QoT, using 10 subcarriers of 2.5 GHz and 5Gbps each, or 5 subcarriers of 5 GHz and 10 Gbps each, or another combination. Note that the factor by which spectrum utilization is reduced (equal to two for the above subcarrier capacity/spectrum choices) can differ for each capacity/spectrum choice. Certain capacity/spectrum choices can affect the QoT and the transmission reach of the connection. We will not focus on the subcarrier capacity/spectrum choices in this study, assuming that these features have already been defined so as to maximize the transmission reach for a particular modulation level and have been included when designing Fig. 2. Moreover, to simplify our analysis, we will assume that a subcarrier always utilizes a constant spectrum F GHz, irrespectively of the modulation format [11][15]. Thus, for constant subcarrier spectrum F , the capacity of the subcarrier is defined by the modulation level choice. The planning algorithms that will be

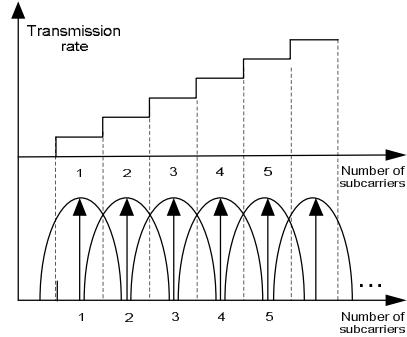


Fig. 1: Variable bandwidth transmission by elastically controlling the number of OFDM subcarriers.

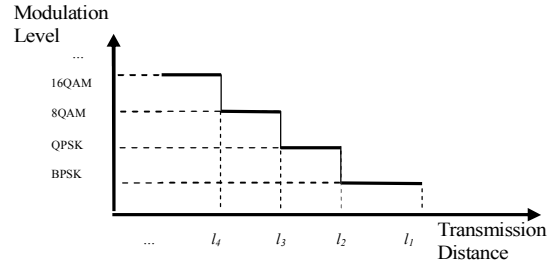


Fig. 2: Higher applicable modulation level with acceptable quality of transmission (QoT) as a function of the transmission distance.

presented are designed for constant spectrum subcarriers but can be extended in a straightforward way for non-constant spectrum subcarriers. A more complete model would use a function of the subcarrier capacity/spacing to constrain the maximum transmission distance. Under the assumption that increasing the subcarrier capacity increases the transmission distance, but, at the same time, the subcarrier granularity deteriorates, there would be a tradeoff in choosing these capacity/spectrum parameters. However, to capture such effects and define an appropriate maximum transmission reach function requires accurate analytical physical layer models, which are not yet available.

Based on the above, a connection requesting a specific rate has two degrees of flexibility, the modulation level and the spectrum. Once these have been determined, the signal transmitted over the optical path is routed through bandwidth variable optical cross-connects (BV OXCs) towards the receiver. In this routing process, only the spectrum domain is essential. Every BV OXC 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 BV OXC has to configure its spectral switching window in a contiguous manner according to the spectral width of the incoming optical signal. MEMS or liquid crystal-based (LCoS) wavelength-selective switches (WSSs) can provide bandwidth variable switching functionality to realize BV OXCs [12]. To avoid interference effects between adjacent optical paths, appropriate spectrum separation, implemented by spectrum guardbands, is required [11]. Fig. 3 presents an example of the routing operation of a spectrum flexible OFDM network.

Fig. 4 presents an example of the utilization of a link in such a spectrum flexible OFDM network. Signals of different optical paths (denoted as “1”, “2” and “3”) are multiplexed in the frequency domain. In this example, we have divided the frequency spectrum into a number of equal frequency slots or subcarriers of F GHz. An optical path utilizes a number of

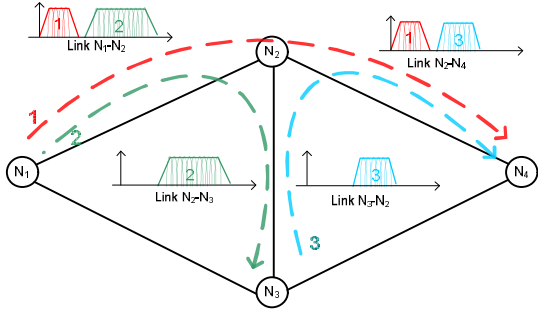


Fig. 3: Spectrum flexible OFDM-base optical network.

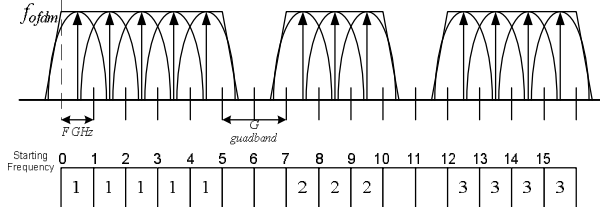


Fig. 4: Spectrum allocation illustrated as a table of subcarrier slots of constant spectrum for the connections using a given link. Spectrum guardbands, each consisting of G subcarriers, separate the path flows so that they can be routed and received with acceptable signal performance.

subcarriers that are allocated using adjacent subcarrier slots. Note that adjacent subcarriers overlap in the spectrum domain, since they are orthogonally modulated, so as to increase spectral efficiency. The elastic data transmission rate using a variable number of low-rate OFDM subcarriers improves the flexibility and granularity of the network when compared to a fixed-grid WDM network. As discussed above, in this study we assume fixed-width spectrum slots, irrespectively of the utilized modulation format, although the proposed algorithms can be straightforwardly extended for non-uniform slot networks.

The employment of OFDM technology does not rule out the usage of WDM, or the opposite. Both technologies can co-exist in a future optical transport network. The employment of OFDM can start gradually in an existing WDM network. As a starting point we envision a fixed-grid WDM system that uses standard fixed-grid OXCs. In such a system, OFDM can be used in the restricted constant-spectrum setting (see the experiments described in Section IV.B.2) and would serve new or increasing-rate connections gradually, replacing some of the WDM transponders. In the case that, instead of fixed-grid OXCs, the network employs bandwidth variable OXCs that enable spectral flexible transmissions (such a network would fully explore the flexibility of OFDM), WDM transmission could be also accommodated. In particular, in such a flexible optical network, WDM transmissions would require a constant e.g. 50 GHz slice of the spectrum. Thus, WDM and OFDM can co-exist in a future optical transport network, and the employment of the one architecture does not rule out the concurrent use of the other.

A. Transmission rate service guarantees

Although the transmission rate of a connection may 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 translates to the requirement for non-overlapping spectrum allocation to all connections for their requested rates. Although planning a network in this way may result in some waste of resources, when the connections under-

utilize their provisioned bandwidth, there are still major gains that can be obtained over the traditional WDM networks. These gains include (i) the high spectrum efficiency because of the orthogonally modulated overlapping subcarriers, (ii) the fine granularity at the low-rate subcarrier level, (iii) the adaptable modulation level, (iv) impairment tolerance due to OFDM properties, and (v) a possible reduction in power consumption by partially deactivating the transmitters, adjusting them to the rate at a specific time. Note that, at a specific time, unused spectrum could 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 information on the traffic time-variations, or by allowing overlapping spectrum allocation based on stochastic traffic models. For example, connections with transmission rates that are complementary 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. Such issues are not explored in the present paper. Future work includes examining semi-static or dynamic problems in OFDM networks where spectrum overlapping is allowed, based on time and/or probabilistic traffic models, as a way to further improve spectrum utilization efficiency.

B. RMLSA requirements

Routing and wavelength assignment (RWA) algorithms devised for fixed-grid WDM systems are not applicable to OFDM networks, even when the modulation level is fixed. To see that, note that the OFDM routing and spectrum allocation (RSA) problem can be transformed into a typical RWA formulation, by viewing a subcarrier in the RSA problem as a wavelength of equal capacity in the RWA problem. Although a typical RWA algorithm is able to find a route for a connection requiring a number of wavelengths (subcarriers in the RSA context), the wavelengths that will be found by the RWA algorithm are not generally going to be contiguous. Allocating contiguous subcarriers is crucial in OFDM networks, since the spectrum of adjacent subcarriers overlap to enable higher spectral efficiency. It is unclear how most RWA algorithms can be modified so as to select contiguous wavelengths.

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 install additional fibers per link. The high number of OFDM subcarrier (of the order of several hundreds) limits the applicability of traditional RWA algorithms. Finally, in the RMLSA problem, we also have to choose a modulation level per subcarrier and connection. This problem in a slightly different setting has been recently examined in traditional WDM networks, and RWA algorithms for mixed line rate (MLR) systems have been presented in [17].

From the above discussion it is clear that RMLSA requires the development of new algorithms that will (i) serve a connection utilizing a contiguous and elastic spectrum, (ii) formulate the problem using variables and constraints that do not depend on the number of subcarriers, and (iii) enable the choice of the modulation level for each connection.

III. ROUTING, MODULATION LEVEL, AND SPECTRUM ALLOCATION (RMLSA) ALGORITHMS

The OFDM 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 OXCs. E denotes the set of (point-to-point) single-fiber links. Let $N=|V|$ and $L=|E|$ denote the number of network nodes and links, respectively.

The spectral granularity of the transmitters and BV OXCs is one subcarrier, corresponding to F GHz of spectrum. The capacity of a subcarrier depends on the modulation level used. We denote the capacity of a subcarrier using single bit per symbol BPSK by C Gbps, and thus QPSK corresponds to $2C$, 8QAM corresponds to $3C$, and so on. We denote by R the modulation level multiplier, having C as the base capacity, that is $R=1,2,3,\dots$, for BPSK, QPSK, 8QAM, etc.

We assume that we are given a function g that relates the length of a path with the higher modulation level that can be used over that path with acceptable quality of transmission (QoT), as depicted in Fig. 2. Even though the QoT generally depends on several physical effects and parameters, e.g. could depend on the number of utilized subcarriers, in this study we will consider the distance as the sole QoT factor that determines the feasibility of a specific modulation level for the transmission over a path.

To route the optical paths through the OXC a guardband of G subcarriers separates adjacent spectrum paths. We assume that elastic OFDM transmitters can be tuned to utilize any 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 (Fig. 4). 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 Fig. 4, connection ‘2’ that requires 3 subcarriers is assigned starting frequency $f_2=7$, assuming that f_{OFDM} corresponds to zero frequency.

The transmission rates of all requested connections are given in the form of a matrix of non-negative integers Λ , called the traffic matrix, and Λ_{sd} denotes the capacity required for the communication between source s and destination d in Gbps. We assume that for commodity (s,d) we utilize a continuous spectrum, that is, a contiguous set of subcarriers. The number of subcarriers allocated to that connection depends on the modulation level, which depends in turn on the chosen path.

Given the above, we want to serve all connections by identifying the route, the modulation level, and the spectrum to be used by each of them, so as to minimize the total spectrum used in the network. In what follows, we will propose algorithms for solving this offline Routing, Modulation Level, and Spectrum Allocation (RMLSA) problem in flexible OFDM-based optical networks. Since offline RMLSA is an NP-complete problem (see Appendix A for a short proof), it cannot be solved efficiently for large input instances and thus heuristic algorithms have to be applied.

A. Combinatorial RMLSA Algorithms

1) Joint RMLSA Algorithm

We initially present an optimal integer linear programming (ILP) formulation [19] that minimizes the utilized spectrum.

For each commodity $s-d$ we pre-calculate k paths, using a variation of the k -shortest path algorithm: at each step, a

shortest path is selected, and the costs of its links are increased (doubled in our experiments) so as to be avoided by the paths found in subsequent steps. The paths obtained in this way tend to use different edges so that they are more representative of the path solution space. Other k -shortest path algorithms are also applicable. We let P_{sd} be the set of candidate paths for $s-d$ and $P=\bigcup_{(s,d)} P_{sd}$ be the total set of candidate paths.

Given function g that relates the length of a path with the highest modulation level that can be used over that path with acceptable QoT (Fig. 2), each pre-calculated path p is mapped to a capacity multiplier $R_p \in \{1,2,3,\dots\}$, considering the base capacity C of single bit per symbol modulation (BPSK).

The problem could also be formulated without using any set of predefined paths, allowing routing over any feasible path. Such an algorithm would give at least as good solutions as the one that uses pre-calculated paths, but would use a much higher number of variables and constraints and scale worse. 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} . Denoting $F_{total} = \sum_{(s,d)} T_{sd} / C$, we have $0 \leq f_{sd} < F_{total}$.

$\delta_{sd,s'd'}$: 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 RMLSA formulation:

minimize c

subject to the following constraints:

- Cost function

$$c \geq f_{sd} + \sum_{p \in P_{sd}} \left[\frac{\Lambda_{sd}}{R_p} \cdot C \right] \cdot x_p, \text{ for all } (s,d) \text{ pairs} \quad (1)$$

- 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 $p \in P_{sd}$ and $p' \in P_{s'd'}$, with p and p' sharing at least one common link l ($\forall (s,d), (s',d') : \exists p \in P_{sd} \cap \exists p' \in P_{s'd'} \cap (l \in p \cap l \in p')$), the following constraints are employed:

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

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

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

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

- Spectrum continuity and non-overlapping spectrum allocation constraints

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

$$f_{sd} + \left\lceil \frac{\Lambda_{sd}}{R_p} \cdot C \right\rceil \cdot x_p + G - f_{s'd'} \leq (F_{total} + G) \cdot (1 - \delta_{sd,s'd'} + 2 - x_p - x_{p'}), \quad (6)$$

$$f_{s'd'} + \left\lceil \frac{\Lambda_{s'd'}}{R_{p'}} \cdot C \right\rceil \cdot x_{p'} + G - f_{sd} \leq (F_{total} + G) \cdot (1 - \delta_{s'd',sd} + 2 - x_p - x_{p'}). \quad (7)$$

When one (or both) of the paths p and p' is not utilized ($x_p \neq 1$ or $x_{p'} \neq 1$), then we do not have to consider the overlapping of their spectrum. In this case, constraints (6) and (7) are deactivated (i.e., they are satisfied for all values of f_{sd} and $f_{s'd'}$), since the right hand side of the constraints takes a value larger than F_{total} , which is always higher than the left hand side.

When both paths p and p' are utilized ($x_p = 1$ and $x_{p'} = 1$), one of the constraints (6) or (7) is activated according to the values of $\delta_{sd,s'd'}$ and $\delta_{s'd',sd}$. In particular, constraint (6) is activated when $\delta_{sd,s'd'} = 1$ (corresponding to $f_{sd} < f_{s'd'}$), in which case (6) becomes:

$$f_{sd} + \left\lceil \frac{\Lambda_{sd}}{R_p} \cdot C \right\rceil + G \leq f_{s'd'},$$

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

$$f_{s'd'} + \left\lceil \frac{\Lambda_{s'd'}}{R_{p'}} \cdot C \right\rceil - f_{sd} \leq F_{total},$$

which is satisfied for all values of f_{sd} and $f_{s'd'}$. In a similar manner, constraint (7) is activated when $\delta_{s'd',sd} = 1$ (corresponding to $f_{sd} > f_{s'd'}$) and constraint (6) is deactivated.

Thus, constraints (6) and (7) ensure that the portions of the spectrum that are allocated to connections that utilize paths that share 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 . By selecting path p , the algorithm also determines the modulation level to be used by the connection, since each path is mapped by function g to a specific modulation format (in particular, the highest modulation level with acceptable QoT over that path).

The number of variables and constraints used by the above ILP formulation depends on the overlapping of the paths considered (and thus depends on the topology and the value of k used). 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 a pair of corresponding paths that share a common link. In the worst case, the formulation would require N^4 Boolean variables $\delta_{sd,s'd'}$, N^4 equality constraints for (3) and $4N^4$ 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 kN^2 Boolean and N^2 integer variables for x_p and f_{sd} , respectively, and N^2 equality constraints of type (1) and N^2 inequality constraints of type (2).

Note that if we did not use a set of pre-calculated paths for each source-destination pair in our formulation, but instead used a multicommodity flow formulation that would allow the routing of a commodity over any path, 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 the flow constraints. Such a number of variables and constraints would result in considerably larger execution times for the algorithms

and is the main reason we chose to use pre-calculated paths in our formulation. Finally, note 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 (RML+SA)

The algorithm presented in this section breaks the RMLSA problem into (i) the routing and modulation level (RML) 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 RMLSA problem might not be found.

a) Routing and Modulation Level Phase

As in the joint RMLSA algorithm described above, we calculate for each commodity $s-d$ a set of candidate paths P_{sd} , and define the Boolean variables x_p for the candidate paths.

ILP Routing and Modulation Level (RML) 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

$$F_l = \sum_{sd} \sum_{\{p \in P_{sd} \mid l \in p\}} \left\lceil \frac{\Lambda_{sd}}{R_p} \cdot C \right\rceil \cdot x_p + GB, \text{ for all links } l \in E$$

- 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 next phase.

b) Spectrum Allocation Phase

Spectrum allocation is done as in the joint RMLSA 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 calculated in the routing and modulation level phase (previous section). Thus, for each connection (s,d) there is only one available path p_{sd} .

ILP Spectrum Allocation (SA) formulation:

minimize c_{SA}

subject to the following constraints:

- Cost function

$$c_{SA} \geq f_{sd} + \left\lceil \frac{\Lambda_{sd}}{R_{p_{sd}}} \cdot C \right\rceil, \text{ for all } (s,d) \text{ pairs}$$

- Starting frequencies ordering constraints

$$\delta_{sd,s'd'} + \delta_{s'd',sd} = 1,$$

$$f_{s'd'} - f_{sd} < F_{total} \cdot \delta_{sd,s'd'},$$

$$f_{sd} - f_{s'd'} < F_{total} \cdot \delta_{s'd',sd},$$

for all p_{sd} and $p_{s'd'}$ sharing at least one common link l .

- Spectrum continuity and non-overlapping spectrum allocation constraints

$$f_{sd} + \left\lceil \frac{\Lambda_{sd}}{R_{p_{sd}}} \cdot C \right\rceil + G - f_{s'd'} \leq (F_{total} + G) \cdot (1 - \delta_{sd,s'd'}),$$

$$f_{s,d'} + \left\lceil \frac{\Lambda_{s,d'}}{R_{p_{s,d}}} \cdot C \right\rceil + G - f_{sd} \leq (F_{total} + G) (1 - \delta_{s',d',sd}),$$

for all p_{sd} and $p_{s',d'}$ sharing at least one common link l .

B. Sequential establishment of Demands

Since the above ILP formulations in both the RMLSA and RML+SA algorithms cannot be solved efficiently for large networks, it is desirable to obtain efficient heuristic algorithms. The approach to be proposed uses a pre-ordering phase and then a heuristic RMLSA algorithm designed for single demands, to sequentially serve the demands one-by-one.

1) Single Demand RMLSA Heuristic Algorithm

We assume that each link $l \in E$ is characterized by a Boolean subcarrier availability vector

$$\bar{u}_l = [u_{li}] = (u_{l1}, u_{l2}, \dots, u_{ld})$$

of length $d = F_{total}$ equal to the maximum number of subcarriers required in the worst case to serve all demands. Element u_{li} records the availability of the i^{th} subcarrier, and is equal to 1 if the i^{th} subcarrier is available, and equal to 0 if it has already been used by a path flow. We can calculate the subcarrier availability vector of path p , \bar{U}_p , by the availability vectors of the links $l \in p$ that comprise it as follows:

$$\bar{U}_p = [U_{pi}] = \&_{l \in p} \bar{u}_l = [\&_{l \in p} u_{li}], \quad (8)$$

where “&” denotes the Boolean AND operation.

The single demand RMLSA 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) , and associate to path p a capacity multiplier $R_p \in \{1,2,3,\dots\}$. To serve a new connection (s,d) that requires capacity Λ_{sd} , the algorithm takes as input the spectrum availability \bar{u}_l of all links $l \in E$. We first use Eq. (8) to calculate the spectrum availability \bar{U}_p of all candidate paths $p \in P_{sd}$. We then search the spectrum availability vector \bar{U}_p of all paths for the first possible placement of $T_{sd,p} = \lceil \Lambda_{sd} / R_p \cdot C \rceil$ subcarriers (along with the required guardbands). In other words, we search for $2G + T_{sd,p}$ continuous 1's in the spectrum availability vector (or $G + T_{sd,p}$ 1's at the starting and ending limits of the spectrum vector). Void filling is performed so that voids of size greater than $2G + T_{sd,p}$ can be utilized. The algorithm selects the path with the lowest indexed starting subcarrier. After selecting the path and the starting frequency, we update the spectrum availability of the links that comprise the selected path by setting the corresponding slot utilization indicator variables equal to 0.

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 up the procedure, especially in the simulated annealing variation of the algorithm, to be described shortly. Note 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 (to reduce spectrum

fragmentation), or select the path with the smallest number of links so as to use less subcarrier resources, etc..

2) Ordering the demands and Simulated Annealing

The heuristic algorithm described above serves the demands in the traffic matrix, one-by-one, in some particular order. The ordering in which the requested connections are served is quite important in this process, and different orderings result in different spectrum utilizations. Two such ordering policies that we have evaluated are the following:

- *Most Subcarriers First (MSF)* ordering: We order the connection demands according to 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 according to the number of links on their shortest paths, and serve first the demand whose shortest path utilizes the largest number of links.

We also use a simulated annealing (SimAn) meta-heuristic to find good orderings that yield good spectrum allocation solutions. The SimAn meta-heuristic works as follows. We start with the MSF ordering and calculate its cost (viewed as “energy” in the SimAn setting) by serving the connections one-by-one, using the single demand heuristic algorithm described in subsection IV.B.1 (this is the “fitness function”). For a particular ordering $((s_1, d_1), (s_2, d_2), \dots, (s_M, d_M))$ of M demands, we define its neighbor as the ordering where (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. We use this random neighbor creation procedure and the single demand heuristic as the fitness function in a typical simulated annealing iterative procedure.

C. Extensions

The algorithms presented in the previous sections can be extended in several directions. An interesting direction is to examine the gains that can be obtained by allocating non-contiguous spectrum to the connections, allowing them to split into subconnections of contiguous spectrum that are individually served. For example, a connection (s,d) with requested capacity Λ_{sd} can be served as multiple subconnections with equal total requested capacity Λ_{sd} . Each of these subconnections can be treated as a separate connection by the proposed algorithms, and individual routing, modulation level and spectrum allocation choices are made for each of them. Clearly, a tradeoff is involved when splitting a connection. On one hand, there is the flexibility of handling a higher number of smaller capacity subconnections, while on the other hand there are penalties related to (i) the additional spectrum guardbands required to route the subconnections and (ii) the need for additional transponders to transmit them. Also, it would be interesting to examine the cost benefits that can be obtained by packing and serving many connections with a single transmitter. Finally, our algorithm could be extended to allow overlapping spectrum allocation based on time or probabilistic traffic models.

D. Lower and Upper Bounds

Since RMLSA is an NP-complete problem and cannot be optimally solved for large input instances, in this paragraph we comment on how we can find lower and upper bounds on the total spectrum required to serve a problem instance.

A lower bound on the total spectrum required can be obtained using the LP relaxation of the joint RMLSA ILP formulation (Section III.A.1). A relatively tighter bound can be obtained by solving the corresponding multicommodity routing and modulation level (RML) problem as presented in the first phase of the RML+SA algorithm (Section III.A.2.a). Essentially, in the multicommodity RML problem the spectrum continuity and non-overlapping spectrum allocation constraints are relaxed. The multicommodity problem can be solved optimally with modern ILP solvers for quite large networks and heavy traffic, to provide lower bounds for the problem.

Useful bounds can also be obtained by relaxing the contiguous spectrum allocation requirement. In the definition of the RMLSA problem we have assumed that a connection is allocated a set of contiguous subcarriers over a single path. This choice is intuitive since OFDM enables the spectrum overlapping of orthogonally modulated adjacent subcarriers so as to increase spectrum efficiency. By breaking a connection into many lower rate subconnections we get a more flexible network planning problem, but also some significant drawbacks (see the discussion in Section III.C). Assuming that we can totally break each connection to single subcarrier demands and serve every subcarrier individually, and viewing a single subcarrier as one wavelength, the problem is transformed into a RWA problem with mixed-line rates MLR (remember that we also choose the modulation level of each subcarrier). We can then obtain lower bounds, by assuming that each subcarrier (wavelength) is routed without the need of spectrum guardbands, and upper bounds by requiring guardbands for each subcarrier. RWA is also NP-hard, but it is a well-studied problem and a number of efficient heuristics exist for it. A problem with this approach is the high number of subcarriers (mapped to wavelengths) present in an OFDM network, which complicates the corresponding RWA problem.

IV. PERFORMANCE RESULTS

We start, in Section IV.A, by comparing the performance of the proposed RMLSA algorithms in terms of the spectrum utilization they achieve and their running times. In Section IV.B, we evaluate the spectrum utilization benefits that can be obtained through the elastic allocation of bandwidth of the envisioned OFDM-based optical network when compared to a typical fixed-grid WDM network. To do so, we initially consider two restricted OFDM network settings, by constraining the two degrees of flexibility available. In particular, we examine the cases where the OFDM network (i) uses a specific modulation format and can elastically allocate the spectrum, or (ii) uses a fixed-grid and can adapt the modulation level of the connections. Finally, we compare the fully flexible OFDM network to a mixed-line-rate (MLR) fixed-grid WDM network.

We used Matlab to implement the RMLSA algorithms, LINDO [22] for ILP solving, and Matlab's built-in simulated annealing meta-heuristic. For all the algorithms we used $k=3$ candidate paths for each source-destination pair, and guardbands of $G=2$ subcarriers. In addition to the performance results on the RMLSA algorithms, we also report on corresponding lower bounds obtained, namely, (i) the optimal solution of the multicommodity RML problem and (iii) the RWA solution assuming that no spectrum guardbands are used, using the LP relaxation algorithm of [18]. These lower bounds

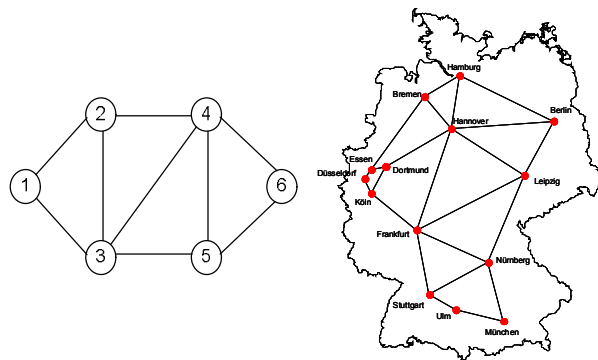


Fig 5: (a) Small network topology, (b) the generic DT network topology, with 14 nodes and 23 undirected links (46 directed links were used in our simulations).

TABLE I. PERFORMANCE RESULTS FOR THE SMALL NETWORK

Load	Algorithm	Average spectrum usage (#subcarriers)	Average running time (sec)
D=4	<i>Lower Bound (RML)</i>	13.8	-
	<i>Lower Bound (RWA)</i>	8.2	-
	RMLSA ILP	13.8	3.25
	RML+SA ILP	13.8	2.72
	MSF+heuristic	14.5	0.17
	LPF+heuristic	14.8	0.17
	SimAn (100 iterations)	13.9	1.94
SimAn (1000 iterations)	13.9	19.1	
D=30	<i>Lower Bound (RML)</i>	59.4	-
	<i>Lower Bound (RWA)</i>	53.7	-
	RMLSA ILP	59.4	503.19
	RML+SA ILP	59.5	6.48
	MSF+heuristic	65.3	0.39
	LPF+heuristic	64.6	0.39
	SimAn (100 iterations)	61.6	3.4
SimAn (1000 iterations)	59.9	41.69	

help us to better understand the performance of the algorithms, by quantifying the effects the various constraints and the involved guardband have on the total spectrum required.

A. Comparison of RMLSA algorithms

In this section we evaluate the performance of the proposed RMLSA algorithms through simulation experiments. The capacity $\Lambda_{s,d}$ required by each connection (s,d) was chosen uniformly between 0 and $D \cdot C$, where C is the base capacity corresponding to a single subcarrier with single bit per symbol modulation level (BPSK), and D is a constant, and in particular $D=4$ and $D=30$ representing low and high load cases. For each traffic instance we measure the spectrum required to serve all connections, by counting the number of subcarriers used, and the corresponding running time of the algorithm. The results were averaged over 50 randomly created matrices. Since we consider the planning phase of the network, that can be performed offline, running times up to a few hours are generally acceptable, but note that our heuristic algorithm executes in a few seconds time for all the examined cases.

The simulations in subsections A.1 and A.2 assume that all connections use the same modulation level, irrespectively of the transmission distance, thus constraining OFDM network flexibility to the spectrum domain. In particular, the multiplier was set to $R_p=1$ for all paths, corresponding to the single bit per symbol BPSK format. Although in this case the RMLSA problem is reduced to the simpler Routing and Spectrum Allocation (RSA) problem, we will still refer to the problem and the algorithms as RMLSA. Then, in subsection A.3, we evaluate

the performance of the full algorithms that, apart from the flexible allocation of spectrum, are also able to choose the modulation level for each connection.

1) Small network experiments

Table I presents the results obtained for the small network illustrated in Fig. 5a. For both light and heavy traffic ($D=4$ and $D=30$) the optimal RMLSA ILP algorithm was able to find solutions in acceptable time. For this small network and both traffic cases, the lower bounds obtained by the multicommodity RML problem were equal to the optimal values found by the joint RMLSA ILP algorithm. This implies that in this small network the spectrum continuity and non-overlapping spectrum allocation constraints are not significant and do not require additional spectrum to be satisfied. From the lower bounds obtained by the RWA heuristic algorithm for light traffic load we can observe the additive effect of the guardbands. The contiguous spectrum allocation plays a more significant role at heavy load than at light load, and increases the gap between the RMLSA solutions and the RWA lower bounds.

Decomposing the problem (RML+SA) reduces the average running time, when compared to the joint RMLSA, while optimality is lost only in one traffic instance of heavy load ($D=30$). Regarding the sequential heuristic algorithm, MSF ordering slightly outperforms LPF. LPF ordering takes into account the paths, which do not play a significant role in this small network. Simulated annealing (SimAn) improves vastly the performance of the sequential algorithm. Especially for light load ($D=4$), the ordering found by simulated annealing resulted in the optimal spectrum utilization in almost all traffic instances. The results presented in Table I for simulated annealing assume 100 and 1000 iterations, since a higher number of iterations showed only marginal performance improvements, while the running time exceeded that of the optimal RMLSA ILP.

2) Realistic network experiments

Table II presents the results obtained for a realistic network, namely the generic Deutsche Telekom (DT) topology consisting of 14 nodes and 46 directed links (Fig. 5b). The optimal RMLSA ILP algorithm was unable to produce good results in reasonable time. Moreover, the results for the decomposed problem (RML+SA ILP) were recorded by stopping the ILP after 2 hours for each traffic instance. In particular, the routing and modulation level (RML) problem was solved optimally quite easily (these results are also reported as RML lower bounds), while the spectrum allocation (SA) problem was quite hard, and we were not able to obtain optimal SA solutions in 2 hours time for the majority of inputs. Note that solving similar problems with RWA algorithms would require the use of a few hundred wavelengths, which are also intractable problems when using optimal ILP algorithms.

From table II we see that the decomposed RML+SA ILP algorithm finds the best solutions (although not the optimal ones within 2 hours). Comparing the lower bound values obtained by RML with the results of the RML+SA algorithm, we observe that the effects of the spectrum allocation (SA) phase are not significant, but increase as the load increases. Satisfying the corresponding SA constraints increases slightly the average number of utilized subcarriers at light load and slightly more at heavy load. As previously, using the RWA lower bounds, we can see the effect of the guardbands at low

TABLE II. PERFORMANCE RESULTS FOR THE REALISTIC DT NETWORK

Load	Algorithm	Average spectrum utilization (#subcarriers)	Average running time (sec)
$D=4$	<i>Lower Bound (RML)</i>	53.7	-
	<i>Lower Bound (RWA)</i>	30.5	-
	RMLSA ILP	-	-
	RML+SA ILP	56.3	7200 (*)
	MSF+heuristic	60.4	1.2
	LPF+heuristic	60.1	1.2
	SA (1000 iterations)	57.0	74.1
	SA (10000 iterations)	56.3	1063.26
$D=30$	<i>Lower Bound (RML)</i>	241.9	-
	<i>Lower Bound (RWA)</i>	210.9	-
	RMLSA ILP	-	-
	RML+SA ILP	252.7	7200 (*)
	MSF+heuristic	264.6	4.5
	LPF+heuristic	264.5	4.5
	SA (1000 iterations)	262.5	207.3
	SA (10000 iterations)	253.0	1168.4

TABLE III. PERFORMANCE RESULTS FOR THE REALISTIC DT NETWORK AND ADAPTIVE MODULATION LEVEL

Load	Algorithm	Average spectrum utilization (#subcarriers)	Average running time (sec)
$D=4$	<i>Lower Bound (RML)</i>	38.2	-
	<i>Lower Bound (RWA)</i>	14.9	-
	RMLSA ILP	-	-
	RML+SA ILP	41.2	7200 (*)
	MSF+heuristic	46.8	1.26
	LPF+heuristic	45.3	1.26
	SA (1000 iterations)	42.4	192
	SA (10000 iterations)	41.4	748.32
$D=30$	<i>Lower Bound (RML)</i>	101.9	-
	<i>Lower Bound (RWA)</i>	75.9	-
	RMLSA ILP	-	-
	RML+SA ILP	127.4	7200 (*)
	MSF+heuristic	145.0	1.82
	LPF+heuristic	143.0	1.82
	SA (1000 iterations)	134.6	299.7
	SA (10000 iterations)	130.30	1380.8

traffic load and the effect of the contiguous spectrum allocation that plays a more significant role as the load increases.

Regarding the sequential heuristic algorithm, LPF slightly outperforms MSF ordering, especially for light load, in contrast to the results of the previous section. In this topology the path lengths differ among the connections. Since connections traversing longer paths utilize more spectrum resources, performance can be improved by serving these connections first, in an empty network, as LPF ordering does. Again simulated annealing (SimAn) improved the performance of the sequential algorithm, producing results close to those of the RML+SA ILP algorithm, especially for light load. For heavy loads, where the variation of requested rates is higher, the ordering in which the connections are considered has a greater effect on the performance of the sequential algorithm, and thus increasing the number of SimAn iterations yields better results. The results obtained for this realistic network by the SimAn algorithm used 1000 and 10000 iterations (as opposed to 100 and 1000 in the previous section), so that the running times are of the same order with those of the ILP.

3) Experiments with variable modulation level formats

In the previous subsections we have assumed that all connections use OFDM with basic single-bit per symbol modulation format and the only flexibility degree was the spectrum domain. In this subsection we examine the

performance of the proposed algorithms when the network supports adaptable modulation levels. That is, we introduce the second degree of flexibility, allowing the adaptation of the modulation level of a connection based on the transmission distance over the chosen path. In these experiments we used the DT network of Fig. 5b with realistic distances between nodes (see deliverable D2.1 in www.diconet.eu/deliverables.asp) and assumed that the BPSK, QPSK, 8QAM, 16QAM formats can be used for transmissions up to 3000km, 1500 km, 750 km, and 375 km, respectively, using the half distance law of [15].

Table III presents the results for light ($D=4$) and heavy ($D=30$) load. As in the experiments with the DT topology of the previous subsection, we do not present results for the optimal RMLSA ILP algorithm, but for RML+SA ILP that was stopped after running for 2 hours per instance. Comparing the results of Table II and Table III the improvements obtained by adapting the modulation levels of the connections become evident. The improvements are more significant under heavy traffic load, in which case the algorithms with adaptive modulation levels are able to find solutions using less than half the number of subcarriers required with single-bit per symbol modulation. Again the RML+SA ILP algorithm gives the best solutions, while the sequential heuristic that uses simulated annealing (SimAn) offers a good and scalable alternative choice. Using adaptive modulation levels seems to increase the effects of the spectrum allocation and the guardbands, since the difference between the best RML+SA values and the lower bounds has increased, compared to the results reported in Table II.

B. Comparing a flexible OFDM with a typical WDM network

In this subsection we evaluate the spectrum gains that can be obtained by employing a flexible OFDM network as opposed to a traditional fixed-grid WDM network. For the OFDM network we used the simulated annealing meta-heuristic with 1000 iterations. We also assume that a subcarrier utilizes $F=5$ GHz spectrum, single bit per symbol modulation (BPSK) capacity per subcarrier is $C=2.5$ Gbps, and the required guardband $G=2$ subcarriers. When OFDM uses adaptive modulation levels we assume, as in Section IV.A.3, that BPSK can be used for transmissions up to 3000km, QPSK can be used up to 1500 km, etc. For the traditional WDM network case we have used the RWA LP-relaxation algorithm presented in [18]. In the RWA formulation we do not consider interference physical impairments, but in subsection 3 in which we utilize a mixed line rate (MLR) WDM system we use transmission reach constraints. We again used the generic DT topology of Fig. 5b and the realistic traffic matrix for this network provided for 2009, as reported in the DICONET project (deliverable D2.1 in www.diconet.eu/deliverables.asp). In this matrix, the average rate between nodes is around 15 Gbps. We uniformly scaled this realistic traffic matrix so as to obtain traffic matrices up to 8 times larger than the reference matrix of 2009.

1) Experiments with fixed modulation level formats

In this set of experiments for the OFDM network we assume that all connections use the same modulation level, irrespectively of the transmission distance, and the only flexibility degree is the elastic spectrum allocation.

For the WDM network, we assumed wavelengths of fixed capacity equal to (a) 40 or (b) 100 Gbps, in a fixed 50 GHz ITU grid, so as to have 0.8 and 2 bit/s/Hz spectral efficiency per wavelength, respectively. The granularity of the OFDM-based network is 5 GHz per subcarrier each of (a) 5 Gbps (using

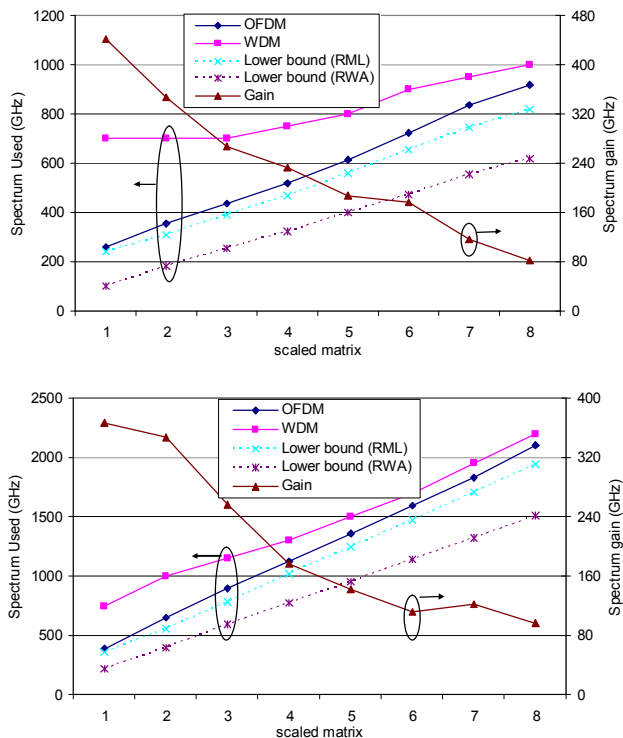


Fig. 6: Spectrum used in a WDM and a spectrum flexible OFDM network with spectral efficiency per WDM 50 GHz wavelength of (a) 0.8 and (b) 2 bit/s/Hz.

QPSK) and (b) 12.5 Gbps capacity (using 32QAM), respectively. Assuming that 10 (8 plus 2 guardband) subcarriers are fitted in 50 GHz, OFDM has the same spectral efficiency per WDM wavelength: (a) 0.8 and (b) 2 bit/s/Hz, respectively. These values were chosen the same for both networks so as to have a fair comparison and evaluate the gains that can be obtained by the spectrum flexibility of the OFDM architecture. However, OFDM with overlapping orthogonally modulated subcarriers is expected to have higher spectral efficiency per wavelength than a typical WDM network.

In Fig. 6a we present the results for case (a), in which the WDM network uses 40 Gbps wavelengths and the OFDM network 5 Gbps subcarriers. The OFDM network is observed to have much better spectral utilization than the WDM network, even though the improvement decreases as the load increases. For the realistic load of 2009 (value of the scale factor equal to 1) the improvement is more than 350 GHz in spectrum that is reduced to 100 GHz for the matrix scaled 8 times. This was expected, since for heavy loads (in the matrix scaled by a factor of 8, the average rate between nodes is 120 Gbps) the finer granularity of the OFDM network (5 Gbps as opposed to 40 Gbps) is not that important, and the elastic allocation of the spectrum utilized by OFDM does not yield significant gains. In Fig. 6b we present the performance results for case (b), in which the WDM network uses 100 Gbps wavelengths and the OFDM network 12.5 Gbps subcarriers. In this case, the OFDM network exhibits even higher performance gains over the WDM network, which remain significant for higher loads, due to the more rigid granularity of 100 Gbps wavelengths.

In both cases, the lower bounds provided by the RML multicommodity problem are quite tight and close to the spectrum utilization of the OFDM network, but as the load increases the spectrum allocation phase becomes more significant and the bounds become less tight. The bounds calculated by RWA heuristic are looser. The difference

between the OFDM performance and RWA bounds increases as the load increases. This difference includes two factors. The first factor is an almost constant effect of the guardbands, and the second is the reduced flexibility of the OFDM network due to the contiguous spectrum allocation that deteriorates the performance of the OFDM network as the load increases.

2) Experiments in a fixed spectrum grid network

Fig. 7 presents the results for the case where the WDM network uses 40 Gbps wavelengths with e.g. QPSK in a 50 GHz grid (Fig. 6a), and the OFDM network uses adaptive modulation levels and a fixed 50 GHz spectrum grid. Note that when QPSK is used in both the OFDM and the WDM network, both networks have equal spectrum efficiency per 50 GHz WDM wavelength. Since here we assume a fixed-grid OFDM network, all OFDM connections are broken to subconnections that fit in a 50 GHz grid. In particular, we assume that a connection utilizes a variable number of subconnections, each using 8 subcarriers (plus 2 guardbands) that form traditional 50 GHz “wavelengths”. The number of subconnections (wavelengths) depends on the chosen modulation level (while the last wavelength can be formed even with less than 8 subcarriers). This OFDM network setting is similar to [16].

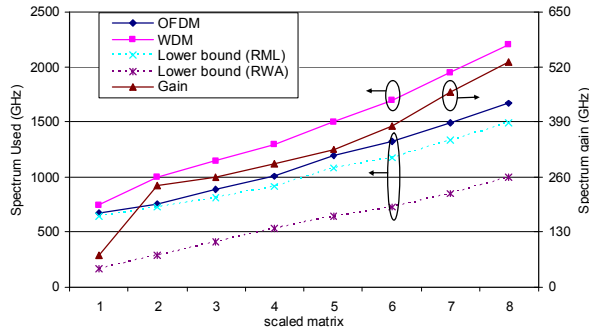


Fig. 7: Spectrum utilization of WDM with 40 Gbps wavelengths vs. an adaptive modulation level OFDM network using fixed 50 GHz spectrum grid.

Contrary to the results presented in Fig. 6, in these experiments the spectrum gains obtained when using the adaptive modulation level OFDM network increase as the traffic load increases. High-rate connections are served over shorter paths using higher modulation levels. As the load and the required transmission rates increase, the use of higher modulation levels, which can be adaptive chosen in OFDM, becomes more efficient.

3) Experiments in a fully flexible OFDM network

In these experiments we evaluate the spectrum gains that can be obtained by a fully flexible OFDM network. We assume that the WDM network uses 40 Gbps wavelengths with QPSK in a 50 GHz grid (as in Fig. 6a). Thus, when QPSK is used in both OFDM and WDM networks, they both have equal spectrum efficiency per 50 GHz WDM wavelength. As mentioned, OFDM provides two flexibility degrees: (i) the elastic spectrum allocation (evaluated alone in Fig. 6a) and (ii) the adaptive modulation level (evaluated alone in Fig. 7). At light loads, the elastic spectrum allocation is quite efficient and thus is the dominant factor of the spectrum gain. Although the spectrum gain of the elastic spectrum allocation decreases as the load increases (Fig. 6a), this is compensated by the improving effects of the adaptive modulation level that become more dominant at heavy load (Fig. 7). In particular, it seems that the

elastic spectrum allocation mechanism complements the adaptive modulation level mechanism. The spectrum gains are much higher than the bare summation of the gains reported by each separate mechanism, an indicative value if these two mechanisms functioned separately.

In Fig. 9 we compare the OFDM with a mixed-line rate (MLR) WDM system. For the WDM network we assumed that 10, 40, 100 Gbps wavelengths are able to transmit with acceptable quality up to 3000, 1500, and 500 km, respectively. The modulation level restrictions of the OFDM network were as previously, resulting in almost the same spectral efficiency per wavelength as the WDM network for the same transmission distance. Comparing the results of Fig. 8 and Fig. 9 we see the improvements obtained in the WDM network when utilizing mixed line rates (the performance of OFDM is the same in both figures). Still, the performance of OFDM with flexible spectrum allocation and adaptive modulation levels is superior and the spectrum improvements are maintained even for heavy traffic loads. Note that adapting the modulation level in the OFDM network can be performed by software, while in the MLR WDM it requires the utilization of different transponders, constraining the adaptability of the network to traffic changes.

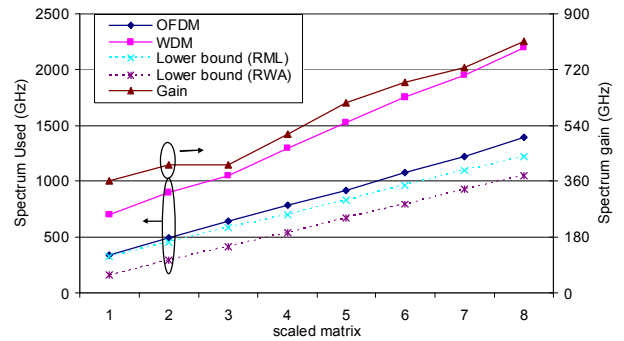


Fig. 8: Spectrum utilization of a WDM network with spectral efficiency per wavelength (50 GHz) of 0.8 bit/s/Hz and an OFDM-based network with flexible spectrum allocation and adaptive modulation levels.

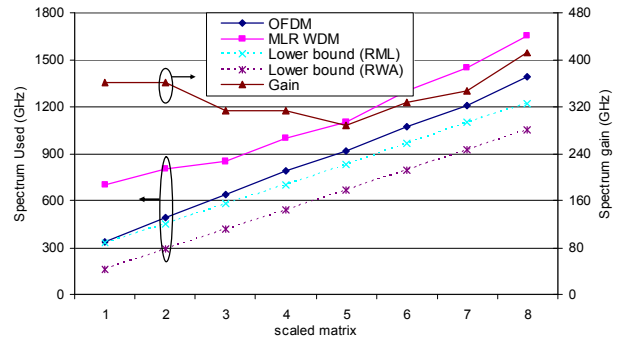


Fig. 9: Spectrum utilization of a mixed-line rate (MLR) WDM network with wavelength capacity of 10, 40 and 100 Gbps and an OFDM-based network with flexible spectrum allocation and adaptive modulation levels.

An OFDM network has higher spectral efficiency per wavelength than a WDM network, because of the overlap of adjacent orthogonally modulated subcarriers. Moreover, due to low subcarrier rates, the reach of the OFDM network is expected to be higher than that of a WDM network. These attributes were not used in the above experiments, (Fig. 7 to Fig. 9), but instead we assumed similar spectral and reach characteristics for both OFDM and WDM networks in order to evaluate the gains that can be obtained just by the flexibility of the OFDM network architecture. When the higher spectral

efficiency and the higher reach of OFDM are taken into account, in addition to the flexibility benefits as presented here, the performance advantages of an OFDM-based network would be even more pronounced. Moreover, it is natural to expect that the flexibility of the proposed architecture would provide even more significant performance gains in a semi- or fully-dynamic traffic scenario. It is in our future plans to examine the performance of the flexible OFDM network under more dynamic scenarios, where the full potentials of the flexible OFDM network would be explored. In any case, the gains presented here, even for static traffic, are substantial and support the applicability of the proposed OFDM architecture as a candidate solution for future transport networks.

V. CONCLUSIONS

Optical OFDM is receiving recent attention as a spectral-efficient modulation format that can provide elastic bandwidth transmission. We considered the problem of planning an OFDM-based optical network where connections are provisioned based on their requested transmission rate (given in a traffic matrix) and assuming no spectrum overlapping between them. We introduced the Routing, Modulation Level, and Spectrum Allocation (RMLSA) problem and presented various algorithms to solve it, ranging from optimal and decomposition ILP algorithms to a sequential heuristic algorithm combined with appropriate ordering policies and a simulated annealing meta-heuristic. Our results indicate that the proposed sequential heuristic with an appropriate ordering discipline can give close to optimal solutions in low running times. Our results show that the OFDM-based networks have significant spectrum benefits over typical fixed-grid WDM networks, indicating that the OFDM architecture offers a promising solution for future high capacity transport networks.

ACKNOWLEDGMENTS

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

APPENDIX – RMLSA COMPLEXITY

In this appendix we sketch the NP-complete proof of the RMLSA problem, by reducing the multiprocessor scheduling problem (MSP) to a related RMLSA problem.

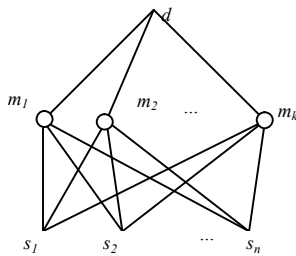


Fig. 10: RMLSA network to solve the MSP problem.

The multiprocessor scheduling problem (MSP) is a known NP-hard problem. An instance of the MSP problem includes a set of tasks $t_i \in T$, $i=1,2,\dots,n$, a set of processors $m_j \in M$, $j=1,2,\dots,k$, and the duration $w(t_i, m_j) \in Z^+$ of each task t_i when executed on processor m_j . The solution of the MSP problem is a “schedule” for T , that is, a function $f: T \rightarrow M$ that maps the tasks to processors. The objective is to minimize the makespan,

that is, the time between the start and the finish of the execution of all the tasks. To solve the MSP as a RMLSA problem we define the following network (Fig. 10).

We assume that we are given a set of connections T . Connection t_i has source node s_i and destination d . Moreover, connection t_i requires capacity Λ_{t_i} that is translated to requiring $w(t_i, m_j)$ subcarriers over path p_{ij} ($s_i \rightarrow m_j \rightarrow d$), assuming that the modulation level of connection t_i over path p_{ij} results in subcarriers of $\Lambda_{t_i}/w(t_i, m_j)$ capacity each. The routing solution of RMLSA problem is a solution to the corresponding MSP. Thus, if we were able to solve the RMLSA optimally, we would be also able to track solutions for MSP. A contradiction, since MSP is NP-complete. Thus, RMLSA is also NP-complete.

REFERENCES

- [1] J. Berthold, A. A. M. Saleh, L. Blair, and J. M. Simmons, “Optical networking: Past, present, and future,” *IEEE/OSA J. Lightw. Technol.*, vol. 26, no. 9, pp. 1104–1118, May 2008.
- [2] S. Yoo, “Optical Packet and Burst Switching Technologies for the Future Photonic Internet” *IEEE/OSA J. of Lightw. Technol.*, 24(12), 2006.
- [3] Y. Chen, C. Qiao, X. Yu, “Optical burst switching: A new area in optical networking research,” *IEEE Network*, 18 (3), pp. 16–23, 2004.
- [4] Q. Yang, Y. Ma, W. Shieh, “107 Gb/s Coherent Optical OFDM Reception Using Orthogonal Band Multiplexing”, NFOEC, PDP7, 2008.
- [5] W. Shieh, X. Yi, Y. Tang, “Transmission experiment of multi-gigabit coherent optical OFDM systems over 1000 km SSMF fiber”, *Electron. Lett.*, vol. 43, pp. 183–185, 2007.
- [6] A. Lowery, L. B. Du, J. Armstrong, “Performance of Optical OFDM in Ultralong-Haul WDM Lightwave Systems”, *IEEE/OSA J. of Lightw. Technol.*, 25(1), pp. 131–138, 2007.
- [7] J. Armstrong, “OFDM for Optical Communications”, *IEEE/OSA J. Lightw. Technol.*, 27(3), 2009.
- [8] M. Jinno, H. Takara, B. Kozicki, Y. Tsukushima, Y. Sone, S. Matsuoka, “Spectrum-efficient and scalable elastic optical path network: architecture, benefits, and enabling technologies” *IEEE Commun. Mag.*, 47 (11), 2009.
- [9] K. Yonenaga, F. Inuzuka, S. Yamamoto, H. Takara, B. Kozicki, T. Yoshimatsu, A. Takada, M. Jinno, “Bit-rate-flexible all-optical OFDM transceiver using variable multi-carrier source and DQPSK/DPSK mixed multiplexing”, *OFC 2009*.
- [10] B. Kozicki, H. Takara, T. Yoshimatsu, K. Yonenaga, M. Jinno, “Filtering Characteristics of Highly-Spectrum Efficient Spectrum-Sliced Elastic Optical Path (SLICE) Network”, *OFC (JWA43)*, 2009.
- [11] M. Jinno, B. Kozicki, H. Takara, A. Watanabe, Y. Sone, T. Tanaka, A. Hirano, “Distance-Adaptive Spectrum Resource Allocation in Spectrum-Sliced Elastic Optical Path Network”, *IEEE Commun. Mag.*, 48(8), 2010.
- [12] R. Ryf, Y. Su, L. Moller, S. Chandrasekhar, L. Xiang, D. Neilson, C. Giles, “Wavelength blocking filter with flexible data rates and channel spacing”, *IEEE/OSA J. Lightw. Technol.*, 23 (1), 2005.
- [13] W. Zheng, Y. Jin, W. Sun, W. Guo, W. Hu, “On the Spectrum-Efficiency of Bandwidth-Variable Optical OFDM Transport Networks”, *OFC (OWR5)* 2010.
- [14] W. Wei, J. Hu, D. Qian, P.N. Ji, T. Wang, X. Liu, C. Qiao, “PONIARD: A Programmable Optical Networking Infrastructure for Advanced Research and Development of Future Internet”, *IEEE/OSA J. Lightw. Technol.*, 27(3), 2009.
- [15] A. Bocoli, B. Schuster, F. Rambach, M. Kiese, C. Bunge, B. Spinnler, “Reach-Dependent Capacity in Optical Networks Enabled by OFDM”, *OFC (OMQ4)*, 2009.
- [16] A. Klekamp, O. Rival, A. Morea, R. Dischler, F. Buchali, “Transparent WDM network with bitrate tunable optical OFDM transponders”, *OFC (NTuB5)*, 2010.
- [17] A. Nag, M. Tomatore, B. Mukherjee, “Optical Network Design With Mixed Line Rates and Multiple Modulation Formats”, *IEEE/OSA J. Lightw. Technol.*, vol 28, no 4, 2010.
- [18] K. Christodouloupoloulos, K. Manousakis, E. Varvarigos, “Offline Routing and Wavelength Assignment in Transparent WDM Networks”, *IEEE/ACM Trans. on Networking*, 18 (5), 2010.
- [19] C. Papadimitriou, K. Steiglitz, “Combinatorial Optimization: Algorithms and Complexity”, Dover publications, 1998.
- [20] K. Christodouloupoloulos, I. Tomkos, E. Varvarigos, “Spectrally/Bitrate Flexible Optical Network Planning”, *ECOC* 2010.
- [21] K. Christodouloupoloulos, I. Tomkos, E. Varvarigos, “Routing and Spectrum Allocation in OFDM-based Optical Networks with Elastic Bandwidth Allocation”, *IEEE Globecom* 2010.
- [22] LINDO API: <http://www.lindo.com/>

Konstantinos Christodoulopoulos (kchristodou@ceid.upatras.gr) received a Diploma in Electrical and Computer Engineering from the National Technical University of Athens, Greece, in 2002, the M.Sc. degree in Advanced Computing from Imperial College London, U.K., in 2004, and the Ph.D. degree from the Computer Engineering and Informatics Department, University of Patras, Greece, in 2009. His research interests are in the areas of protocols and algorithms for optical networks and grid computing.

Ioannis Tomkos (itom@ait.edu.gr) is with the Athens Information Technology Center (AIT) since Sep 2002. In the past, he was senior scientist (1999 - 2002) at Corning Inc. USA and research fellow (1995 - 1999) at University of Athens, Greece. At AIT he founded and serves as the Head of the "High Speed Networks and Optical Communication (NOC)" Research Group that was/is involved in many EU funded research projects within which Dr. Tomkos is representing AIT as Principal Investigator and has a consortium-wide leading role (e.g. Project Leader of projects DICONET and ACCORDANCE, Technical Manager of EU projects SOFI and TRIUMPH, Chairman of the EU COST 291 project). Dr. Tomkos received in 2006 the title of "Distinguished Lecturer" of IEEE Communications Society for the topic of transparent optical networking. He has co-authored about 300 peer-reviewed archival articles (over 170 IEEE sponsored), including about 100 Journal/Magazine/Book publications. Dr. Tomkos has served as the Chair of the International Optical Networking Technical Committee of IEEE and the Chairman of the IFIP working group on "Photonic Networking". He is currently the Chairman of the OSA Technical Group on Optical Communications and the Chairman of the

IEEE Photonics Society Greek Chapter. He has been General, Technical Program, Subcommittee, Symposium Chair, or/and member of the steering/organizing committees for major conferences (e.g. OFC, ECOC, GlobeCom, ICC, ONDM, BroadNets, etc.) in the area of telecommunications/ networking (about 100 conferences/workshops). He is a member of the Editorial Boards of the Journal of Lightwave Technology, the Journal of Optical Communications and Networking, the IET Journal on Optoelectronics, and the International Journal on Telecommunications Management. He is a Fellow of the IET.

Emmanouel (Manos) Varvarigos (manos@ceid.upatras.gr) received a Diploma in Electrical and Computer Engineering from the National Technical University of Athens in 1988, and the M.S. and Ph.D. degrees in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology in 1990 and 1992, respectively. He has held faculty positions at the University of California, Santa Barbara (1992-1998, as an Assistant and later an Associate Professor) and Delft University of Technology, the Netherlands (1998-2000, as an Associate Professor). In 2000 he became a Professor at the department of Computer Engineering and Informatics at the University of Patras, Greece, where he heads the Communication Networks Lab. He is also the Director of the Network Technologies Sector (NTS) at the Research Academic Computer Technology Institute (RA-CTI), which through its involvement in pioneering research and development projects, has a major role in the development of network technologies and telematic services in Greece. His research activities are in the areas of high-speed networks, protocols and architectures, distributed computation and grid computing.