Algorithmic Challenges in Flexible Optical Networks

E. A. Varvarigos and K. Christodoulopoulos
Computer Engineering and Informatics Department, University of Patras, Greece, and
Computer Technology Institute and Press "Diophantus", Patra, Greece

Abstract— **We discuss some basic algorithmic issues in flexible optical networks, highlighting the algorithmic challenges and differences to fixed-grid WDM networks and outline algorithms for planning and operating flexible optical networks. In the planning problem we assume that we are given the traffic matrix and the transponders' feasible configurations that account for the physical layer, and the objective is to serve the traffic and find a solution that minimizes the maximum spectrum used and the cost (number and type) of transponders. The offline RSA algorithm that we outline serves demands for their requested rates by choosing the route, breaking the transmissions in more than one connection and placing regenerators, if needed, and allocating spectrum to them. We then turn our attention to operating a flexible optical network. We assume that the spectrum allocated to a connection can be adapted so as to follow the time-varying required transmission rate. We outline a framework to orchestrate spectrum sharing and a RSA algorithm to serve the demands so as to minimize the blocking caused by the traffic fluctuations.**

I. INTRODUCTION

Increasing the capacity and improving the efficiency of optical transport networks has been an important research challenge for many years. To cope with the capacity requirement increases of more than 30% per year [1], WDM systems target the employment of higher rate and improved distance transmissions. However, the rigid granularity of WDM systems leads to inefficient capacity usage, a problem expected to become more significant with the deployment of the higher channel rate systems.

A technology that has received a lot of attention lately is *flexible* or *elastic* optical networking [2]-[12]. Flexible optical networks assume the use of tunable transponders and a flexible spectrum grid or *flex-grid*. Flex-grid's granularity is much finer than that of standard WDM systems: the spectrum is divided into spectrum slots of 12.5 GHz that can be combined to create channels that are as wide as needed. Tunable optical transponders, also called bandwidth variable transponders (BVT), have lately been proposed, and can adapt several transmission parameters, such as the modulation format, the spectrum and the transmission rate used [9].

Both in standard fixed-grid WDM and in flexible optical networks establishing connections involves the solution of some sort of resource allocation problem, resources being the transponders, the regenerators, and the spectrum units (wavelengths or spectrum slots) on the links and the cross-connects. In traditional WDM networks, resource allocation corresponds to the Routing and Wavelength Assignment (RWA) problem, used to select the paths and wavelengths to create connections called *lightpaths*. Accounting for physical layer impairments and regeneration placement, traffic grooming, etc, are added to the basic RWA problem. In flexible optical networks the related problem is the Routing and Spectrum Allocation (RSA), which assigns routes and combines spectrum slots to create *flexpaths* [6]-[8]. Apart from the difference in allocating spectrum resources, when adaptable transponders are present in flexible networks, there are a number of transmission parameters that can be controlled that directly or indirectly affect the resource allocation decision. This is because transmission parameters interrelate the transmission reach with the spectrum used, the rate, and other parameters. Thus, transmission configuration has to be included in the RSA problem, and the multiple degrees of freedom of flexible networks and their interdependencies make connection establishment in such networks more complicated than in fixed-grid WDM networks.

In this paper we spot the differences between flexible optical networks and fixed-grid WDM networks from an algorithmic perspective and outline our solutions for planning and operating flexible networks.

II. ALGORITHMIC ISSUES IN FLEXIBLE OPTICAL NETWORKS

Network planning typically occurs before a network is deployed and focuses on how to better accommodate the current and foreseen network traffic. More common in optical transport networks is to plan the network assuming that the fibers are already installed, which is the reality for tier-1 and tier-2 operators. In this case the topology is given and the purpose is to decide the equipment (transponders, regenerators, switches) to be purchased and where to deploy it. The objective is to minimize the equipment cost and the resources used for serving the given traffic. In the network operation phase, the demands are generally processed upon their arrival, one or a set at a time, and the traffic is accommodated using the equipment already deployed in the network or when necessary some additional equipment. Therefore, the operation process must take into account any constraints posed by the current state of the network. Following the above, the algorithms for optical transport networks can be broadly classified into (i) static (planning) and (ii) dynamic (operating) algorithms. Following is a short list of issues that have to be addressed in flexible networks:

- Accounting for physical layer impairments
- Routing, spectrum allocation, and transmission options
- Serving dynamic traffic fluctuations

III. ACCOUNTING FOR PHYSICAL LAYER IMPAIRMENTS

Optical networks have evolved from opaque (point-to-point) to more transparent connections, as a way to reduce CAPEX and OPEX costs. In the latter case, optical switches are configured to transparently handle transit traffic; the signal remains in the optical domain, saving on the cost of transponders used in the past to terminate and retransmit traffic at intermediate hops. Since optical connections may span over many and long links, physical layer impairments (PLIs), such as noise, dispersion, interference, accumulate and affect the quality of transmission (QoT). Accounting for PLIs

978-1-4799-2358-8/14/$31.00 ©2014 IEEE

is a challenge for algorithm designers, especially with respect to their exact modelling and the interdependencies introduced.

PLIs affect both fixed-grid WDM and flexible networks, but there are distinct differences between the two cases. Flexible networks are expected to use coherent detection and DSP, implying that impairments, particularly those related to dispersion, will be substantially reduced or fully compensated. However, the additional degrees of flexibility available in flexible networks make the minimization of these effects more complicated. On the other hand, PLIs, even though more significant, can be accounted for quite accurately in WDM networks, since fewer parameters are involved (non-tunable transponders and constant guardband) and analytical models successfully capture these effects.

To formulate the PLIs and the transponders' tunability in a flexible network we assume that each flexpath has a specific optical reach, defined as the length it can transmit to with acceptable QoT (e.g., BER). The optical reach depends not only on the flexpath's transmission configuration, but also on the adjacent interfering flexpaths, their transmission configurations and guardbands used. The number of combinations of possible configurations can be huge; also, analytical models for PLIS may not capture all effects or experimental measurements may be limited for some of the options. So, it seems that the only viable solution is to resort to some sort of simplification that captures PLIs in a coarser but safe manner, reducing the parameters and the solution space without eliminating good solutions.

In the algorithms we presented in [6], we used such a simplification that fits well with the above described requirements. The transmission reach of a flexpath is calculated assuming that it suffers worst case interference by other adjacent flexpaths for a given transmission configuration and guardband distance. To be more specific, assume that a flexible transponder of cost c can be tuned to transmit r Gbps using bandwidth of b spectrum slots and a guardband of g spectrum slots from its adjacent spectrum flexpaths to reach l km distance with acceptable QoT. This defines a *physical feasibility function* $l=f_c(r,b,g)$ that captures PLIs and can be obtained experimentally or using analytical models [3][4]. Fig. 1 shows an example of a physical feasibility function without displaying (for illustration purposes) the guardband g, assuming a transmitter capable of transmitting up to 600Gbps in 50GHz. Note that defining the rate r and spectrum b incorporates the choice of the modulation format used.

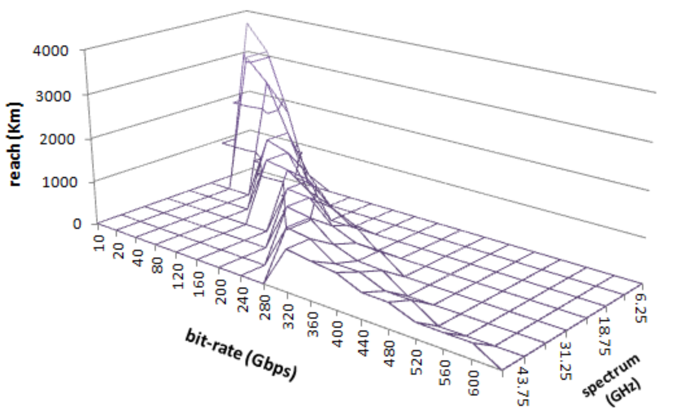

Figure 1: Transmission reach as a function of the rate and spectrum used for a specific flexible transponder.

Using the functions f_c of the available transponders (BVT) we define (*reach-rate-spectrum-guardband-cost*) transmission *tuples*, corresponding to feasible configurations of the BVT. The term "feasible" is used to signify that the tuple definition incorporates PLI limitations, while the cost parameter is used when there are BVTs of different capabilities and costs. The above definition is very general and can be used to describe any type of flexible or even fixed-grid optical network, including mixed-line-rate (MLR) networks [13].

The above methodology can be used to enumerate the feasible transmission options and incorporate in their creation the physical layer effects. The planning algorithms for flexible network to be outlined in the next section takes these as input when examining the options for serving a demand.

IV. Routing and Spectrum Allocation

The problem of establishing connections in fixed-grid WDM networks is typically referred to as the Routing and Wavelength Assignment (RWA) problem, known to be NP-complete. Connection establishment in flexible networks is more complicated for several reasons. First, in contrast to WDM networks where each connection is assigned a single wavelength, in flexible networks spectrum slots can be combined to form variable width channels, leading to the so-called Routing and Spectrum Allocation (RSA) problem. Additionally, the transponders' (BVTs') adaptability yields many transmission options, each with different transmission reach and spectrum used. Algorithms that try to capture, at varying degree, the problem of jointly allocating recourses and selecting the transmission configurations are referred to as Distance Adaptive Spectrum Allocation or as Routing, Modulation Level and Spectrum Allocation (RMLSA).

Both RWA and RSA include as a subproblem the placement of regenerators in the network. Given the BVTs' limited transmission reach, due to PLIs, regenerators are used to establish lengthy connections. However, in contrast to WDM networks where the capabilities of the BVT and regenerators are given, in flexible networks the transmission reach depends on the transmission configuration of the tunable transponder and can be controlled. Thus, regenerator placement in flexible networks also involves the choice of the BVTs' configurations, making it more complicated than in fixed-grid networks.

We now outline a RSA algorithm that takes as input the feasible transmission options of the flexible transponders and the traffic matrix. The demands are served for their requested rates by choosing the route, breaking the transmissions in more than one connection and placing regenerators, if needed, and allocating spectrum to them. The objective is to serve the traffic and find a solution that is Pareto optimal with respect to the maximum spectrum used and the cost (number and type) of transponders used.

A. Static RSA Algorithm Outline

We developed ILP algorithms to plan flexible optical networks [6], but since the problem is NP-hard, we also devised heuristics to find solutions for real problem instances.

Both ILP and heuristic algorithms use a pre-processing phase for calculating the set of path-transmission tuple pairs that are considered as candidate solutions to serve the demands. To do so, for each demand we pre-calculate k paths. Then for each path we find the transmission tuples that have acceptable reach and we form the feasible *path-transmission tuple pairs*. For a path-transmission tuple pair (p,t), given the capacity Λ_{sd} required for demand (s,d) and the rate and reach specified in the transmission tuple t, we calculate the number of connections $W_{p,t}$ and the set of transparent sub-paths $R_{p,t}$ if connections are regenerated, to serve demand (s,d) using (p,t). Thus, a demand served using (p,t) is realized by one or more translucent connections, each comprising of one or more transparent flexpath: (p,m,t,i), $i \in \{1,2,\ldots, W_{p,t}\}$ and $m \in R_{p,t}$.

The set of path-tuple pairs that are candidate to serve the demands are passed to the RSA algorithm, whose role is to choose a path-tuple pair for each demand and assign spectrum to the flexpaths of that path-transmission tuple (recall that a demand can be broken up into $W_{p,t}$ connections, which can be regenerated according to $R_{p,t}$). The number of connections to break each demand and the regeneration points are chosen in the pre-processing phase.

The objective of the devised algorithms is to minimize a weighted sum of the cost of transponders and the total spectrum used. Note that the transponders' cost criterion includes both the type and number of transponders. We used a weighting coefficient W to control the significance given to the two optimization criteria. The objective cost is calculated by multiplying the spectrum used by W and the transponders cost by (1-W) and summing these two. Values of W close to 0 (or close to 1) make the transponders' cost (or spectrum usage, respectively) the dominant optimization criterion. The heuristic serves the demands one-by-one in a particular order, and simulated annealing is used to find good orderings.

B. *Performance of proposed static RSA algorithm*

We compare the performance of a flexible optical network to that of a fixed-grid mixed-line-rate (MLR) WDM network. For the flexible network we assumed the use of a single type of flexible OFDM transponder that supports transmissions up to 50 GHz and modulates up to 64 QAM, so as to transmit up to 400 Gbps. The (reach-rate-spectrum-guardband-cost) tuples used as input to these experiments were obtained from studies on physical layer impairments for optical OFDM flexible networks [4]. For planning the MLR WDM network we also used the heuristic developed for flexible networks, since it is general and can be applied to such networks as well, by defining appropriate tuples. We assumed a MLR system that utilized 4 types of transponders with the following characteristics: (3200 km,10 Gbps,50 GHz,0,1), (2300 km, 40 Gbps,50 GHz,0,2.5), (2100 km, 100 Gbps,50 GHz,0,3.75), and (790 km, 400 Gbps, 50 GHz,0,5.5). The unit cost is taken as the cost of a 10 Gbps transponder. The MLR system employs four transponders of different capabilities and costs, while we assumed that the flexible OFDM network has a single type but tunable transponder. For a fair comparison, we set the cost of the OFDM flexible transponder to 5.5, so that both the OFDM and the 400 Gbps MLR transponders have the same maximum spectral efficiency and cost. We used the DT network topology for the comparison [6], so that the results obtained are representative of real networks We extrapolated future traffic demands for the DT network from 2012 until 2022, assuming that each year the traffic is uniformly increased by 33% (as observed to be the case for the last few years [1]). The average demand capacity for 2012 is 36.5 Gbps (max 115Gbps) and 690Gbps for 2022 (max 2145 Gbps).

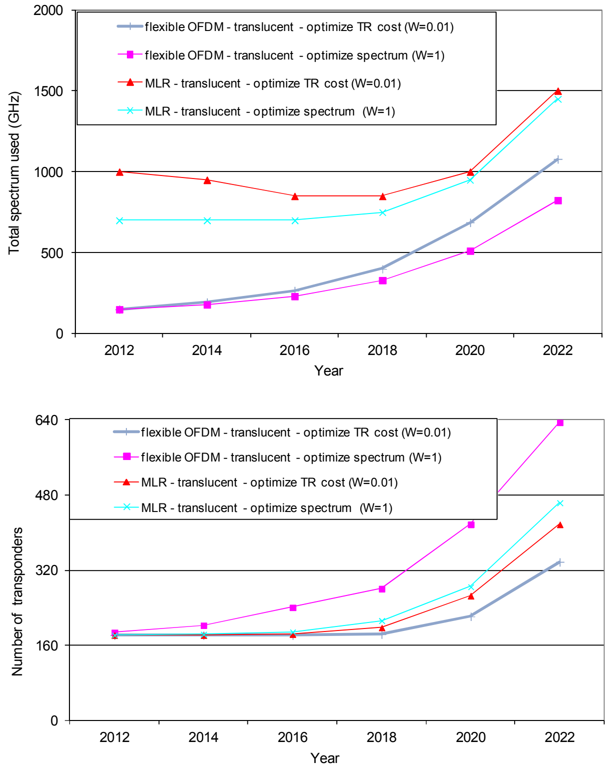

Figure 2: (a) Maximum spectrum used (GHz), and (b) transponders' cost, for the flexible and the MLR network, for optimizing the maximum spectrum used (W=1) and the transponders' cost (W=0.01) on the DT network.

Fig. 2a and b present the results obtained for the two types of networks, flexible OFDM and fixed-grid MLR, and for two different choices of the weighting coefficient in the objective function, namely W=1 (maximum spectrum used minimization) and W=0.01 (transponders' cost minimization). In Fig. 2a we see that the flexible network uses much lower maximum spectrum than the MLR network. This was expected since in the flexible network the connections are established utilizing exactly the amount of spectrum they require, while in the MLR case they always utilize 50 GHz per wavelength and some connections utilize low spectral efficiency transponders (e.g. 10 or 40 Gbps transponders). The *MLR–optimize TR cost* case starts at the year 2012 using high spectrum, because it uses these low spectral efficiency but cheap transponders to serve the traffic (in this case we only optimize the transponders' cost). As the years and the load increase, the MLR network gradually starts employing more the higher spectral efficiency transponders when optimizing both the spectrum (for obvious reasons) and the transponders' cost. This is because, as traffic increases, it becomes more cost-effective to utilize a single high rate transponder than many low rate ones. As the load increases the maximum spectrum used in the *MLR–optimize TR cost* case decreases and then starts to increase again, since after a certain point (year 2018) almost all transponders installed are efficient ones. This is the reason that the performance of *MLR–optimize TR cost* and *MLR–optimize spectrum* cases converge in both Fig. 2a and b. With respect to the transponders cost (Fig. 2b), the *MLR-optimize TR cost* achieves the best performance for light load (see above comments), but after year 2018, it becomes slightly worse than the *OFDM-optimize TR cost* case. At light load, all demands are served by single and transparent connections. As high rate demands appear, the algorithms utilize higher rate transponders (in the MLR network) or higher rate configuration tuples (in the flexible network) but also break some demands into multiple connections and start employing regenerators to enable the use of higher spectral efficiency connections. The finer granularity and the more transmission options of the flexible transponders in the *flexible-optimize spectrum* case can lead to small gains in the transponders' cost, which we observe at heavy loads in Fig. 2b.

The high cost of planning the flexible OFDM network at light load is due to the use of powerful but expensive transponders that are not fully utilized, a problem that would be ameliorated if more than one type of flexible transponders with different performance/cost capabilities were used. The evolution cost is an interesting study, left for future work.

V. DYNAMIC NETWORK OPERATION

The network is typically initiated with an offline/planning algorithm assuming an oversubscribed traffic matrix, to absorb short term fluctuations (e.g. daily cycles) and avoid frequent network upgrades. Thus, the network is operated in an incremental manner, with new connections added sporadically, when utilization between endpoints exceeds a certain percentage, and existing connections rarely (if ever) terminated. Flexible networks using adjustable transponders (BVT) require a different approach than WDM systems as their operation will be more dynamic, having time scales at which optical connection rate changes occur probably orders of magnitude smaller than in fixed-grid networks. Flexible networks can bring the optical layer closer to the IP layer, making the IP layer able to "dial"/control the bandwidth that it uses.

Dynamic traffic variation in flexible networks can be accommodated at two different levels. The first level is the establishment of new connections or the termination down of existing ones, as in fixed-grid networks. Given the high capacity that flexible transponders (BVT) are expected to transmit (designs of 400 Gbps or higher have appeared in the literature [9]), relatively long periods of time will pass until a new connection is established, longer than in WDM networks. A second level is to absorb changes in the requested rate that are short- or medium-term by adapting the BVT, e.g., tuning the modulation format and/or the number of spectrum slots they use, a feature not available in WDM systems.

Fig. 3 describes a generic approach for operating a flexible network. The offline algorithm used to initialize the network, or the online algorithm that subsequently adds flexpaths, assigns to each flexpath a path and a reference frequency. A flexpath occupies a certain amount of spectrum slots around that reference frequency, and traffic variations can be absorbed by the BVT by tuning the modulation format or expanding/ contracting the spectrum they use. Slots that are freed by a flexpath can be assigned to different flexpaths at different time instants, obtaining statistical multiplexing gains.

To enable the dynamic sharing of spectrum, we need what we call Spectrum Expansion/Contraction (SEC) policies to regulate how this is performed. We give examples of such policies in the following paragraph. An RSA operation is performed again when a SEC policy cannot absorb traffic variations by granting additional spectrum slots, or when the requested rate exceeds the transponder capabilities. Then, the RSA algorithm is called to route the excess traffic over a different flexpath, or reroute the entire connection (to save in guardbands).

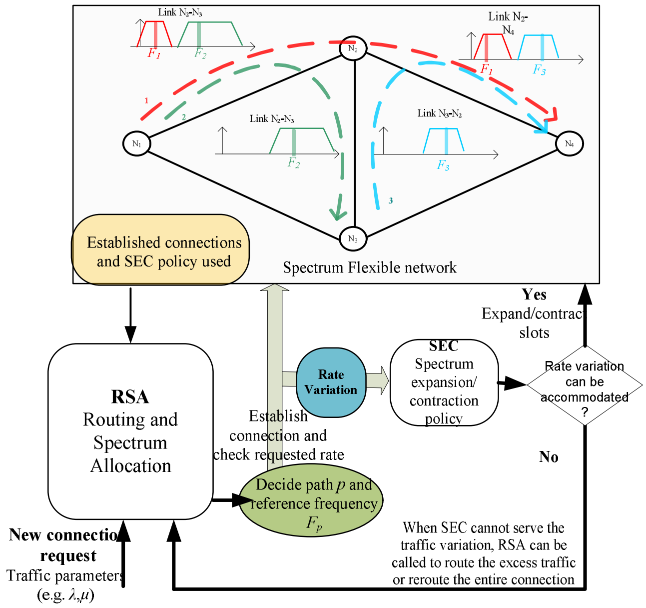

Figure 3: Flow diagram of a generic approach to operate a flexible network.

A. *RSA for time-varying traffic*

We outline simple Spectrum Expansion Contraction (SEC) policies and an RSA algorithm to serve time-varying traffic.

1) *Spectrum Expansion Contraction Policies*

Under the proposed spectrum sharing framework, a connection shares the spectrum slots with its upper and bottom spectrum-adjacent connections. A first policy is called the Constant Spectrum Allocation policy, in which each connection is given a specific amount of spectrum and does not share it with other connections. This policy forms the baseline for the other policies that enable dynamic spectrum sharing among connections.

According to the Dynamic High expansion-Low contraction (DHL) policy, a connection over path p wishing to increase its transmission rate first uses its higher spectrum slots until it reaches a slot already occupied by an upper spectrum-adjacent connection on some link of p. Then, if additional bandwidth is needed, it expands its lower slots until it reaches a slot that is occupied by some bottom spectrum-adjacent connection on some link of p. If the connection needs to increase further its rate and there is no higher or lower free slot space, blocking occurs (for the excess rate). Note that the DHL policy performs indirectly slot defragmentation, since it fills the free higher spectrum slots in every chance it gets. When a connection decreases its spectrum slots due to a reduction in its rate, we first release lower spectrum slots and, if these have been reduced to zero, we release higher slots.

Another SEC policy is the Dynamic Alternate Direction (DAD) policy which aims at the symmetrical use of spectrum around the reference frequencies. With DAD, a connection wishing to increase its transmission rate, alternates between using its higher and lower spectrum slots starting from its higher slots, until it reaches a slot already occupied by an upper or bottom, respectively, spectrum-adjacent connection. Then, if additional slots are needed, it expands towards the other direction, in which case the symmetry is lost. After that it always examines if it can expand towards the direction that uses fewer slots, using slots that were freed by other connections in the meantime. Blocking occurs if the connection needs more slots and there is no higher or lower free slot space. When a connection decreases its slots due to a reduction in its rate, we first release spectrum slots from the direction that has used more slots, and once we have an equal number of higher and lower slots, we decrease the lower spectrum slots. Thus, both expansion and contraction processes are designed to yield symmetrical spectrum utilization so that the central frequency of the connections does not frequently change and remains close to the related reference frequency when congestion is low. More advanced SEC policies that consider the utilization of the spectrum adjacent connections are proposed in [15].

We developed an exact analysis for calculating the blocking probability of each connection and of the whole network for the CSA, the DHL and the DAD policies. These models assume that spectrum slots for each connection are generated according to a Markovian Birth and Death process. We assume that we are given the traffic characteristics of the connections (arrival rates $\boldsymbol{\lambda}$ and average durations $\boldsymbol{\mu}$) and their reference frequencies around which they expand/contract their spectrum. Since in the proposed policies the spectrum is shared only between spectrum-adjacent connections, the blocking states for a connection depend on the utilization of its bottom and upper spectrum-adjacent connections over all the links of the path it follows. In turn, the utilization of these connections depends on the utilization of their bottom and upper spectrum-adjacent connections, and so on. Thus, the interdependence among the connections is quite complicated and cannot be simplified in the general case, except for the simple CSA policy. In the simple CSA policy each connection does not depend on the others and thus its blocking can be calculated using the Erlang-B formula. For the DHL and DAD policies we have to resort to some sort of approximation assumption that limits the interdependence between the connections so as to consider just the spectrum adjacent connections. This enables us to calculate the blocking probability of each connection and the overall probability of the network in a computationally efficient manner.

2) RSA algorithm for time-varying traffic

In the dynamic scenario with time-varying traffic rates that we adopted, the connections expand/contract their utilized spectrum around their reference frequency so as to follow the traffic variations, in the way determined by the SEC policy used. As expected, network performance does not only depend on the SEC policy, but also on the RSA algorithm used, whose role is to assign the routes and reference frequencies so as to minimize the average blocking of the network.

To solve the time-varying RSA problem we transform it into a *static* RSA problem, solve the static problem, apply the blocking models outlined above to calculate the network blocking of this RSA solution, and iteratively search for better solutions. The term *static* is used here to refer to the problem that takes as input a traffic matrix with specific number of required slots for all connections and solves the joint optimization problem for all connections.

To formulate the related static problem, we initially assume that the no spectrum sharing is performed (CSA policy). The blocking performance for the used SEC policy that enables spectrum sharing will be always better or equal to that. Calculating the blocking of CSA policy for each connection is reduced to the Erlang-B formula. We assume we are given a blocking threshold B, that is considered acceptable (e.g., $B=10^{-6}$) and we use Erlang-B formula to calculate for each connection p the number N_p of spectrum slots for which the CSA blocking is acceptable. We use the set of N_p values for all connections, as the traffic matrix in a *static* Routing and Spectrum Allocation (RSA) algorithm to find a path and a reference frequency slot for all connections. We denote by $T^* = \max_p(F_p + N_p)$ the highest slot allocated to a connection by the static algorithm. If the system can support T^* subcarriers slots: $T^*<T$, where T is the number of slots supported by the system, the algorithm finishes and we have found an acceptable solution with the CSA policy that does not require spectrum sharing at all. We can use heuristic algorithms as the one outlined for static traffic in the previous section. If the RSA algorithm does not find a solution within T slots, we iteratively increase the related acceptable blocking threshold B and repeat the above.

A static RSA solution is acceptable if it utilizes less than the T slots supported by the system. After obtaining an acceptable static RSA solution we take into account the specific SEC policy used and apply the corresponding blocking model to calculate the average network blocking. However, we do not stop the first time we find an acceptable solution within T slots, but we search for different static RSA solutions with the same numbers of required slots or keep decreasing the number of required slots until we find K solutions that are acceptable and select the one with the lowest network blocking.

It is clear that the problem of finding the RSA solution that minimizes network blocking is very complicated. In the general case, the network blocking depends on the paths, the reference frequencies and ordering of the connections, the SEC policy, and the traffic parameters. The proposed algorithm solves the time-varying RSA problem *indirectly*. It solves a related static problem, considering the CSA policy, and then applies the blocking models developed for the particular SEC policy used to estimate the performance under time-varying traffic. In the future we plan to analyze more SEC policies and also work on time-varying RSA algorithms that will incorporate more *directly* the policy blocking models.

B. Performance Results

We present performance evaluation results for serving traffic with time-varying rates in a spectrum-flexible network under the proposed framework, for the SEC policies and the RSA algorithm outlined above.

We performed experiments using the 14-node DT network. We assumed that communication is performed among all source-destination pairs in the network. Spectrum slot requests for each source-destination pair p are generated with a Poison process of rate λ_p and their duration is exponentially distributed with mean $1/\mu_p=1$. The arrival rate λ_p for the slot requests of each connection p is drawn from an exponential distribution with mean λ.

We graph the blocking performance of the CSA, the DHL, and the DAD policies. For the same traffic scenarios, we conducted full network simulation experiments and we also graph the corresponding blocking probability returned by the simulations for 10^7 slot requests. For comparison purposes, we also present the blocking performance of a network that does not follow the framework and SEC policies, but supports the full sharing of all spectrum slots among the connections. This type of network can be viewed as a typical WDM network, with the additional constraint of having to use spectrum guardbands between spectrum-adjacent connections.

In Fig. 4 we graph the network blocking performance as a function of the total load (in Erlangs) assuming T=250 spectrum slots are available in the network. We observe that the proposed analytical models for calculating the blocking probabilities of the CSA, DHL and DAD policies are in very close agreement with the corresponding simulation results. The calculations of the average network blocking performance (for all connections) for a given RSA solution for the DHL policy took around 10 sec, and 30 sec for the DAD policy, while the related simulations required around 15 minutes, proving that we are able to calculate the network blocking probabilities in an accurate and quick manner.

The network blocking probability of the DAD and DHL policies is lower than that obtained for the CSA policy by more than one order of magnitude, in most cases. This is the gain that we obtain by enabling spectrum slot sharing among spectrum-adjacent connections, as done with these dynamic policies. The DAD policy achieves better performance than the DHL policy. In the DHL policy each connection utilizes more heavily its higher spectrum slots, and the statistical multiplexing gains are achieved by sharing the spectrum mainly with its bottom adjacent connections. The DAD policy is more symmetric and connections share spectrum more efficiently with both their upper and bottom adjacent connections achieving higher multiplexing gains. The path length has both positive and negative effects in the blocking performance. The longer paths deteriorate the spectrum allocation process and result in higher fragmentation of spectrum. The dynamic spectrum sharing performed by the DHL and DAD policies, resolves partly the defragmentation problem. However, the longer the path becomes, the more connections compete for the same (spectrum) resources, and after a point efficient spectrum sharing becomes difficult. So the placement of adjacent connections in the network (role of RSA algorithm) becomes very important on networks with longer paths. The performance of the WDM network and the corresponding RWA algorithms utilizing $T/2$ wavelengths is worse than the solutions that follow the proposed framework.

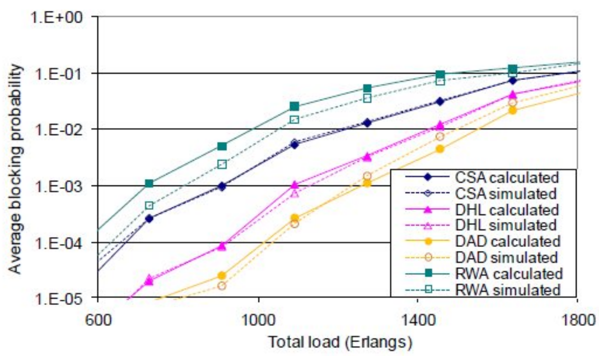

Figure 4: Blocking performance of the different SEC polices for the flexible optical network and a WDM network.

VI. CONCLUSIONS

The introduction of flexible optical networking raises the need for the development of new algorithms that are considerably more complicated than those used in traditional WDM networks. We outlined some of the basic differences between these two networking paradigms and presented our solutions for establishing connections in a flexible optical network under static and time varying traffic.

ACKNOWLEDGEMENTS

This work was partially funded by the European Community's FP7 through IP IDEALIST under grant agreement nº 317999.

REFERENCES

[1] Cisco white paper, "Cisco Visual Networking Index: Forecast and Methodology, 2010–2015"

[2] O. Gerstel, M. Jinno, A. Lord, S. J. Ben Yoo, "Elastic Optical Networking: A New Dawn for the Optical Layer?", IEEE Communications Magazine, 50 (2), Feb, 2012

[3] A. Klekamp, R. Dischler, R. Buchali, "Limits of Spectral Efficiency and Transmission Reach of Optical-OFDM Superchannels for Adaptive Networks", IEEE Photonics Technology Letters, 23 (20), 2011.

[4] R Borkowski, F. Karinou, M. Angelou, V. Arlunno, D. Zibar, D. Klonidis, N. G. Gonzalez, A. Caballero, I. Tomkos, I. T. Monroy, "Experimental Study on OSNR Requirements for Spectrum-Flexible Optical Networks", Journal on Optical Communications and Networking, 4 (11), 2012

[5] 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 Communications Magazine, 48 (8), pp.138-145, 2010.

[6] K. Christodoulopoulos, P. Soumplis, E. Varvarigos, "Planning Flexgrid Optical Networks under Physical Layer Constraints", Journal of Optical Communications and Networking, 2013.

[7] L. Velasco, M. Klinkowski, M. Ruiz, J. Comellas, "Modeling the Routing and Spectrum Allocation Problem for Flexgrid Optical Networks", Springer Photonic Network Communications, 24, pp. 177-186, 2012.

[8] A. N. Patel, P. N. Ji, J. P. Jue, T. Wang, "Routing, wavelength assignment, and spectrum allocation algorithms in transparent flexible optical WDM networks", Optical Switching and Networking (OSN), 9 (3), 2012.

[9] A. Autenrieth, et. al., "Evaluation of Technology Options for Software-Defined Transceivers in Fixed WDM Grid versus Flexible WDM Grid Optical Transport Networks", Photonics Network Symposium, 2013.

[10] K. Christodoulopoulos, I. Tomkos, E. Varvarigos, "Time-Varying Spectrum Allocation Policies and Blocking Analysis in Flexible Optical Networks", Journal on Selected Areas in Communication, 31 (1), 2013.

[11] T. Takagi, H. Hasegawa, K. Sato, Y. Sone, B. Kozicki, A. Hirano, M. Jinno, "Dynamic routing and frequency slot assignment for elastic optical path networks that adopt distance adaptive modulation", Optical Fiber Communication Conference (OFC), 2011.

[12] A.N. Patel, P.N. Ji, J.P. Jue, W. Ting, "Defragmentation of transparent Flexible optical WDM (FWDM) networks", Optical Fiber Communication Conference (OFC), 2011.

[13] A. Nag, M. Tornatore, B. Mukherjee, "Optical Network Design With Mixed Line Rates and Multiple Modulation Formats," IEEE/OSA Journal of Lightwave Technology, 28 (4), pp.466-475, Feb. 2010.

[14] E. Palkopoulou, M. Angelou, D. Klonidis, K. Christodoulopoulos, A. Klekamp, F. Buchali, E. Varvarigos, I. Tomkos, "Quantifying Spectrum, Cost, and Energy Efficiency in Fixed-Grid and Flex-Grid Networks", Journal of Optical Communications and Networking, 4 (11), 2012.

[15] E. Palkopoulou, L. Stiakogiannakis, D. Klonidis, K. Christodoulopoulos, E. Varvarigos. O. Gerstel, I. Tomkos "Dynamic Cooperative Spectrum Sharing in Elastic Networks", OFC 2013.