Trading off Transponders for Spectrum in Flexgrid Networks

[1]K. Christodoulopoulos, [2]P. Soumplis, [2]E. Varvarigos
1: School of Computer Science and Statistics, Trinity College Dublin, Ireland. Email christok@tcd.ie
2: Department of Computer Engineering and Informatics, University of Patras, Greece. Emails: {soumplis,manos}@ceid.upatras.gr

Abstract: We propose algorithms for planning flexgrid networks under physical layer impairments. Using an optimization function that accounts for both the spectrum used and the transponders cost, we observe tradeoffs between these two optimization parameters.
OCIS codes: 060.0060, 060.4251, 060.4254, 060.4256

1. Introduction

To cope with the increased capacity requirements in transport networks recent research efforts are focusing on optical architectures that support variable spectrum connections as a way to improve network efficiency. Spectrum flexible, elastic, or flexgrid are examples of the terms used to describe solutions that migrate from the fixed 100 or 50GHz grid WDM systems to variable-spectrum systems [1]. A flexgrid network has spectrum granularity finer than that of standard WDM systems and can also combine the spectrum units (referred to as slots) to create wider channels on an as needed basis. A bandwidth-variable transponder uses just enough spectrum to serve the demand and every bandwidth-variable OXC on the path establishes a cross-connection with sufficient spectrum to create an appropriately sized end-to-end optical connection. Several networking paradigms adopting the flexgrid approach have emerged in the past few years. SLICE [2] uses optical OFDM, which distributes the data on several low data rate subcarriers that can overlap, since they are orthogonally modulated, to achieve high spectral efficiency. Single-carrier systems may also operate in a flexgrid manner, such as the Flexible-WDM (FWDM) architecture considered in [4].

Planning a flexgrid optical network [2]-[5] is typically performed by serving the demands for their requested capacities, assumed to be known in advance, by elastically allocating spectrum to them, satisfying the spectrum continuity and non-overlapping assignment constraints. The problem is typically referred to as the *Routing* and *Spectrum Allocation* (RSA), or when the modulation level can be also chosen flexibly as *Routing*, *Modulation Level* and *Spectrum Allocation* (RMLSA). [2] presents a scheme to allocate the minimum spectral resource based on the paths distance. In our previous work [3], we provided an optimal RMLSA algorithm based on integer linear programming (ILP) and a heuristic to solve large problem instances. [4] presents an alternative ILP formulation and a load balancing heuristic. The planning of an FWDM network is studied in [5].

We consider the planning problem of a flexgrid optical network under physical layer impairments. Physical layer effects are incorporated in the definition of the *feasible* transmission configurations of the transponders, described by (*rate-reach-spectrum-guardband-cost*) *tuples*. Given the transponders' feasible configuration tuples and the traffic matrix, we formulate the planning problem of a flexgrid network considering both the use or not of regenerators. Demands are served for their requested rates by choosing the route, breaking the transmission in multiple connections, placing regenerators if needed, and allocating spectrum to them. Connections are separated by appropriate spectrum guardbands so that physical layer interference is kept at acceptable levels. The objective is to serve the traffic and find a solution that is Pareto optimal with respect to the spectrum usage and the cost of the transponders used. We devised algorithms for planning transparent (without regenerators) and translucent (with regenerators) networks, which extend previous solutions [2]-[5] in a number of aspects. The devised algorithms consider the physical layer in more details than the simplistic reach-modulation level transmission constraints previously used. Also, previous algorithms considered only a single connection per demand (source-destination) and only transparent connections, while the new algorithms decide also on how to break the demands into multiple connections and place regenerators, if needed. Finally, the objective of previously proposed algorithms was to minimize the spectrum used, while the new algorithms consider both the used spectrum and the cost of used transponders in a multi-objective optimization formulation.

Using realistic transmission specifications we compare the performance of a flexgrid OFDM to that of a Mixed Line Rate (MLR) network and verify the achieved gains. Moreover, by examining different optimization weights for the spectrum used and the cost of transponders, we observe the tradeoffs between these two optimization criteria.

2. Problem description and proposed algorithms

We consider a slotted flexgrid network with transponders that have two flexibility degrees, enabling: (a) the selection of the modulation format and (b) the choice of the spectrum (in contiguous spectrum slots) they use. By adapting these, the flexible transponder can be tuned to transmit at a specific *rate* over a specific *reach* using a specific amount of *spectrum* (in slots) and requiring a specific *guardband* (in slots) from the spectrum-adjacent connections to exhibit acceptable transmission quality, what we call a *feasible* (*rate-reach-spectrum-guardband-cost*) *transmission tuple*. The cost parameter is used when we have different types of transponders with different capabilities. Note that the rate/spectrum parameters of a tuple incorporate the choice of the transmission's modulation format. Also the definition of the tuples allows variable guardband connections. Given the transponder capabilities (parameters such as the maximum baud rate, the modulation format, the spectrum used and the bpd, can be the limiting factors), and since the modulation format and the spectrum used are selected from discrete sets, we have certain *feasible* transmission configurations (tuples) for the transponders, which incorporate in their definitions the physical layer impairments, rendering the algorithms developed impairment aware (IA). The problem definition and the proposed algorithms are general and applicable to flexgrid as well as fixed-grid networks. The only requirement is the ability to express the feasible transmission options of the transponders as (rate-reach-spectrum-guardband-cost) tuples. For example for a MLR network we have to consider tuples in which the spectrum is always constant (50 GHz), the rate-reach-cost parameters capture the physical layer transmission limitations for the different types of available transponders, and the guardband is set to zero. The planning problem is defined as follows: Given the network topology, the traffic matrix, and the feasible transmission configurations of the transponders, the objective is to serve the traffic and minimize a function of the spectrum used and the cost of the transponders. Fig.1 shows an instance of the planning problem.

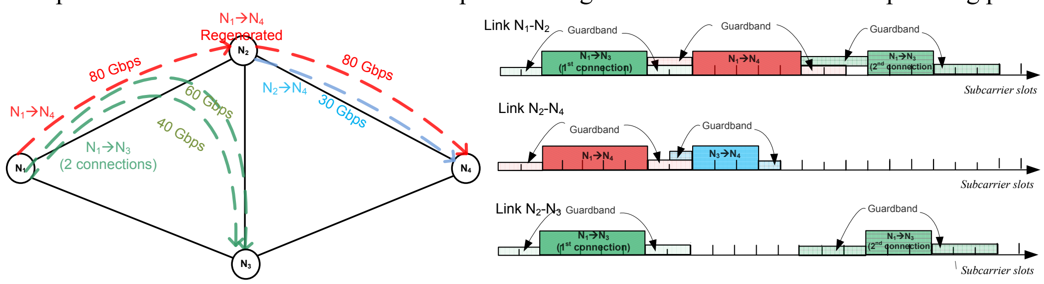

Fig. 1: Instance of the planning problem in a flexgrid network and spectrum slot usage of the 3 used links.

2.1. Algorithms

We developed ILP algorithms to plan transparent (without regenerators) and translucent (with regenerators) networks. Since the planning problem is NP-hard [3], we also devised heuristic algorithms to find solutions for real problem instances. Both ILP and heuristic algorithms use a pre-processing phase for calculating the set of path-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 tuples that have acceptable reach and define what we call *path-tuple pairs*. For a path-tuple pair (p,t), given the traffic Λ_{sd} to be transferred for demand (s,d) and the rate and reach specified in 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) with (p,t). Thus, in the translucent network setting, a transmission using path-tuple pair (p,t) is realized by one or more translucent connections, each comprising of one or more transparent flexgrid lightpaths: (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 connections of that path-tuple (recall that a demand can be broken up to $W_{p,t}$ 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 ILP and heuristic 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 follows the approach of [3] (it serves the demands one-by-one in a particular order, and simulated annealing is used to find good orderings), with a number of enhancements. It uses the path-tuple pairs as input, accounts for multi-connection establishment and regeneration, and also for non-constant guardband connections.

3. Performance Results

We used the 14-node DT network topology and, starting with a realistic current traffic matrix, we scaled it up assuming a uniform increase of 34% per year to obtained matrices for years 2012 to 2022 (6.5 to 115 Tbps, respectively). The transmission specifications were based on physical layer studies for optical OFDM reported in [6]. We compare the performance of the OFDM flexgrid network to that of a MLR network. To solve the RWA MLR problem, we used the same heuristic algorithm for the flexgrid network (see discussion in Section 2). We assumed a MLR system that utilizes transponders with the following (rate-reach-cost) characteristics: (10 Gbps-3200 km-1), (40 Gbps-2300 km-2.5), (100 Gbps-2100 km-3.75), (400 Gbps-790 km-5.5). We assumed that in the flexgrid network we have a single type of OFDM transponder with maximum rate of 400 Gbps, and cost and spectral efficiency equal to those of the 400Gbps WDM transponder, to have a fair comparison. We present the spectrum usage and the cost of transponders (recall that different transponders have different costs) for *W*=0.01 and *W*=1, and then for values of *W* in-between. Setting *W*=0.01 optimizes mainly the transponders' cost, while setting *W*=1 optimizes solely the spectrum used. We avoided using *W*=0, since this removes the combinatorial nature of the problem and the algorithm selects for each demand the path-tuple pair with the lowest cost irrespective of the other demands.

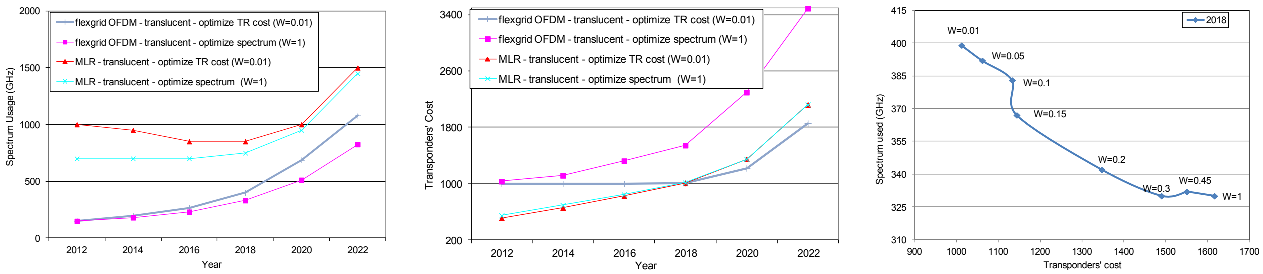

Fig. 2: (a) Spectrum used and (b) Transponders' cost as a function of time (load). (c) Tradeoff between spectrum used and transponders' cost in the flexgrid network for year 2018.

From Fig 2a we see that the flexgrid network uses significantly less total spectrum than the MLR network. This was expected since the flexgrid network has finer spectrum granularity, while in the MLR network the connections utilize always 50 GHz wavelengths and often utilize lower cost but also lower spectral efficiency (i.e., 10 or 40 Gbps) transponders. The *MLR–optimize TR cost* case starts with high spectrum usage for year 2012, because it uses these low spectral efficiency but cheap transponders (remember that in this case we optimize the transponders' cost). As years (and load) increase, the MLR network gradually starts employing the more spectrally efficient transponders when optimizing both the spectrum (for obvious reasons) and the transponders' cost (as traffic increases, it becomes more cost-efficient to utilize a single high rate transponder instead of many low rate ones). So, as the load increases the spectrum usage of the *MLR–optimize TR cost* case decreases and then starts to increase again and converges to the *MLR–optimize spectrum*, since after year 2018 almost all spectrally inefficient transponders have been replaced by efficient ones. This is the reason the *MLR–optimize TR cost* case yields the lowest cost at light load, an advantage that is lost after 2018. The differences in spectrum usage between the flexgrid and the MLR network decrease slightly as the load increases, but they remain significant even for the highest load examined. As load increases, the cost of the MLR and the flexgrid optical networks converge and after a point the *flexgrid-optimize TR cost* case becomes better than the *MLR-optimize TR cost* case. The finer granularity and more transmission options of the OFDM transponders lead to these gains. The high cost of the flexgrid network at light load is because it uses powerful but expensive transponders whose capabilities are not fully utilized, a problem that would be ameliorated, if more than one type of flexgrid transponders with different performance/cost capabilities were used. However, from the operator's point of view, it might make sense to place powerful and tunable transponders at an earlier stage, when these have reasonable prices, and tune them to serve the increased traffic demands at later years. In the MLR case, the cheap-low rate transponders have to be replaced later by more efficient ones. This additional cost is not accounted for here, since we don't study the incremental evolution of the network. Fig. 2c presents the Pareto optimal front obtained by ranging the optimization coefficients *W* between 0.01 and 1 for the flexgrid network and year 2018. To reduce the used spectrum ($W\rightarrow 1$) a larger transponders' cost is encountered, indicating that some transponders use transmission tuples with high modulation format (more bits/symbol) but low reach and low rate. On the other hand, minimizing the transponders' cost ($W\rightarrow 0$) selects tuples with the maximum total rate that might not use the higher possible modulation format. In Fig. 2c we see a significant tradeoff between the spectrum used and the cost of the transponders. Depending on the actual market prices of these two parameters, we can pick the solution that achieves the minimum CAPEX.

4. Conclusions

We devised planning algorithms that account for the physical layer impairments that are general and can be used for flexgrid but also for fixed-grid optical networks. Using realistic transmission specifications we verified the gains that can be obtained by a flexgrid OFDM as opposed to a MLR network. Using an optimization function that accounts for both the spectrum usage and the transponders cost (number and type), we observed significant tradeoffs between these optimization parameters. Depending on the actual market prices, we can pick the optimal planning solution.

5. References

[1] O. Gerstel, et. al., "Elastic Optical Networking: A New Dawn for the Optical Layer?", IEEE Com. Mag 2012
[2] M. Jinno, et. al, "Distance-adaptive spectrum resource allocation in spectrum-sliced elastic optical path network", IEEE Com. Mag. 2010.
[3] K. Christodoulopoulos, et. al., "Elastic Bandwidth Allocation in Flexible OFDM-Based Optical Networks", JLT, 2011.
[4] M Klinkowski et. al "Off-line algorithms for Routing, Modulation Level, and Spectrum Assignment in elastic optical networks", ICTON, 2011
[5] A Patel et. al. "Routing, wavelength assignment, and spectrum allocation algorithms in transparent flexible optical WDM networks" OSN, 2012
[6] A Klekamp et al "Limits of Spectral Efficiency and Transmission Reach of Optical-OFDM Superchannels for Adaptive Networks", PTL, 2011.