Optimally-driven Online Reservations in Elastic Optical Networks

P. Soumplis (1), P. Kokkinos (1), E. A. Varvarigos (2)

(1) Computer Engineering & Informatics Dept, University of Patras, Greece, {soumplis, kokkinop} @ceid.upatras.gr
(2) Electrical & Computer Engineering Dept, National Technical University of Athens, Greece, vmanos@central.ntua.gr

Abstract *We present an admission control and solution validation-based online, routing, spectrum allocation approach for elastic optical networks. This is parameterized periodically by an optimal offline mechanism that drives the blocking of traffic requests, which if served would negatively affect network performance.*

Introduction

Advances in optical communications and network programmability enable the creation and offering of network capacity services. Such services enable the on demand reservation of raw network capacity, in the form of optical lightpaths and for a specific duration, interconnecting branch offices, datacenters or devices around the world[1].

In particular, elastic optical networks with their fine spectrum granularity[2] and the use of elastic transponders can lead to the efficient and dynamic use of the available spectrum, based on the actual traffic requirements. Elastic optical networks along with the Software Defined Networking (SDN) paradigm, which allows the separation of the control and the forwarding plane, provide the ability to an operator to decide on the specific characteristics of its network dynamically and from a central point of control.

Actually, today we are already witnessing some proof of this; e.g., Pacnet network in the Asia-Pacific region provides scalable bandwidth and software-enabled intelligence, allowing customers to dynamically provision bandwidth in 25 Gb/s, 37.5Gb/s, 50 Gb/s or 100 Gb/s increments, in a matter of minutes, with deactivation at the end of a customer-specified period[3].

The immediate or in advance or timed reservation of capacity has been studied considering generic networks but also elastic optical ones[4],[5]. Most works formulate the specific problem using an Integer Linear Programming (ILP) model, which is too computationally intensive to solve and then propose a respective heuristic.

In practice, for the realization of the described capacity services, network operators require efficient, in terms of time and performance, algorithmic approaches to decide in almost real time, the way incoming requests will be served.

In our work, we present an approach based on admission control and solution validation, for the routing and spectrum allocation in elastic optical networks. This makes use of an optimal offline mechanism for parameterizing the way the online algorithm will operate. To do so, it blocks connections that if served would negatively affect the performance of the network in the future (waste of network resources, increasing the fragmentation of the network, etc). This approach can be used for both immediate and advance reservation of network resources. The idea of admission control for advance reservations was being studied some years ago for general networks[6], while since then it has also been applied in other fields such as that of ad auctions[7] that require real-time service, and optimality in the performed decisions. We show through simulations that this optimally-driven approach for the admission control of the requests and for the solution validation, for online reservations in elastic optical network resources, can improve the performance, while reducing the total execution time.

Optimally-driven admission control and solution validation for online reservations

Fig. 1 illustrates the operation of the proposed approach in time, along with the required modules and their interactions.

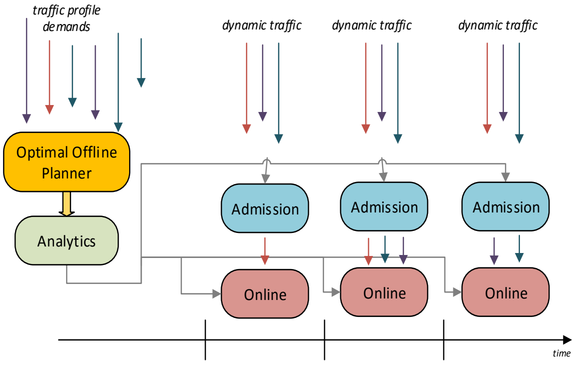

Fig. 1: Optimally-driven admission control for online reservations

An *Optimal Offline Planner* (or *Planner*) receives as input historical traffic data (e.g., from the previous day) and makes an offline network resource allocation for these demands. The Planner is executed at long periodic intervals (e.g., once a day), targeting the exploration of large solution spaces and the identification of optimal or near-optimal solutions, in comparison to the respective Online mechanism. The decisions taken by the Planner are then analysed by the *Analytics* module, so as to extract parameters that will drive the online operation of the Admission and Online modules.

The Admission module that runs prior to the advanced reservation mechanism (the Online module) blocks connection requests that may negatively affect that network state, as this state

978-1-5386-5624-2/17/$31.00 ©2017 IEEE

has been defined/decided by the Planner. We assume that time is partitioned in timeslots defining resource reservations granularity. The Admission and the Online modules run at the start of each timeslot, in a period shorter than that of the Planner and the Analytics modules, serving the demands that have arrived during the previous slot. The Online module serves the demands that are "accepted" by the Admission module, reserving the respective network resources, trying to maintain the critical network performance parameters, e.g., network utilization as close as possible to the optimal. As a result, a solution may be rejected based on the Analytics' inputs and a new solution will be searched. The Online module must maintain a balance between the achieved performance and the execution time, exploring a smaller solution space than the Offline module.

We should also note that in order for this approach to operate efficiently, it is necessary that the daily traffic patterns remain semi-static. In case it is identified (e.g., by a network monitoring entity) that the network traffic pattern has changed, the operator (or an automatic mechanism) may choose to re-run the Optimal Offline Planner, so as to improve the online decisions.

It is evident from the above that many design decisions need to be investigated, during the actual realization of the presented approach (Fig. 1). Till today, several offline planning and online operation algorithms have been presented in the literature for elastic optical networks that can be utilized in this process[4],[5]. Also, the periodicity with which the Planner, the Analytics, the Admission and the Online modules are executed affects efficiency of the approach. In addition, the input to the Planner can either be the traffic profile from the previous day, as already mentioned, or one that represents the status of the network at steady state. Moreover, the Analytics module's extracted parameters that are fed to the Admission and Online modules are critical. Different parameters can be extracted that reflect different views of the network's status (e.g. utilization, fragmentation, accepted connections).

Realization of the approach

For the realization of the presented optimal-driven approach, we assume there is a number of classes of users that request capacity services. Each class has particular statistical characteristics, regarding the arrival rate, the average requested bandwidth and requested reservation duration, expressed in time slots.

For the Optimal Offline Planner, we have developed an Integer Linear Programming (ILP) mechanism to address the spectrum and path allocation problem, which finds the optimal solution based on the received input traffic matrix. This traffic matrix is composed of requests belonging to different, the starting time of each request and the reservation time. This mechanism is the extension of the ILP for dynamic traffic[5], with additional time constraints. In particular, Boolean variable $x_{\theta,p,t}$ is introduced, which is equal to 1 if transmission over path *p* using transponder configuration *t* is used to serve connection *θ* (belonging to the set of existing connections $\overline{\Theta}$ or to the new demands set Θ), and equal to 0, otherwise. We use two additional constraints: The first is about already established demands $\theta \in \overline{\Theta}$, to ensure that they will be served for the requested number of time slots.

$$\sum_{\theta\in\overline{\Theta},(p,t)\in Q_{\overline{\Theta}}} x_{\theta,p,t} = 1 \qquad (1)$$

The second is about new demands, $\theta \in \Theta$, which can be either accepted or blocked:

$$\sum_{\theta\in\Theta,(p,t)\in Q_{\Theta}} x_{\theta,p,t} \geq 0 \qquad (2)$$

What is more, a weighted objective function is used that tries to minimize the number of blocked connections, ensuring fairness among demands of different classes. To do so, different weighting coefficients w_θ are used for the respective classes.

$$\max\Big(\sum_{\theta\in\Theta,(p,t)\in Q_{\Theta}} w_\theta \cdot x_{\theta,p,t}\Big) \qquad (3)$$

The output of the ILP is analysed and various parameters are extracted and used, in order to guide the Admission and Online processes. These parameters attempt to describe the network status, in steady state, that if achieved will render the allocation of resources near optimal. In practice, the extracted parameters for the Admission process include for each class the percentage of requests served or the total capacity reserved. Furthermore, the utilization of the network links is used to guide the Online process, in order to search for better solutions when needed and maintain a state utilization as close as possible to that planned by the ILP algorithm. To do so, the online operating Admission process blocks requests that pass the capacity limits identified by the ILP for the respective class. Also, a request may be blocked, even if there is available spectrum, suggesting that its acceptance may negatively affect the performance of the network in the future.

The Online mechanism serves new demands one by one at the start of the time slot. Already established demands remain unaffected while demands that expire are removed from the current state of the network, freeing resources/spectrum. Initially, based on the requested demand, as well as its source destination pair, the Online algorithm selects *k*-shortest paths. Then, it breaks the demand to connections based on the capabilities of the transponders as well as the length of the path where the connections will be established, also determining the nodes where regenerators will be placed (namely path-tuple pairs). Then it tries to establish the connections searching among voids of spectrum with proper size. The objective is to select an appropriate spectrum void that minimizes a weighted objective function, which depends on the number of the utilized spectrum slots. Each solution found is checked against the network links' utilization extracted by the ILP. If no valid solution is found, the Online algorithm either searches for a new one or performs push pull and rerouting techniques to free the required spectrum and finally serve the demand.

Performance Results

For our simulations we used a six node network topology consisting of 18 links and link lengths that vary between 200 Km and 600 Km. Spectrum slots were taken to occupy 12.5 GHz, while a network link supports 320 spectrum slots. We assumed the use of a single type of elastic transponder that transmits up to 400 Gbps. The (reach-rate-spectrum-guardband) transponders capabilities were obtained from[5]. Demands are generated at each node according to a Poisson process with arrival rate λ and an exponentially distributed duration with mean $1/\mu_i$ time units for class i, with the destination uniformly chosen among all nodes. The duration of the time slot is taken to be one time unit. We run our algorithms 5000 until connections were established. Each request's demand rate derives from the following discrete uniform distribution: [100, 200, 300, 400] Gbps.

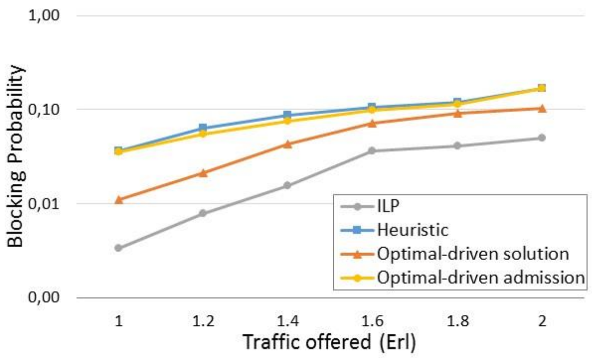

Fig. 2: Blocking probability for different traffic loads.

In what follows we compare the *ILP*, the *Heuristic*: Online without admission control and solution validation, the *Optimal-drive solution*: Online with admission control and solution validation and the *Optimal-driven admission*: Online with admission control.

Fig.2 presents the achieved blocking probability as the offered traffic increases. The optimal-driven admission heuristic's performance is not improved much and it is close to the heuristic. On the other hand, the optimal-driven solution heuristic performs better and closer to the optimal one.

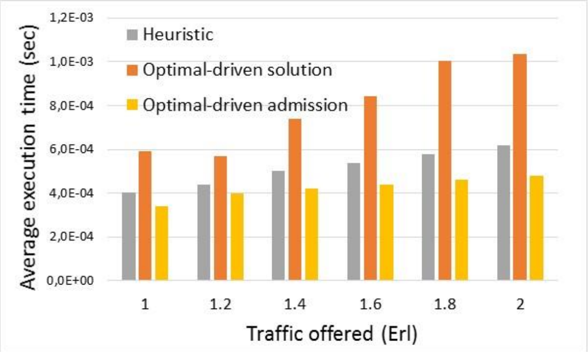

Fig. 3: Average execution time of the algorithms per connection for different traffic loads.

Fig. 3 shows the average execution time of the algorithms per connection for different traffic loads. The optimal-driven admission heuristic that blocks demands before any processing, reduces the average execution time in comparison to the Heuristic. On the other hand, the optimal-driven solution heuristic, increases the execution time, introducing a trade-off between this and the improvement in performance achieved by searching for better solutions

Conclusions

SDN-enabled elastic optical networks are able to offer on demand capacity services for point to point (e.g., between datacenters) and end-to-end (for 5G) demands. We show that it is possible to use an optimal-ILP's algorithm outputs as input to a heuristic, so as to improve its performance and execution time. In that case it (i) blocks connections that if served could negatively affect the performance of the network in the future and (ii) evaluates the quality of the proposed heuristic solutions, searching for better ones when needed.

Acknowledgements

This work was partially supported by the EC through the Horizon 2020 Nephele project (g.a. 645212).

References

[1] P. Kokkinos, et al., "Survey: Live Migration and Disaster Recovery over Long-Distance Networks", ACM Computing Surveys, Vol. 49, No. 2, 26 pages, 2016.

[2] K. Christodoulopoulos, et al., "Elastic Bandwidth Allocation in Flexible OFDM-based Optical Networks", IEEE/OSA Journal of Lightwave Technology, 2011.

[3] Pacnet – network virtualization, http://bit.ly/2pPX7r5, Retrieved February 2017.

[4] Wei Lu, et al., "On Hybrid IR and AR Service Provisioning in Elastic Optical Networks", IEEE/OSA Journal of Lightwave Technology, Vol. 33, pp.4659-4670, 2015

[5] P. Soumplis, et al., "Dynamic connection establishment and network re-optimization in flexible optical networks," Optical Network Design and Modeling, pp. 61-66, 2014.

[6] D. Wischik and A. Greenberg, "Admission control for booking ahead shared resources," INFOCOM, pp. 873–882, 1998.

[7] D. Parkes, et al. "Optimize-and-dispatch architecture for expressive ad auctions", ACM Conference, 2005.