

A Framework for Dynamic Multi-layer Resource Allocation and Optimization in Application-Centric Networking

Ćiril Rožić^{(1)*}, Marco Savi⁽²⁾, Chris Matrakidis⁽¹⁾,
Dimitrios Klonidis⁽¹⁾, Domenico Siracusa⁽²⁾, and Ioannis Tomkos⁽¹⁾

⁽¹⁾Athens Information Technology (AIT), 44 Kifisias Ave., 15125 Marousi, Greece

⁽²⁾CREATE-NET, Via alla Cascata 56/D, 38123, Povo, Trento, Italy

*croz@ait.gr

Abstract: In an SDN-based network, connection requests can be accommodated according to application requirements. We devise a framework where such requirements drive IP and optical network resource allocation, dynamic optimization, and instantiation through an SDN orchestrator.

OCIS codes: (060.4256) Networks, network optimization, (060.4251) Networks, assignment and routing algorithms

1. Introduction

Recently, Software-Defined Networking (SDN) has been enabling service providers to use the existing optical network resources more intelligently and efficiently. This advance allows for assigning bandwidth where and when needed, and coping with the increase in the number of traffic flows. The IP/optical network thus approaches the realm of computer-like programmability.

In our project [1] we use SDN principles to address the challenge of diversifying network applications. Specifically, it is now becoming insufficient to merely provide bandwidth for an application, because the application may instead have very specific *service requirements* such as maximum latency in addition to bandwidth. To address this, we build an SDN *network orchestrator* as the connecting component and intelligence between the network applications and the underlying network infrastructure [2]. The orchestrator first exposes a set of supported service requirements called *network primitives* which the network can support. The application then chooses among these, thus expressing its *intent*. The orchestrator computes a path through the network and requests the resources (IP ports and spectrum slices) from the network through an SDN control platform such as OpenDaylight [3] or ONOS [4], which in turn establish connections through the IP and optical controllers [5]. In case the orchestrator cannot satisfy the intent due to insufficient network resources, it negotiates with the application until the intent can be satisfied. This mechanism forms the basis of our implementation of an *intent-based network* [6].

In such a dynamic situation, it is key that the orchestrator features a good resource allocation and network optimization component. In this paper we propose a modular algorithmic strategy to dynamic resource application jointly in both the IP and optical layer. Our scheme can then be implemented as part of an SDN orchestrator.

2. Resource Allocation and Optimization Framework and Modules

IP and optical layers differ in terms of (i) capacity provisioning, (ii) cost, (iii) power consumption, and (iv) connection setup speed. Our study focuses on (iii) and (iv). To establish a new optical connection, two transceivers have to be turned on, which consume a few hundred W of power. In addition, optical switch ports must be activated, although their energy consumption is typically a few W each. A lightpath has to be set up between the two transceivers, which can take tens of seconds. Although shorter setup times have been achieved, with the currently available equipment, optical and IP layer setup times are different by at least order of magnitude. Finally, note that a new IP layer connection will use only the previously-setup optical connections as its constituent IP link.

Because the IP layer is more time-responsive and power-friendly, whenever possible, we will accommodate a service request in the IP. Otherwise, if the demanded service requirements cannot be satisfied, we will resort to optical. This logic guides our resource allocation procedure, which is composed of three modules: IP Provisioning (IPP-M), IP Optimization (IP-OPT-M), and Optical Provisioning (OPP-M). IPP-M tries to satisfy the intent request by using available resources at the IP layer only. IP-OPT-M changes the existing routes in the IP. OPP-M adds new lightpaths, which are then used as IP links.

In Fig. 1, after a service request is received by the orchestrator, it tries to set up an application-aware path by using IPP-M. The path is then requested from the IP controller, unless there have been too many successful accommodations using IPP-M. In that case, or if IPP-M fails, IP-OPT-M is called to rearrange the IP layer connections to free up unneeded transceivers. This IP layer re-optimization can reduce power, or optimize the network for another objective. Otherwise, if IPP-M fails, the request is passed to IP-OPT-M, and if that fails, to OPP-M, which calls IP-OPT-M to select the paths for new optical connections.




Fig. 1. Resource allocation procedure modules.

In our study the service requirements are the two endpoints, bandwidth, maximum latency, and whether a service requires a dedicated lightpath. Next, we describe some algorithmic details of each module.

IPP-M first constructs an Auxiliary Graph (AG) by creating a copy of the IP layer topology, which consists of the sufficiently capacitated links. Next, we set the AG link weights to equal the physical length of the IP links (lightpaths). Then we produce candidate paths using K_{IP} -Shortest Path (SP) algorithm on the AG. The first path (if any) that meets each service requirement in the intent is selected. Finally, the demanded bandwidth is allocated over the traversed IP links.

IP-OPT-M tries to improve the routing over the network by minimizing a cost function. This is done by trying alternative routings for each established connection over the network and accepting changes only when the cost function is improved. This is a local optimization approach, in the sense that the solution found converges to a local optimum. However, since the connections over the network continuously change, a later call of IP-OPT-M will break out of this local optimum. The cost function tries to satisfy three main objectives: (i) that no traffic requests are unsatisfied, (ii) that the number of links is minimum, and (iii) that the utilization of the links is balanced. This is done using the following formula:

$$C_0 \sum_{\text{unrouted demands}} \text{request_size} + C_1 \times \text{number_of_links} + C_2 \sum_{\text{link}} \text{link_traffic}^2 \text{ with } C_2 \ll C_1 \ll C_0$$

Each term in the formula is designed to satisfy each of the above criteria, and in the same order. The cost is calculated in a way that guarantees that the values of the three terms never overlap, i.e. the minimum of the first term is higher than the maximum of the second term. C_2 can be set to one and the others adjusted accordingly. An important consideration in the algorithm selection for IP-OPT-M was that the changes suggested should be implemented on the network in a hitless way. In order to facilitate this, a list of changes is maintained in the order that they can be implemented, and only the changes that are possible without affecting other connections are considered. When the final, improved, state is determined, these changes are implementable on the network in the order specified without any service interruption.

OPP-M is similar to IPP-M. It constructs an AG from the IP topology but augments it by also considering *potential IP links*. For each node pair, a potential IP link is added if and only if there is no IP link for the pair in the AG. Each potential IP link is associated to a *potential lightpath*. To build such a lightpath, the K_{WDM-SP} First-Fit (FF) routing and spectrum assignment (we assume no wavelength conversion) is run. Then, when the K_{WDM-SP} FF algorithm has found free spectrum for the associated potential lightpath, the spectrum is reserved but not allocated yet. Next, the K_{IP-SP} algorithm is run over the AG and remove the candidate paths not meeting the service requirements. Next, the network, now consisting of the existing and remaining potential IP links, is sent to IP-OPT to find the optimal IP routes. Finally, the potential lightpaths that carry traffic are allocated the spectrum and turned on, and the bandwidth for the request is allocated in the IP layer.

3. Numerical results

To show the benefits of this approach, we implemented a scenario over the Spanish national core network with assumed elastic optical transmission (variable bit-rate, modulation format, and spectrum allocation). We specify three distinct traffic classes: high-priority, latency-sensitive, and best effort. For the high-priority class, requests are routed separately over the network, down to the optical layer. This class does not share resources with other classes and uses the minimum bit-rate modulation format. For the other two classes, resource sharing is allowed so the modulation format offering the largest bit-rate is selected, leaving space to accommodate other requests in the IP layer. The percentages of traffic in the three classes we chose for this study are 10%, 15% and 75% respectively.

Fig. 2(a) shows the progression of the number of IP links established in the network as a function of time, comparing the results we get with and without IP-OPT-M. Using IP-OPT-M results in a reduction around 10% in the IP links needed to serve the exact same request sequence. We also show in the same figure the performance of a hypothetical optimization module that simultaneously rearranges all connections on the network to find a better solution, but without guaranteeing a hitless way to move from one state to the other. However, we see that only a

small additional improvement is achieved with this non-realistic approach. Fig. 2(b) shows the measured latency distribution for all demands, showing that the service requirements are satisfied for the latency-sensitive classes.

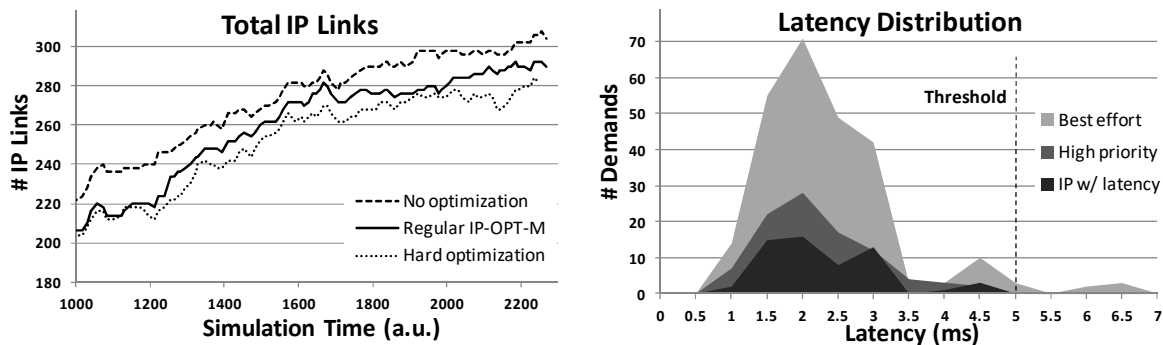


Fig. 2. (a) Number of IP links in the network with and without IP-OPT-M and comparing with an optimizer that ignores the requirement for hitless transition to the final state; (b) Distribution of actual latency values over the network for the three classes.

Another issue we considered was whether to remove idle lightpaths (to save power) or to leave them available (to quickly satisfy new requests). In Fig. 3 we show what happens with four different policies, which are the two extremes mentioned and the option to periodically remove idle lightpaths after either every 10 or every 20 calls to OPP-M. The periodic removal causes batches of lightpaths to be removed simultaneously (shown as breaks in the curve). However, even the slow periodic removal converges to the same minimum as in the “every round” case. At the same time, the need for establishing additional links varies with the removal frequency, making it a useful parameter for tuning the balance between energy efficiency and time responsiveness of the network.




Fig. 3. Progression of Total IP links and OPP-M added links over time for three different idle lightpath removal strategies.

In such a dynamic framework, the time responsiveness of the orchestrator is important. Our runtimes are tens of milliseconds for IP-OPT-M and low hundreds of milliseconds for OPP-M per request on an Intel i5-based PC. These are small fractions of the time needed to establish a new lightpath.

4. Conclusions

We proposed a network resource allocation and optimization framework based on the properties of the IP and optical layers, which can be implemented on a network orchestrator. Simulation on a test network with realistic application demands shows that our strategy provides reasonable resources for each application class. Periodic re-optimization of the IP routes helps reduce the power consumption of the network without causing service interruptions. In addition, infrequent removal of lightpaths to save energy was shown to be effective. The frequency of removal can be tuned so that the slow operation of lightpath setup is not needed too often.

The research leading to these results has received funding from the European Commission within the H2020 Research and Innovation programme, ACINO project, Grant Number 645127, www.acino.eu. We thank Dr. Victor Lopez of Telefonica I+D for providing network data.

5. References

- [1] Application-Centric IP/Optical Network Orchestration (ACINO), <http://www.acino.eu/>
- [2] V. Lopez, et al., "The role of SDN in application centric IP and optical networks", EuCNC, 2016
- [3] OpenDayLight, <http://www.opendaylight.org/>
- [4] ONOS, <http://onosproject.org/>
- [5] O. Gerstel, V. Lopez and D. Siracusa, "Multi-layer orchestration for application-centric networking", Photonics in Switching, 2016
- [6] ONF webinar, <https://www.sdxcentral.com/articles/contributed/network-intent-summit-perspective-david-lenrow/2015/02/>