

Application aware multilayer control and optimization of elastic WDM switched optical networks

Ioannis Tomkos^{(1)*}, Ćiril Rožić⁽¹⁾, Marco Savi⁽²⁾, Pontus Sköldström⁽³⁾, Victor Lopez⁽⁴⁾, Mohit Chamania⁽⁵⁾, Domenico Siracusa⁽²⁾, Chris Matrakidis⁽¹⁾, Dimitrios Klonidis⁽¹⁾, Ori Gerstel⁽⁶⁾

⁽¹⁾Athens Information Technology, Kifissias Avenue 44, 15 125 Marousi, Greece

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

⁽³⁾RISE Acreo AB, Isaffordsgatan 22, 164 40 Kista, Sweden

⁽⁴⁾Telefonica I+D/GCTO, Ronda de la Comunicación, Madrid, Spain

⁽⁵⁾ADVA Optical Networking, Fraunhoferstr. 9a, 82152 Martinsried, Germany

⁽⁶⁾Sedona Systems, 17 Hatidhar St. Raanana Israel;

*itom@ait.gr

Abstract: In dynamic networks with diverse application requirements, Software Defined Networking (SDN) principles enable application-aware in-operation planning. Project ACINO built a network orchestrator as the connecting component between network applications and the underlying network infrastructure.

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

1. Introduction

Backbone networks are undergoing a clear trend towards simplification. IP/MPLS technology became the current standard in packet core networks, while WDM/OTN networks with GMPLS control plane emerged as the transport network, by combining the scalability of WDM technologies with the dynamicity provided by a dynamic control plane. Even though IP/MPLS and WDM/OTN still represent two significantly different domains, previous work [1] demonstrates that a rational combination of packet and optical resources can deliver significant CAPEX savings. However, planning cycles in telecom networks are still based on single layer planning tools and manual operations. Network operators estimate service demands based on the inputs of their marketing department and the forecasted network traffic growth. The IP network design assesses if enough capacity exists in the network to accommodate the traffic under normal and failure conditions and the optical planning team designs the optical layer, finding a route for the optical spectrum requests over lightpaths and defining the positioning of transponders and regenerators. Once the overall design phase is completed, the deployment phase starts with the purchase of the hardware (typically new ports and transponders to increase the capacity), the installation and the validation.

Applications can require stringent constraints in terms of latency and jitter (e.g. financial transactions), bandwidth (e.g., video over IP), reliability, security, etc. Despite this diversity, traffic from multiple applications is aggregated and routed over the same path at the packet layer (since the operators want to re-use their multi-service networks) according to global parameters. This can easily lead to inadequate service offering or ineffective resource-utilization. In fact, if requirements are implicit and not known in detail by the transport network, optimization techniques are limited to few (relevant) parameters (e.g. bandwidth). As a consequence, transport pipes providing the bare minimum amount of resources potentially violate specific parameters that have been not taken into account. A smarter application-aware approach is required to offer a better adaptation of network services to applications needs.

Given the dynamicity of network operation and the diversity of application service requirements, it is natural to employ Software Defined Networking (SDN) to enable *application-aware in-operation planning*. In project ACINO [2], we use SDN principles to build a *network orchestrator* as the connecting component and intelligence between the network applications and the underlying network infrastructure. As shown in Fig. 1, the orchestrator first receives the service requirements from the application and a current picture of the underlying IP/optical network, and then computes an application-aware path through the network.




Fig. 1. ACINO SDN orchestrator architecture and communication with network applications and infrastructure.

2. Application-aware in-operation planning

In-operation network planning [3] refers to the ability of automatically reconfiguring or re-optimizing the network in near real time. It is an online process, unlike traditional planning processes that are performed offline over long time-scales. In this process, the knowledge of the network topology and status is exploited to reactively adapt to network changes, such as network failures (re-configuration), or to proactively re-arrange IP and optical network connections according to operators' policies (re-optimization). The main benefit of in-operation network planning is the ability to better tune the state of the network to the real short-term needs of its traffic, rather than aiming to be ready to serve a higher amount of traffic that may not materialize in several months.

Current approaches [3] are bandwidth-driven, i.e., they focus only on the capacity needed to support short-term traffic demands. This is sufficient to achieve CAPEX and OPEX savings in best-effort network scenarios, but also limiting in terms of the characteristics of the services being offered to the applications that ultimately drive the network traffic.

In contrast, application-aware planning considers the optimum allocation of IP and Optical layer resources according to the applications' *service requirements* that must be met through a joint IP-Optical multi-layer resource allocation process. This resource allocation process aims to (i) accommodate network application service requests as quickly as possible, (ii) optimize the use of network resources, including power, and (iii) be implementable as part of an orchestrator deployed on a real network. At the same time, we observe that in IP/optical networks the two layers differ in terms of (a) capacity provisioning, (b) cost, (c) power consumption, and (d) connection setup speed. Regarding (a), at least some applications do not necessarily require their own optical channel, so (iv) support for grooming application flows in the IP layer should be supported. Next, (b) can be handled implicitly by choosing optimization objectives so that they directly affect the network cost. Regarding (c) and (d), to establish a new optical connection, power-consuming equipment such as linecards, transceivers and transponders have to be turned on. By contrast, a new IP layer connection (e.g. an MPLS path) will use only previously-setup optical connections as its constituent IP links. The extra power needed to send packets down existing IP links is much smaller than the power required to set up new ones. IP layer connections are thus more (v) time-responsive and (vi) power-friendly.

Given the above considerations, an application-aware approach to in-operation planning, where explicit application requirements are known to the control plane of the network, can be useful to safeguard the service characteristics that are critical to applications, while still achieving significant savings for operators [4].

3. A framework for application-aware dynamic multi-layer resource allocation and optimization

The ACINO orchestrator is based on several open source projects which have been extended to support application-aware operations and planning. The main component is the Open Network OS (ONOS) [5], an open source SDN controller (left in Fig. 2). It is responsible for interacting with applications that require network resources and determining their needs (through a REST API), and control the underlying IP/MPLS and optical network elements, either directly or via technology/vendor specific controllers (using *Providers*). It provides many other functions as well, e.g. a novel *intent*-based mechanism to describe service requirements, e.g. to create a path between two nodes in the network. It also provides all the functionalities expected from an SDN controller, such as network topology discovery, GUI and CLI for controlling the network, etc. ONOS provides a Java API which lets third parties develop modular applications that can be plugged into the controller to provide additional functionality and interact with the core of the controller.

The other main component we use is Net2Plan [6] (to the right in Fig. 2), an open source network planning tool. Net2Plan provides an extensive framework to develop and test algorithms and can be extended through a plugin system that lets third parties develop additional features or algorithms that can be dynamically loaded by the main application. It also provides both CLI and GUI to let the user visualize and modify topology, routes, and demands in the network, as well as execute algorithms.

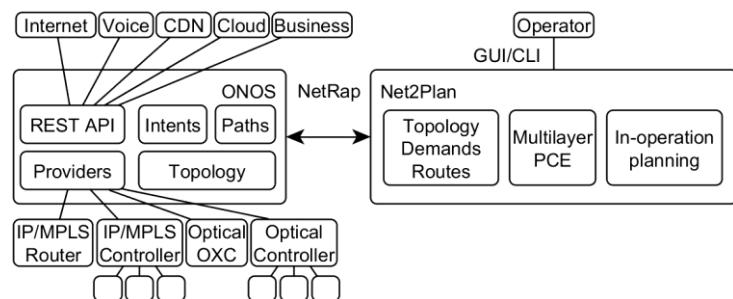


Fig. 2: Implemented network orchestration architecture with ONOS and Net2Plan.

Between the two components we employ [7] a REST-based protocol called NetRap, which synchronizes the topology, traffic demands, and routes between the two components. It also lets ONOS request path calculations from Net2Plan and retrieve the resulting path(s). Net2Plan can use it to directly push new paths or re-configure existing ones in ONOS, without a preceding request from ONOS, to execute

calculated network plans. This allows Net2Plan to act both as a multi-layer path-computation engine (PCE), calculating optical and IP/MPLS paths on-demand as service requests (referred to as intents) arrive to ONOS, and as an in-operation planning tool controlled manually by an operator through the built-in GUI. This also brings additional benefits; algorithms developed for planning simulations can be directly used to control the network, bridging the gap between simulation and reality. Conversely, any modifications made to algorithms to deal with e.g. hardware constraints can go directly into simulations, making them more realistic.

The application-aware in-operation planning multi-layer resource allocation and optimization framework is designed following the considerations given in the previous section. Considerations (i), (iv), (v) and (vi) dictate that we accommodate service requests in the IP layer whenever possible. Otherwise, if the demanded service requirements cannot be satisfied by IP, we will resort to optical. Considering (ii), we will purposely optimize the network, but, considering (i) and (vi), not move established lightpaths around. Considerations (i) and (iii) call for a fast, rather than exhaustive and complex, procedure. Finally, (iii) calls for hitless (make-before-break) schemes, or else accommodating new applications may disrupt the running ones.

The framework, implemented in Net2Plan, is depicted in Fig. 3 and is composed of three modules: IP Provisioning (IPP), IP Optimization (IP-OPT), and Optical Provisioning (OPP), jointly implementing the “multi-layer PCE” and “in-operation planning” features shown in Fig. 2. In Fig. 3, after a service request is received by the orchestrator, it tries to set up a path by using IPP. The path is then requested from the IP controller, unless there have been too many successful accommodations using IPP. In that case, or if IPP fails, IP-OPT is called to rearrange the IP layer connections. In addition to making space for the new service requests, this IP layer re-optimization can reduce power, or optimize the network for another objective. If IP-OPT fails, the request is passed to OPP, which sets up new lightpaths (IP links). In doing that, OPP may call IP-OPT to help select the paths for new optical connections.

Using the same framework with the Net2Plan event generator to generate requests, we run simulations on a test network with realistic application demands [8] showing 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 in terms of the resources savings and network responsiveness. The frequency of removal can be tuned so that the slow operation of lightpath setup is not needed too often. In other tests we also found a significant energy, cost, and responsiveness difference between a network offering separately routed connections, and one that has a less exclusive service where packets can be mixed together. Such findings may help a network operator decide which services to offer.

Acknowledgments: This work was supported by the European Community’s H2020 Research and Innovation program, within the ACINO project, Grant Number 645127.

4. References

- [1] O. Gerstel, et al.: Multi-Layer Capacity Planning for IP-Optical Networks, in IEEE Communications Magazine Feature Topic “Advances in Network Planning”, January 2014, Vol. 52, Issue. 1, pp. 44-51.
- [2] F. Pederzoli, et al.: SDN application-centric orchestration for multi-layer transport networks (Invited), in proc. IEEE International Conf. on Transparent Optical Networks ICTON, paper Th.A4.2, Trento, Italy, July 2016.
- [3] L. Velasco, D. King, O. Gerstel, R. Casellas, A. Castro and V. López: In-operation Network Planning, in IEEE Communications Magazine Feature Topic “Advances in Network Planning”. January 2014, Vol. 52, Issue. 1, pp. 52-60.
- [4] V. López, D. Klonidis, D. Siracusa and J.P. Fernandez-Palacios: On the Benefits of Multilayer Optimization and Application Awareness, in Journal of Lightwave Technology, Volume: 35, Issue: 6. PP. 1274-1279. March. 2017.
- [5] ON.LAB: ONOS, visited 2017-04-19, *online*: <http://onosproject.org/>
- [6] P. Pavon-Marino, et al.: Net2plan: an open source network planning tool for bridging the gap between academia and industry. IEEE Network 29.5, 2015.
- [7] P. Sköldström et.al. Making powerful friends: Introducing ONOS and Net2Plan to each other (Invited) ICTON, July 2017 (paper We.C3.2)
- [8] Č. Rožić, et al.: A Framework for Dynamic Multi-layer Resource Allocation and Optimization in Application-Centric Networking, in Proc. Optical Fiber Conference (OFC), Mar 2017 (paper W3I.4).

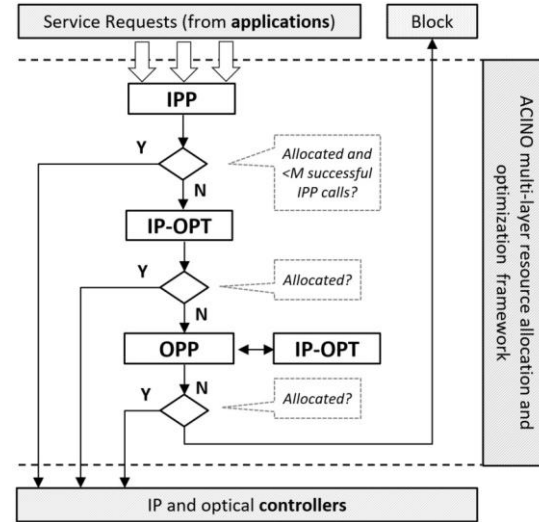


Fig. 3. Resource allocation and network optimization procedure.