Disaggregating Optical Nodes in a Multi-Layer SDN Orchestrator for the Integration of an In-Operation Planning Tool

Federico Pederzolli[1], Mohit Chamania[2], Michele Santuari[1], Thomas Szyrkowiec[2], Chris Matrakidis[3], Ćiril Rožić[3], Dimitrios Klonidis[3], Victor Lòpez[4], Domenico Siracusa[1]

[1] *FBK CREATE-NET, Via alla Cascata 56/D Povo, Trento, Italy;*
[2] *ADVA Optical Networking, Fraunhoferstr. 9a, 82152 Martinsried, Germany;*
[3] *Athens Information Technology (AIT) center, 44 Kifissias Av., 15 125, Marousi-Athens, Greece;*
[4] *Telefònica I+D/GCTO, Ronda de la Comunicacin, Madrid;*

fpederzolli@fbk.eu

Abstract: Optical disaggregation can provide the intermediate models required by In-Operation planning to compute feasible configurations in IP/Optical networks. We demonstrate disaggregation on a real SDN-orchestrated testbed, and quantify its benefits and costs.

OCIS codes: (060.4256) Networks, network optimization; (060.4258) Networks, network topology

1. Introduction

The advent of SDN has disrupted the traditional network provisioning and planning ecosystem with the introduction of dynamic service provisioning and, more recently, automated network operation. Traditional network operations work on largely static semi-permanent demands and, as a consequence, network planning typically deals with long operation time scales (~1-2 years) and the provisioning of services in this timescale uses existing (typically over-provisioned) network resources. As we move towards dynamic services with increasing capacity demands, operating over-provisioned static network infrastructure is not an economically viable option.

To address challenges with dynamic services in typical multi-layer networks, *in-operation planning* [1] was introduced as a concept to dynamically compute, validate and deploy batch changes to network configuration in order to optimize network metrics (installed services, link utilization, etc.). In order to perform in-operation planning, a multi-layer in-operation planning tool must be able to collect information about the network configurations, services, and provide potential solutions for migrating from the current to a more efficient configuration. *Disaggregated models*, where single nodes are substituted with subgraphs consisting of (potentially) multiple nodes and links, designed to mimic inherent constraints (shared resources, capacity constraints, etc.) of each original node, are an ideal method to convey the complex constraints imposed by optical hardware.

In this paper we present the first, to the best of our knowledge, implementation of the automatic disaggregation of a real network, and use it to transparently encode ROADM restrictions at the interface between an external multi-layer in-operation planning tool and an SDN multi-layer network orchestrator [2]. We detail the challenges we faced implementing it, and show that, despite adding some overhead, it considerably improves the accuracy of in-operation planning results.

2. Challenges for In-Operation Planning in SDN-controlled multi-layer networks

In a multi-layer Packet/Optical scenario, in-operation planning can be used to firstly re-optimize the packet layer, operating possibly bulk changes to accommodate new service requests, and if required, re-optimize the configuration of the optical layer, moving or establishing new lightpaths to better suit a new overall configuration. As it has been presented in [3], introducing in-operation planning can improve energy consumption without impacting service availability.

To produce feasible solutions, a multi-layer in-operation planning tool requires detailed knowledge about the capabilities and constraints of the underlying network layers; without, for example, knowledge about directionality or contention issues within optical nodes, it is likely that a significant fraction of configurations computed by the planning procedure would be unfeasible in practice (due to e.g. wavelength contention). This is a challenge, since SDN orchestrators prefer to use generic network element representations, while pushing complexity of the device internals, especially in the optical layer, down to the device/sub-network controllers.

The inherent constraints of a device can be represented using a disaggregated model, which are ideal for interfacing external tools with a multi-layer SDN orchestrator because they can easily express intricate constraints, such as fixed filters, directionality or contention issues in ROADM configurations, without requiring any external data structure or special operation. In addition, the representation does not need to be altered to reflect the evolution of the network's state (e.g. color-based connectivity in a connectivity matrix), as that is already part of the graph. Finally, such a model makes representing complex networks running a mixture of newer and older nodes and technologies easy to both create and parse.

3. Implementation and System Integration

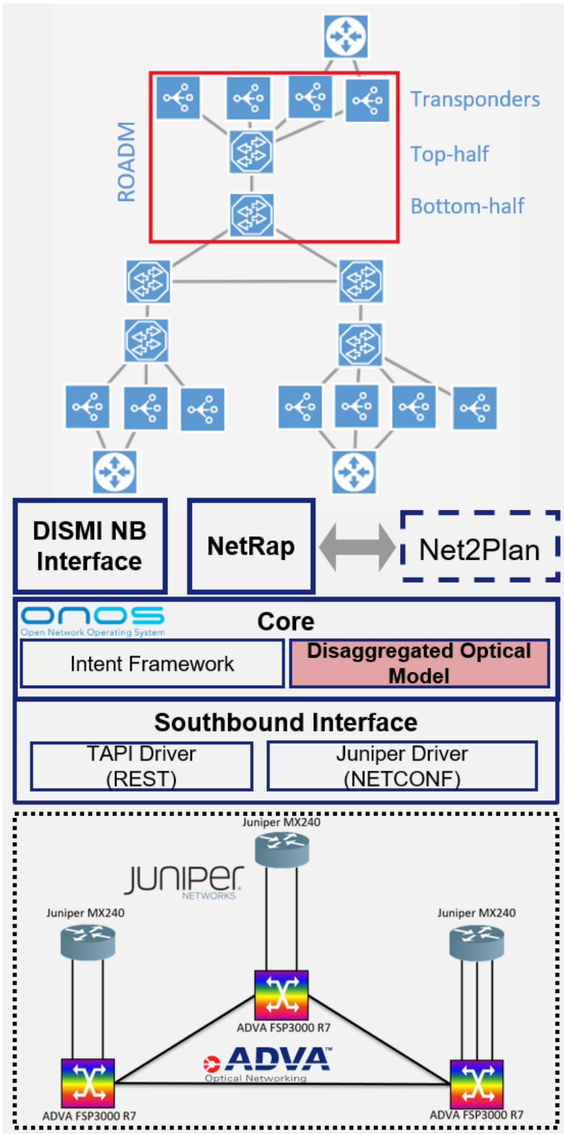

Fig. 1: Orchestrator architecture (center), testbed (bottom) and disaggregated (top) topologies.

We implemented our in-operation planning algorithms, described in more detail in [3], in an open-source network simulation tool called Net2Plan (N2P) [4]. Our ONOS-based network orchestrator is outlined in [2] and in Fig. 1. At the top, the DISMI interface [5] handles service requests from customer applications, while NetRap [6] handles the translation of ONOS' flat model into Net2Plan's multi-layer one and encodes service requests coming from DISMI. Within the ONOS core, an enhanced intent framework handles the translation of service requests fed back by N2P into multi-layer configuration, which is installed in our testbed network composed of Juniper Routers and ADVA ROADMs. These are, respectively, controlled directly by the orchestrator via a NETCONF driver, and indirectly via a Transport API (TAPI) [7] interface towards ADVA's Network Management System (NMS).

In order to model a wide variety of constraints over NetRap, rather than defining complex and custom data structures for each possible case we chose to encode them directly into the exchanged topology graph. As shown on top of Fig. 1, constrained ROADM nodes are disaggregated into sub-graphs, which expose the inherent constraints of the device. Consider the highlighted section in the figure, representing a single directionless ROADM. Each transponder (4 in total) is represented as a single node at the top, and the colors used by each transponder are defined in its ports. Two optical switches, connected by an optical link are used to mimic constraints on add-drop configurations, where no two transponders can use the same wavelengths.

Note that the internal representation (disaggregation mechanism) of the nodes can be specific to an SDN implementation, therefore, we modified the ADVA NMS to annotate the standard representation in TAPI (single node per optical device) with additional parameters indicating the device's nature/capabilities (e.g. "CD" or "CDC") and ports. This object is passed to ONOS, where it is processed by a disaggregator module that uses the annotations to transform a single node into multiple annotated devices representing transponders, top and bottom halves of a ROADM, etc. (see red square at the top of Fig. 1). Each of these is stored into ONOS' network model, which further requests back through the disaggregator and TAPI the details of each individual device, including e.g. client and line ports, and lastly configures the intermediate links between all new sub-objects. Observe that this bootstrap procedure, while lengthy, depends exclusively on the (average) number of sub-objects within a monolithic node.

Once ONOS' model is initialized, in-operation planning requests are forwarded to N2P along with the disaggregated topology (using incremental changes to improve performance) via NetRap. There, no modification to the planning algorithms is required to handle diverse optical nodes, as their properties are encoded in the network graph, although the higher number of nodes does lead to a slight increase in computation time.

4. Evaluation of disaggregation and results

Intuitively, the decomposition of nodes into complex sub-graphs, with each new node in the sub-graph representing a network element that is individually managed by the SDN orchestrator, requires a non-negligible overhead compared to simply treating a node as a monolithic entity managed by an underlying controller/NMS. To quantify this, we measured the time needed to perform device and port discovery using both our disaggregator and the regular integrated mode of ONOS, on our testbed network (depicted at the bottom of Fig. 1). We performed 100 measurements, and reported the resulting median, 1st and 3rd quartiles, minimum and maximum times in Fig. 2 (left). Despite the fact that port

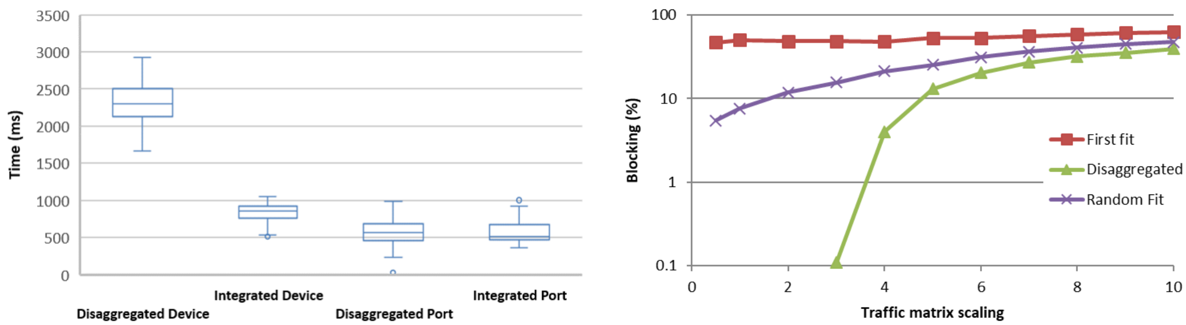

Fig. 2: Disaggregated vs. Integrated device and port discovery times (left), and Blocking Probability vs. load for different wavelength assignment algorithms and network models (right).

discovery is mostly unaffected by the disaggregation process, the device discovery phase is about 2-3 times slower when disaggregating (due to the feedback and the fact that ONOS polls device details sequentially), although the process still only requires less than 3 seconds to discover the full testbed topology (Fig. 1, bottom).

Despite this limited overhead, by using the disaggregated model we can avoid the planning of unfeasible configurations. To quantify this advantage, we ran our multi-layer in-operation planning algorithms [3] on the Spanish National Topology comprising 30 nodes and 56 bidirectional links [8], assuming that CD ROADMs are used at every node, 80 channels are available at each fiber and 100 Gb/s transceivers are used.

Fig. 2 (right) depicts the percentage of blocked demands due to both lack of resources and invalid path-lambda combinations as a function of the scaling factor used to multiply the traffic matrix from [3]. Three scenarios are tested, two that ignore the contention limitations using either “first fit” or “random fit” wavelength assignment and the “disaggregated” model that correctly represents ROADM limitations. As expected, if CD ROADMS are used without properly encoding the resulting contention limitations in either the model or the routing algorithm, then even at low loads a large fraction of demands is assigned an unfeasible path and lambda combination, resulting in needless blocking. This is very pronounced when the commonly used “first fit” wavelength assignment is used, since the lower numbered wavelengths will be used first, leading to close to 50% blocking due to contention, while “random fit” reduces it somewhat since it uses a wider range of colors. In contrast, the disaggregated model has considerably lower blocking, only due to lack of resources. Note that the wavelength assignment method used with the disaggregated model has no measurable effect in this test, so only the “first fit” results are shown.

5. Conclusions

We described a multi-layer SDN orchestrator implementing disaggregation to expose optical constraints to external applications. We showed that operating under assumptions of unconstrained switching in otherwise constrained optical networks leads to significant penalties in terms of blocking, which can be avoided by the use of disaggregation, at the cost of a small time penalty during discovery (bootstrapping) of the network topology.

Acknowledgement: 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.

References

1. L. Velasco et al., “In-operation Network Planning”, in IEEE Comm. Mag. 2014.
2. M. Chamania et al., “Intent-based in-flight service encryption in multi-layer transport networks,” OFC 2017.
3. Ć. Rožić et al., “A framework for dynamic multi-layer resource allocation and optimization in application-centric networking,” OFC 2017.
4. P. Pavon-Marino et al., “ Net2plan: an open source network planning tool for bridging the gap between academia and industry”, IEEE Network 2015.
5. P. Sköldström et al., “DISMI - an intent interface for application-centric transport network services,” ICTON 2017.
6. P. Sköldström, et al., “Making powerful friends: Introducing ONOS and Net2Plan to each other,” ICTON 2017.
7. Open Networking Foundation, “Functional Requirements for Transport API,” ONF TR-527 (2016);
8. F. Rambach et al., “A Multilayer Cost Model for Metro/Core Networks”, IEEE/OSA JOCN 2013.