Integrated OpenFlow–GMPLS Control Plane: An Overlay Model for Software Defined Packet over Optical Networks

Siamak Azodolmolky, Reza Nejabati, Eduard Escalona,
Ramanujam Jayakumar, Nikolaos Efstathiou, Dimitra Simeonidou
High Performance Networks Group (HPNG), School of Computer Science and Electronic Engineering, University of Essex
Wivenhoe Park, Colchester, Essex, CO4 3SQ, UK
Email: {sazod, rnejab, eescal, jayaka, nefsta, dsimeo}@essex.ac.uk

Abstract: A novel software-defined packet over optical networks solution based on the OpenFlow and GMPLS control plane integration is demonstrated. The proposed architecture, experimental setup, and average flow setup time for different optical flows is reported.
OCIS codes: (060.4250) Networks; (060.4510) Optical communications

1. Introduction

Future Internet is characterized by global delivery of packet switched traffic driven by high-performance network-based applications such as ultra high definition (UHD) video on demand streaming and cloud computing. Bandwidth demands resulting from traffics generated by these types of applications, when aggregated for transport over core networks can be met only with high-capacity WDM circuit switched optical networks. A key challenge for network operators to support future Internet applications, is the efficient end-to-end delivery of packet switched traffic in terms of network operation, control and management when traversing from a packet switched domain (i.e. campus and metro) through a circuit switched domain (i.e., WDM optical core network).

Generalized Multi Protocol Label Switching (GMPLS) developed in IETF as a generic network control plane framework is used for managing physical path and core tunnelling technologies of the Internet and Telecom service providers. The GMPLS architecture extends MPLS (Multi Protocol Label Switching) to encompass control and management of time-division (e.g., SONET/SDH, PDH, G.709), wavelength, and spatial switching (e.g., incoming port or fibre to outgoing port or fibre). The GMPLS control plane, due to its support for various optical transport technologies as well as its capability for dynamic and on demand lightpath provisioning, is widely being considered by operators as the control plane of their next generation core optical networks.

The OpenFlow paradigm has recently been proposed as a control framework that supports programmability of network functions and protocols (i.e., software defined networking) by decoupling the data plane and the control plane, which are currently vertically integrated in many networking equipments (e.g., routers, switches, access points) [1]. OpenFlow adopts the concept of flow based switching and network traffic control for intelligent, user controlled and programmable network service provisioning with the capability to execute any user defined routing, control and management application in software outside the data path, in the OpenFlow controller. OpenFlow technology is being adopted by most of the major packet switching vendors as the key technology enabler for realization of software defined networking and there are currently a number of products available from different vendors supporting OpenFlow technology especially on campus and metro class products.

OpenFlow control framework is a promising technology for integrating the control and management of packet switched domain and optical circuit switched domain. It provides a framework for development of innovative functionalities and protocols thanks to its support for software defined networking. There have been several attempts and proposals to control both circuit switched and packet switched networks using the OpenFlow protocol [2,3]. In the work proposed in this paper, a novel software-defined packet over optical networks solution is presented, which is enabled by the integration of OpenFlow and GMPLS control planes. It describes the control plane architecture, workflow and a proof of concept experimental demonstration of the proposed solution. To the best of our knowledge, this is the first time that integration of OpenFlow and GMPLS control plane for software-defined packet over optical networks has been proposed and experimentally demonstrated.

2. Integrated OpenFlow-GMPLS control plane

The underlying principle of OpenFlow is to treat traffic as flows and to have the control functionality taken out of the networking equipments to a centrally managed or a distributed controller (i.e., OpenFlow controller), while retaining only data plane on the equipment. In the OpenFlow control framework a network is managed by a network-wide operating system running on top of a controller that controls the data plane with OpenFlow protocol. The OpenFlow controller is a server, which has the capabilities to host different network management and control applications to effectively manage the network in a centralised or distributed way. OpenFlow abstracts each data plane switch as a flow-table. The control plane makes decisions as to how each flow is forwarded (reactively as new flows are detected, or proactively in advance), then caches its decision in the data plane's flow-table. In a packet switched network a flow can be defined in a flexible way as a combination of any L2, L3, L4 headers of a packet as shown in Figure 1a.

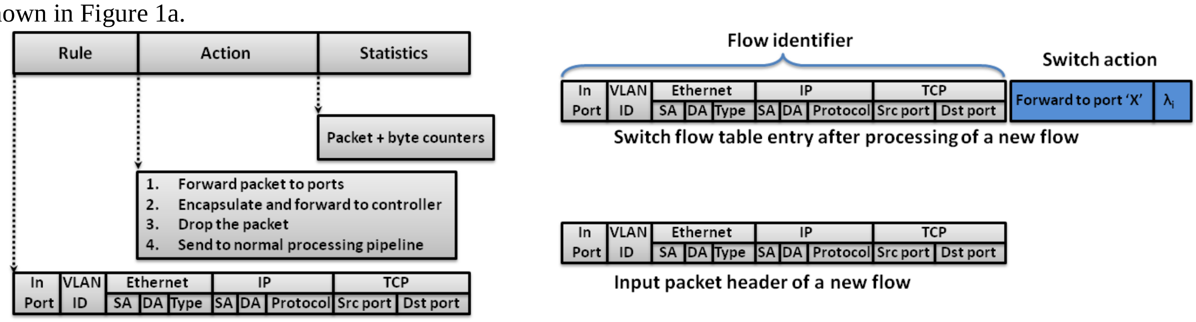

Fig. 1: a) Flow table structure within an OpenFlow switch (left), b) Extended flow table entry (right).

In an OpenFlow controlled packet switched network, incoming packets are matched against the flow definitions; if there is a match, a set of actions are performed, and statistics will be updated. Packets that do not match any flow-table entry are (typically) encapsulated and sent to the controller. The controller can decide how to process the packet and then insert its decision in the data plane (i.e., a new rule in the flow-table). Thus, consequent packets in the flow are processed according to this new rule, without contacting the controller. Therefore, while each packet is switched individually, the flow is the basic unit of manipulation within the switch. Note that by switch we are refereeing to a generic OpenFlow enabled switch. The fields of the flow table are shown in Figure1a. The rule field defines the fields from the header to be matched in the packet. Action defines the way the packet should be treated depending on the rule. Statistics field is used to gather packet statistics, which are used by the network applications to make dynamic and/or automatic decisions.

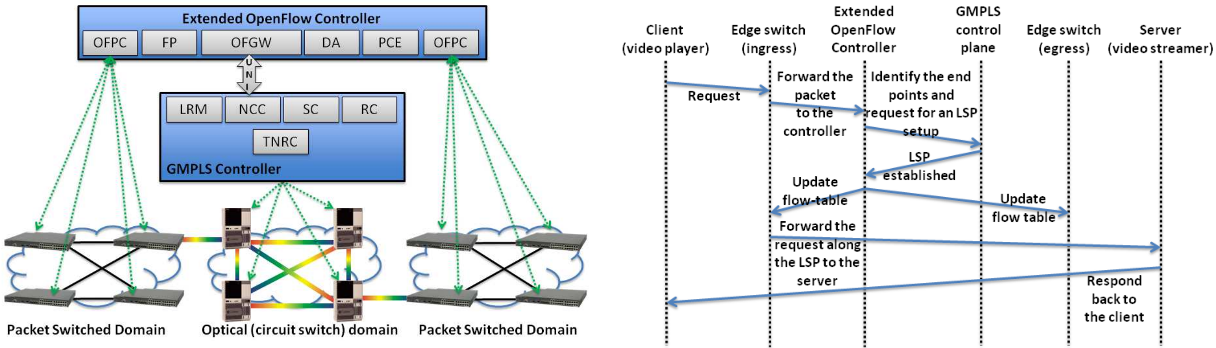

Fig. 2: a) Integrated OpenFlow-GMPLS control plane (left), b) Timing diagram for the experimental demonstration (right).

Figure 2a shows the architecture of the proposed integrated OpenFlow-GMPLS control plane. It utilizes an extended packet switch OpenFlow controller integrated with a GMPLS control plane in an overlay model. The GMPLS control plane follows the standard ASON model and its building blocks include: the network connection controller (NCC) responsible for handling and processing the connection requests, the signalling controller (SC), which implements the RSVP-TE protocol and is responsible for handling the GMPLS signalling, the routing controller (RC) comprising OSPF-TE protocol and a path computation algorithm for calculating the end to end path, the link resource manager (LRM) which is responsible for monitoring and information collection regarding status of network elements and the transport network resource controller (TNRC) which provides the interface between GMPLS and network elements for their configuration and monitoring. The extended OpenFlow controller comprises: the flow processor (FP), which is responsible for processing new flows and creating flow rules and updates for flow tables in switches, the path computation element (PCE), responsible for path computation for each flow within each packet switched domain, the discovery agent (DA), which is responsible for discovering the network topology and connectivity including end points in each packet switched domain and the OpenFlow Gateway (OFGW) that provides the interface between the OpenFlow controller and the GMPLS control plane through a user networl interface (UNI) signalling interface. OpenFlow protocol controller (OFPC) is responsible for interfacing Openflow Controller and OpenFlow enabled switched for flow table configuration and monitoring of switches.

In the proposed model the packet switched domain and the GMPLS optical domain are controlled by an extended OpenFlow controller. The extended controller provides functionality for requesting optical connectivity from the optical domain. It is assumed that the edge packet switch, which interconnects the optical domain to the packet switched domain, is equipped with tuneable WDM interfaces and is connected to an add/drop port of the ingress/egress nodes of the optical domain as shown in Figure 2a. Initially the flow tables of the switches are empty. The first packet of the request will be forwarded to the extended OpenFlow controller. Based on the destination address of the request, the extended OpenFlow controller identifies the end points of the lightpath in the circuit switched domain and requests a lightpath from GMPLS control plane using UNI. Once the lightpath is established the extended controller then updates the flow table of the ingress and egress switches, finalizing the establishment of an end-to-end optical flow path. The extended OpenFlow controller also maintains a list of established lightpaths in terms of lightpath Identifiers. The corresponding timing diagram for this process is depicted in Figure 2b.

3. Experimental setup, demonstration scenario and results

The experimental setup is shown in Figure 3a. It comprises two OpenFlow enabled L2 packet switched domains and one optical fibre switched domain. One of the packet switched domains (client side), is equipped with three NEC IP8800 switches and the other one, the server domain, is equipped with one NEC IP8800 switch. The optical switching domain comprises 4 Calient DiamondWave optical switches. Each optical switch is controlled by a GMPLS controller and the GMPLS control plane is distributed among four GMPLS controllers. The client and server packet switch domains are controlled by the proposed extended OpenFlow controller, which is connected to the GMPLS control plane through UNI interface. For simplicity the two NEC switches at the boundary of optical domain are equipped with fixed wavelength transponders at different wavelengths and each port of the optical switch carries only one wavelength.

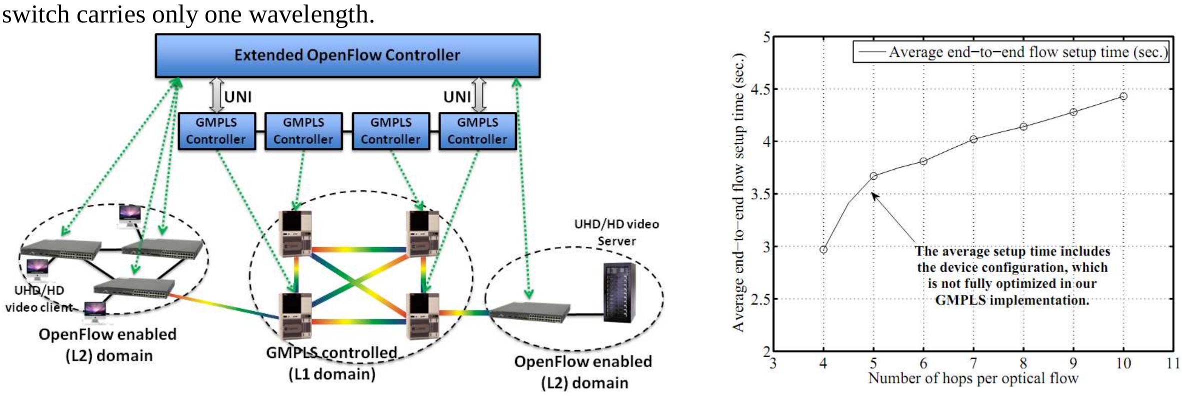

Fig. 3: a) experimental test-bed setup (left), b) Average end-to-end flow setup time vs. number of hops per optical flow (right).

Once a client sends a new request for video streaming (connectivity to the video server) to one of the L2 switches in client domain, the switch will not be able to find a flow entry for the new request and forwards the request to the extended OpenFlow controller. This controller resolves the destination address (i.e., server domain) and requests from GMPLS via UNI a new optical lightpath between client and server domains. GMPLS returns the acknowledgment for the established LSP with ingress port/wavelength and egress port/wavelength. The OpenFlow controller updates the flow tables of the switches in the client and server domains with appropriate port for the new connection. It also maintains the lightpath identification. Finally the client receives acknowledgment of the request and the video connectivity (streaming) will be established. The interaction of OpenFlow and GMPLS control plane extends the benefits of the software defined networking to the circuit switched (optical) domain. For instance, the extended OpenFlow controller can route the flows over different paths based on different level of quality of service. The average end-to-end flow setup time for different optical flow paths are depicted in Figure 3b. As the number of hops per optical flow increases, more time is required to establish an end-to-end optical flow path. The setup times include both signalling and device configuration delays, which is not fully optimized in our GMPLS implementation.

4. Conclusions

In this work, we experimentally demonstrated the integration of OpenFlow and GMPLS control planes. The proposed overlay model extends the functionality of a typical OpenFlow controller in a way to properly interface with GMPLS control plane. We also reported on the end-to-end flow setup time as a function of number of hops per optical flow path. This work is partially supported by EU FP7 funded project OFELIA (www.fp7-ofelia.eu).

5. References

[1] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, J. Turner, “OpenFlow: enabling innovation in campus networks,” ACM SIGCOMM Comp. Comm. Rev., vol. 38, no. 2, pp. 69-74, April 2008.

[2] S. Das, G. Parulkar, P. Singh, D. Getachew, L. Ong, N. McKeown, “Packet and Circuit Network Convergence with OpenFlow,” Optical Fiber Conference (OFC/NFOEC10), San Diego, March 2010.

[3] V. Gudla, S. Das, A. Shastri, G. Parulkar, N. McKeown, L. Kazovsky, “Experimental Demonstration of OpenFlow Control of Packet and Circuit Switches,” Optical Fiber Conference (OFC/NFOEC’10), San Diego, March 2010.