E-WDM: Wavelength Switching in Hybrid Networks Without Wavelength Selective Switches

G. Guelbenzu[1], W. Miao[1], N. Calabretta[1], O. Raz[1]

[1] Institute for Photonic Integration, Eindhoven University of Technology, Netherlands, G.Guelbenzu@tue.nl

The use of wavelength multiplexing in optical circuit switches effectively multiplies the capacity per port. However, switching groups of wavelengths together may lead to underutilized links. The solution is the introduction of wavelength switching with granularity at wavelength level. In this paper we present a novel solution, E-WDM. It is a SDN control technique that utilizes the electronic switches in hybrid data centers to dynamically assign wavelengths to the high capacity links of optical circuit switches. The technique is demonstrated in our hybrid setup, integrating central controller, servers, electronic switches, and our 1.5 ms centralized reconfiguration time fast OCS.

Introduction

Optical circuit switches (OCS), in topologies such as *Helios*[1] and *C-Through*[2], are able to switch groups of *m* wavelengths among *n* ports, effectively connecting $m{\cdot}n$ devices with an *n* port switch. This leads to a reduction by *m* of the number of fibers and optical switches required. However, if those groups of *m* wavelengths are statically assigned to groups of *m* servers, the traffic patterns are limited, and the high-capacity WDM (Wavelength Division Multiplexing) links may be underutilized. For instance, the *m* servers of one group must communicate with the *m* servers of other group, and servers can only communicate with other peers having the same assigned wavelength.

The solution to these problems of OCS exploiting WDM links is the dynamic assignment of wavelengths to the links. Several approaches have been suggested, usually presenting a different implementation of a WSS (Wavelength Selective Switch). For instance, a common alternative is build a WSS based on broadcast and select networks, and SOA (Semiconductor Optical Amplifier) arrays[3]. Other option builds wavelength switching cards with an optical cross-connect, and an array of lasers with fixed wavelengths[4,5]. Others require AWG (Arrayed Waveguide Gratings), TWC (Tuneable Wavelength Converters), and TL (Tuneable Lasers)[6,7].

In this paper we demonstrate a different approach using FOX (Fast Optical Circuit Switch), and E-WDM. Our FOX is a SOA-based fast OCS operating at 1300 nm. The number of ports of FOX is effectively duplicated by connecting two wavelengths to each port. It constitutes the core optical unit of our hybrid 10G data center demonstrator, providing 20G WDM links from side-to-side of the data center. E-WDM is a control technique to dynamically assign the wavelengths of the high-capacity links in hybrid data centers exploiting WDM. E-WDM does not require additional and costly WSS in the data center. It achieves wavelength switching by tight integration of electronic packet switches (EPS), OCS and servers by means of a central SDN controller. As a result, the utilization of the high capacity links is optimized for flexible traffic patterns that can be switched in around 70 ms, with the FOX being controlled in only 1.5 ms.

Hybrid Data Center Demonstrator

A diagram of our demonstrator, inspired by Helios, is shown in Fig. 1. The eight servers are (virtually) organized in four racks. Each rack has a top of the rack switch, and two servers. A second layer of electronic switches aggregates the traffic of several racks into clusters, and converts the single-wavelength 10G links of the EPS tier to double-wavelength 20G WDM links. Finally, the third layer of optical switches constitutes the optical core performing the inter-cluster connections. In this case, FOX is the only switch of the optical core, but a more realistic scenario would include several optical switches. Our demonstrator is a fully centralized data center, where the SDN central controller orchestrates the behaviour of EPS, FOX, and servers.

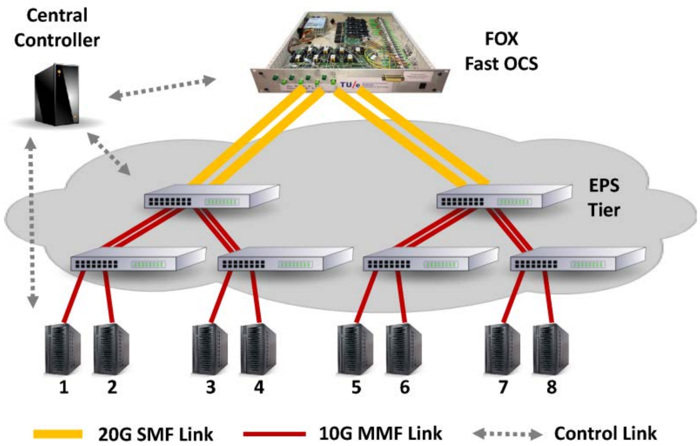

Fig. 1: Hybrid Data Center setup

E-WDM: Flexible Wavelength Switching

In a static configuration scenario, servers 1, 3, 5, 7 could have been assigned one of the 10G wavelengths, and servers 2, 4, 6, 8 could have been assigned the other 10G wavelength. With such configuration, then servers 1-2 could communicate with servers 5-6 or 7-8 depending on the FOX configuration. Servers with different wavelengths could not communicate. For instance, the communication of server 1 with servers 6 or 8 would not be possible.

Fig. 2: E-WDM examples

E-WDM utilizes the electronic switches in hybrid networks to dynamically assign wavelengths to the servers. In order to do that, the central controller installs different OpenFlow rules in the EPS according to the desired traffic pattern. For instance, Fig. 2 left shows an example where the first layer of EPS is configured to perform intra-group wavelength swapping, interchanging the wavelengths assigned to servers 1-2, and 3-4.

As a result, servers 1-2 communicate with servers 6-5 (or 8-7), and servers 3-4 connect with servers 8-7 (or 6-5). The example of Fig. 2 right visualizes another case, where the second layer of EPS is configured to perform inter-group wavelength swapping. Thus, servers 1-2 communicate with servers 5-7 (or 7-5), and servers 3-4 connect with servers 6-8 (or 8-6).

Reconfiguration Time of Fully Centralized Hybrid Data Centers

When the traffic pattern changes, it is required a reconfiguration of the switches. Fig. 4 shows the reconfiguration time needed by the operations performed during a set of a hundred reconfiguration events.

The configuration of the electronic packet switches takes most of the reconfiguration time. The removal of previous flow rules requires 15 - 20 ms, and the insertion of new flows in the switches around 35 ms. Hence, important savings are achieved if these operations are minimized. The central controller reconfigures the optical switch (FOX) in approximately 1.5 ms. The communication is carried through a permanently established TCP (Transmission Control Protocol) socket, where most of the time is needed by the use of I2C (Inter-IC-Bus) as interface to control the SOAs of the fast optical switch. Finally, the servers require less than 15 ms to stop and start the traffic.

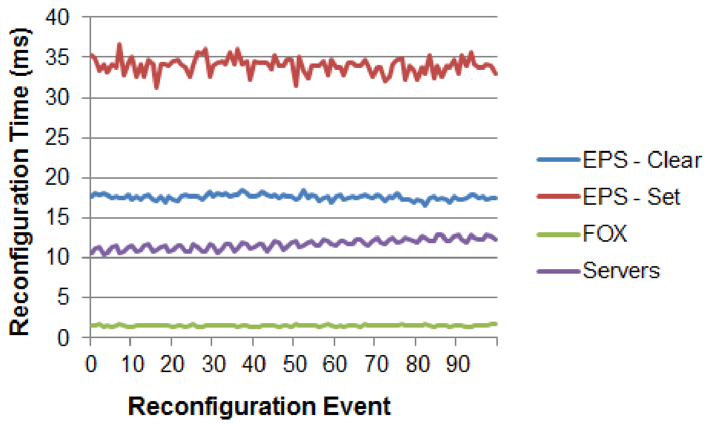

Fig. 4: Partial contributions to reconfiguration time

Design Trade-Offs Discussion

Regarding the data plane of fast optical switches, traditional WSS (Wavelength Selective Switches) based on SOAs (Semiconductor Optical Amplifier) require $m \cdot n^2$ SOAs to build n-ports WSS with m-wavelengths per port that can connect m·n servers. These WSS provide granularity at wavelength level, but unfortunately, they scale only up to ten of ports due to the large number of SOAs required. This work has demonstrated with FOX, our SOA-based OCS, that it is possible to switch a number of wavelengths per SOA, and recover the loss of granularity with E-WDM. Our demonstrator switches two wavelengths per SOA, connecting 2·n servers with n^2 SOAs. Hence, this approach connects more servers than the traditional method with equal number of SOAs, or requires fewer SOAs to connect certain amount of servers. The restriction to only two wavelengths is due to the optical bandwidth of the SOAs, and the 20 nm channel spacing of the CWDM (Coarse WDM) transceivers. However, the introduction of DWDM (Dense WDM) transceivers with 0.8 nm channel spacing would scale to at least 50 channels per port.

Regarding the control plane, FOX chooses simplicity of design by delegating the decision tasks to the central controller. Many commercial OCS follow this approach, like those based on fiber to fiber alignment[8], or based on MEMS[9] (MicroElectroMechanical Systems). In this manner, the optical switch does not require additional complexity like optical transceivers, electronic buffers, or processing units to receive, store, and decode traffic. The disadvantage of this approach is that optical switches cannot take local decisions, and require always the intervention of a central controller.

Finally, regarding the SDN control software, FOX uses a very simplified protocol to communicate with the central controller. As a result, it achieves faster reconfiguration times than the electronic switches using OpenFlow. In our proof-of-concept, a permanent TCP socket is established between FOX and the central controller, and a few bytes stream is enough to reconfigure the optical switch. This contrasts with the complexity of the OpenFlow protocol that requires several (and longer) transactions to delete and add flows every time is needed a reconfiguration of the electronic switches.

Conclusions

In this work we have presented FOX and E-WDM. FOX is a fast OCS based on SOAs that effectively duplicates the number of ports by switching two wavelengths per SOA. E-WDM is a control technique achieving wavelength switching in hybrid data centers integrating OCS and EPS. A hybrid data center demonstrator has been built, verifying the feasibility of FOX and E-WDM. The centralized reconfiguration of FOX requires approximately 1.5 ms, mainly due to the I2C interface controlling the SOAs. The reconfiguration of the whole network requires around 70 ms, being EPS the main contribution with 55 ms.

Acknowledgements

This work has been supported by the FP7 European Project COSIGN (FP7- 619572).

References

[1] N. Farrington, A. Vahdat, et al, “Helios: a hybrid electrical/optical switch architecture for modular data centers,” ACM SIGCOMM Comput. Commun. Rev., vol. 40, no. 4, p. 339, 2010.

[2] G. Wang, D. G. Andersen, M. Kaminsky, K. Papagiannaki, T. S. E. Ng, M. Kozuch, and M. Ryan, “c-Through : Part-time Optics in Data Centers,” ACM SIGCOMM Comput. Commun. Rev., vol. 40, no. 4, pp. 327–338, 2010.

[3] a. Wonfor, H. Wang, R. Penty, and I. White, “Large Port Count High-Speed Optical Switch Fabric for Use Within Datacenters [Invited],” IEEE/OSA J. Opt. Commun. Netw., vol. 3, no. 8, pp. A32–A39, 2011.

[4] P. G. Raponi, N. Andriolli, I. Cerutti, and P. Castoldi, “Two-step scheduling framework for space-wavelength modular optical interconnection networks,” IET Commun., vol. 4, no. 18, p. 2155, 2010.

[5] O. Liboiron-Ladouceur, I. Cerutti, P. G. Raponi, N. Andriolli, and P. Castoldi, “Energy-efficient design of a scalable optical multiplane interconnection architecture,” IEEE J. Sel. Top. Quantum Electron., vol. 17, no. 2, pp. 377–383, 2011.

[6] K. Xia, Y. Kaob, M. Yangb, and H. Chao, “Petabit optical switch for data center networks,” Tech. Report, Polytech. Inst. New NYU, pp. 1–9, 2010.

[7] K. I. Sato, H. Hasegawa, T. Niwa, and T. Watanabe, “A large-scale wavelength routing optical switch for data center networks,” IEEE Commun. Mag., vol. 51, no. 9, pp. 46–52, 2013.

[8] “Polatis - Series 7000.” [Online]. Available: http://www.polatis.com/series-7000-384x384-port-software-controlled-optical-circuit-switch-sdn-enabled.asp.

[9] “Calient - Photonic Switches.” [Online]. Available: http://www.calient.net/products/s-series-photonic-switch/.