

Experimental Demonstration of Backplane Architectures for Programmable Optical Nodes

Original

Experimental Demonstration of Backplane Architectures for Programmable Optical Nodes / GARRICH ALABARCE, Miquel; Magalhaes, E.; Carvalho, H.; Bianco, Andrea; Giaccone, Paolo; Zervas, G.; Simeonidou, D.; Gonzalez, N. G.; Oliveira, J.; Oliveira, J.. - STAMPA. - (2014). ((Intervento presentato al convegno ECOC 2014 tenutosi a Cannes, France nel September 2014 [10.1109/ECOC.2014.6964053]).

Availability:

This version is available at: 11583/2560139 since:

Publisher:

IEEE - INST ELECTRICAL ELECTRONICS ENGINEERS INC

Published

DOI:10.1109/ECOC.2014.6964053

Terms of use:

openAccess

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Experimental Demonstration of Backplane Architectures for Programmable Optical Nodes

Miquel Garrich^{(1)*}, Eduardo Magalhães⁽¹⁾, Heitor Carvalho⁽¹⁾, Andrea Bianco⁽²⁾, Paolo Giaccione⁽²⁾, Georgios Zervas⁽³⁾, Dimitra Simeonidou⁽³⁾, Neil G. González⁽¹⁾, Juliano Oliveira⁽¹⁾, Júlio Oliveira⁽¹⁾

⁽¹⁾CPqD Foundation, Rod. Campinas/Mogi Mirim, km 118.5, Campinas, Brazil *miquel@cpqd.com.br

⁽²⁾Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, 10128 Torino, Italy

⁽³⁾High Performance Networks Group, Dep. Electrical Electronic Engineering, University of Bristol, U.K.

Abstract Programmable optical nodes supporting heterogeneous traffic require optical backplanes with a high port count. We present two backplane architectures to enhance modularity, compare their scalability in terms of available cross-connections and we experimentally validate both proposals in a SDN scenario.

Introduction

The dramatic growth of bandwidth-hungry services and applications increases traffic heterogeneity in all network segments. Thus, future elastic optical networks may require high levels of reconfigurability, flexibility and adaptability. Research trends driven by carriers' requirements on network functions virtualization (NFV¹) are pushing hardware infrastructures capable to play different roles according to service requirements. Moreover, current proposals on software defined networking (SDN²) offer high levels of network functions programmability by decoupling the data plane and the control plane. In this context, architecture on demand³ nodes with modular and pluggable structures, or more in general, network function programmable (NFP⁴) nodes are suitable options for networks with high levels of uncertainty, because they allow enhanced control for network operations. Fig. 1 shows a NFP node consisting of an optical backplane that interconnects inputs, outputs, and modules (e.g. couplers, WSSs). NFP nodes synthesize specific architectures suited to the switching requirements of the input traffic by interconnecting, in a suitable way, the modules using the optical backplane cross-connections.

Besides their inherent flexibility, these node design concepts create new challenges, some of which are addressed, for the first time to the best of our knowledge, in this paper. To support heterogeneous traffic requests that may require different types of optical processing (e.g. space/frequency/time switching³) there is a need of an optical backplane with a high port count. However, available commercial optical switches range from 192⁵ to 320⁶ ports, thus limiting the number of available cross-connections for the specific architecture instances and the number of pluggable modules. Therefore, several optical backplane switches must be interconnected together to overcome this limitation.

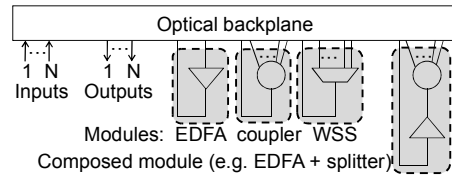


Fig. 1: Network function programmable (NFP) node

Our major contributions in this work are twofold: i) the proposal of two backplane architectures and their scalability analysis in terms of number of available cross-connections to synthesize architectures and ii) a small-scale experimental validation of both proposals.

Backplane architectures

Here we present and compare two architectures for the composition of a large optical NFP nodes with enhanced modularity (allow from pay as you grow backplane model) and high availability. Let k be the number of ports for each optical switch (e.g., $k = 192$ for the Polatis⁵ or $k = 320$ for the Calient⁶ optical switches). We denote as inlet and outlet the two fiber terminations of each port of the switch devoted to transmit and receive optical signals.

Fig. 2(a) presents the *unidirectional* backplane architecture, where U optical switches with k ports are connected in a unidirectional fashion. More precisely, the N input ports of the NFP node are connected to N inlets of the first switch. Then, N outlets of the first switch are connected to N inlets of the second switch. This process is repeated until N outlets of the switch $U - 1$ are connected to N inlets of the switch U . Finally, N outlets of the switch U are associated to the N output ports of the NFP node. By construction, optical signals are constrained to pass through all the optical switches. This backplane architecture offers a number of available cross-connections

$$X_U = U \cdot k - N \cdot (U - 1) \quad (1)$$

where the second term of the subtraction evaluates the number of ports dedicated to the

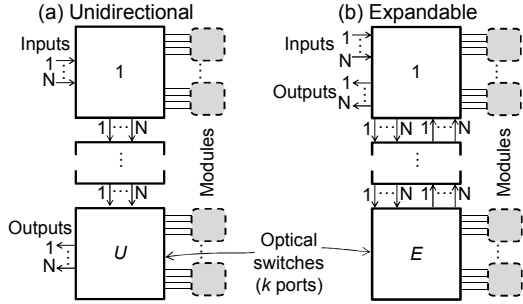


Fig. 2: Backplane architectures for NFP nodes

interconnection between switches. Therefore, the subtracted ports are not available for the modules to be connected neither for the synthesis of architectures.

Fig. 2(b) shows the *expandable* backplane architecture, where E optical switches of k ports are bidirectionally connected. More precisely, the N input and output ports are connected to the first optical switch. In addition, N outlets of the first switch are connected to N inlets of the second switch and N outlets of the second switch are connected to N inlets of the first switch. This process is repeated until N outlets of the switch $E - 1$ are connected to N inlets of switch E and N outlets of the switch E are connected to N inlets of the switch $E - 1$. This backplane architecture offers a number of available cross-connections

$$X_E = E \cdot k - 2N \cdot (E - 1) \quad (2)$$

Backplane architectural trade-offs

We consider as a reference a commercially available 320-port 3D-MEMS switch with a power consumption of $50W^6$. Fig. 3 shows an example of the power consumption of the two backplane architectures as a function of the available backplane cross-connections X according to (1) and (2), for $N = 20$ and $k = 320$. Similar behavior was observed for other parameter settings, not reported here for the sake of space.

On the one hand, the adaptable nature that characterizes the expandable architecture permits to increase in a step-wise fashion the power consumption (i.e. the number of used backplane switches $E = \{1, 2, 3, 4, 5\}$) according to the number of available cross-connections. Indeed, once the NFP node is operating, this architecture uses the optical switches in an incremental manner because switches in the first stages of optical backplane are the first ones to be completely used. Therefore, the expandable architecture permits to adapt E to the traffic request offering a clear benefit for resource dimensioning purposes because additional optical switches may be powered on

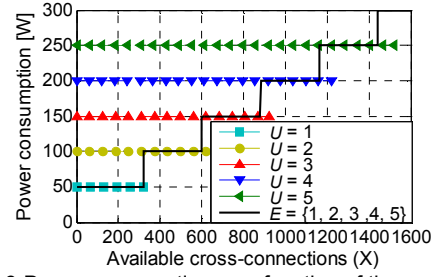


Fig. 3: Power consumption as a function of the available cross-connections for the unidirectional and expandable backplane architectures with $k = 320$ and $N = 20$

only when required. In addition, the connection of additional optical switches does not compromise already established optical links through the NFP node.

On the other hand, the number of backplane switches in the unidirectional architecture U must be set in a resource dimensioning study carried out before the NFP node is deployed and used. Indeed, once U is set and the node is operating, the connection of additional optical switches would compromise already established optical links. However, for very limited ranges of supported cross-connections X (e.g. $1160 \leq X \leq 1220$ and $1440 \leq X \leq 1520$) a higher power consumption (i.e. a higher number of backplane switches) is required by the expandable architecture compared to the unidirectional one. Indeed, the unidirectional supports more cross-connections for a given number of optical switches due to the lower number of ports used to interconnect them. Finally, note that given the same number of backplane switches $E = U$, the expandable permits an arbitrary utilization of the modules that belong to different optical switches, whereas in the unidirectional case this would not be possible. More complex tree-based composition topologies could be considered as well, but either they are equivalent from the performance point of view to the ones considered in this work, or their possible advantages are limited. Their investigation is left for future work.

Experimental setup and results

In our experiment we validate the proposed backplane architectures for the NFP node with an aggregate traffic of 8.96 Tb/s. The transmitter (Fig. 4(a)) is composed of 80 continuous wave (CW) lasers with 50-GHz channel spacing. Each CW is modulated by four multiplexed lines of 28 Gb/s (PRBS $2^{31}-1$) that generate 112-Gb/s DP-QPSK channels. We add the 80 modulated channels at three ROADMs of CPqD's network testbed² (Fig. 4(b)) and we send them towards the NFP node according to the considered scenarios. The SDN controller sets the network

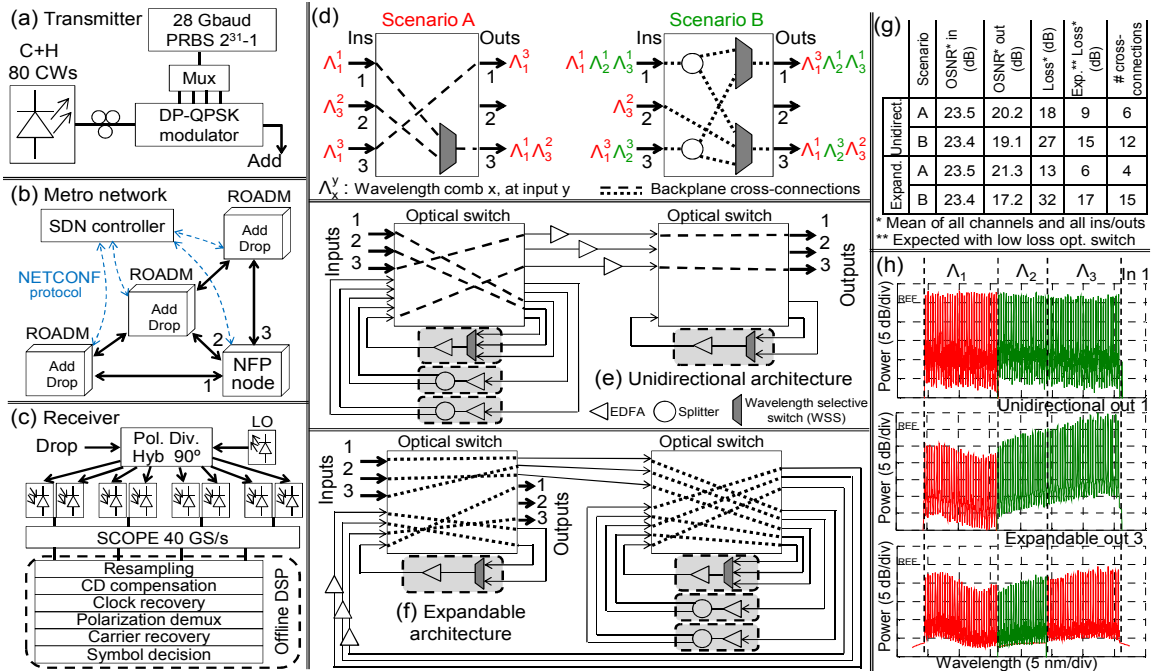


Fig. 4: (a) Transmitter, (b) metro network and (c) receiver setups for the validation of the NFP in (d) two scenarios using (e) the unidirectional and (f) the expandable backplane architectures. (g) Summary of results and (h) spectra at input 1 and outputs 1 and 3 for the unidirectional and expandable architectures respectively in scenario B.

parameters via NETCONF⁷. At the receiver (Fig. 4(c)), the polarizations of each dropped 112-Gb/s channel are fed to 40-GS/s real time scope and digital signal processing algorithms are used offline. We consider two scenarios shown in Fig. 4(d) in which the spectra is divided in aggregates of channels $\Lambda_1 = \{1529.62, \dots, 1541\}$, $\Lambda_2 = \{1541.38, \dots, 1548.94\}$ and $\Lambda_3 = \{1549.38, \dots, 1561.04\}$ (central wavelengths in nm). Scenario A presents a single aggregate of channels at each input. Conversely, scenario B considers channel aggregates of A plus loopback traffic so as to synthesize a node with more optical functionalities. The NFP node is implemented with two 8×8 planar lightwave circuits (PLCs with 5 dB of loss per cross-connection) that interconnect inputs, outputs and composed modules. The unidirectional and expandable architectures (Figs. 4(e) and 4(f)) show the required cross-connections for scenarios A and B respectively.

Fig. 4(g) shows our obtained results in terms of optical signal-to-noise ratio (OSNR) and optical loss. The unidirectional architecture presents high OSNR at the output ports and lower loss compared with the expandable architecture in scenario B. Indeed, for switching cases with high number of cross-connections, the unidirectional architecture is convenient due to the reduced connectivity between backplane optical switches. However, a clear benefit is observed for the expandable architecture against the unidirectional one in scenario A.

In particular, when all channels of a given input are switched to the same output fewer cross-connections, modules and backplane switches are required. Indeed, only the first optical switch is used in the expandable for scenario A reducing losses and OSNR degradation.

Conclusions

We presented two backplane architectures (unidirectional and expandable) for NFP nodes and compared their scalability in terms of available cross-connections. We validated both architectures in a small-scale experimental demonstration under a SDN controller. The expandable architecture offers higher modularity which provides clear benefits for resource dimensioning and power saving purposes.

Acknowledgements

This work is supported by FUNTEL, FINEP, CPqD Foundation under the project 100GETH, the EC FP7 grant no. 317999: IDEALIST, and the EPSRC grant EP/I01196X: The Photonics Hyperhighway.

References

- [1] www.etsi.org/technologies-clusters/technologies/nfv
- [2] J. Oliveira, et al., "Towards Software Defined Autonomic Terabit Optical Networks", *OFC*, W2A.35, 2014.
- [3] N. Amaya, et al., "Introducing Node Architecture Flexibility for Elastic Optical Networks", *JOCN*, 2013.
- [4] Y. Yan, et al., "FPGA-based Optical Network Function Programmable Node", *OFC*, W1C.1, 2014.
- [5] Polatis' Series 6000 Switch, www.polatis.com
- [6] Calient's S320 Optical Switch, www.calient.net
- [7] IETF rfc on NETCONF, tools.ietf.org/html/rfc6241