

Strategies for the Interconnection of Heterogeneous Optical Networks

Ignacio de Miguel, Iria Prieto*, Ramón J. Durán, Óscar González de Dios, Santiago García-Jiménez*, Eduardo Magaña*, Antonio Ferrero, Mikel Izal*, Noemí Merayo, Daniel Morató*, Rubén M. Lorenzo**
Universidad de Valladolid, E.T.S.I. de Telecomunicación, Campus Miguel Delibes, 47011 Valladolid, Spain

** Public University of Navarre, Campus de Arrosadía s/n, 31006 Pamplona, Spain*

*** Telefónica I+D, Core Network Evolution, GCTO, Don Ramón de la Cruz 82, Madrid, Spain*

e-mail: ignacio.miguel@tel.uva.es, iria.prieto@unavarra.es

ABSTRACT

Optical networks are facing increased levels of heterogeneity, from types of services to the range of technologies involved. In this paper, we focus on the inter-domain heterogeneity scenario and consider the case of a set of interconnected network domains, each one using its own optical networking technology (wavelength-routing, optical burst-switching or wavelength-routed optical burst switching). We discuss different approaches to deal with such heterogeneity in order to provide end-to-end connections traversing different domains, including the use of hierarchical PCE (Path Computation Element) mechanisms and interfacing strategies between optical burst switching and wavelength routing-based networks.

Keywords: heterogeneous optical networks, multi-domain scenarios, hierarchical PCE, wavelength routing, optical burst switching.

1. INTRODUCTION

There is a very wide range of technological alternatives for optical core networking. For instance, different network operators may employ different technologies, such as SDH, carrier-grade Ethernet or IP/WDM based on wavelength-routed optical networks (WRONs), and even optical burst switching (OBS). Besides that, each carrier will usually employ different wavelength channels and different granularities (data rates) from those used in other networks. These issues, and especially if each carrier employs different switching paradigms, makes the physical interconnection of the networks a very difficult task. Moreover, it should be noticed that carriers are generally reluctant in sharing topology information, and they aim at controlling their resources according to their own policies. Hence, the efficient interconnection of different carriers' networks is an open issue.

In this paper, we focus on such scenario and consider the case of a set of interconnected network domains, each one using its own optical switching technology (like WRON or OBS), and discuss strategies for the efficient interconnection of those heterogeneous networks.

First of all, routing in multi-domain networks is one of the main issues to be solved. Its complexity comes from the size of a multi-domain network, and also from the reluctance of carriers to share network information, so that decisions may be made with incomplete information. Methods for constraint-based path computation have been developed, mainly relying on the use of Path Computation Elements (PCEs) [1] (or equivalent devices). PCEs are network entities which hold topology information and are polled by the nodes to determine the path from a source to a destination node. A domain may have one or more PCEs, and although they have access to local domain resource/policy databases, their inter-domain visibility may vary. Peer-to-peer mechanisms (where PCEs of neighboring domains interact with each other to exchange routing information) and hierarchical schemes have been proposed as approaches to solve the topology visibility issue [2]. Moreover, when considering the connection of heterogeneous optical networks, the knowledge of the technologies used in each domain should also be taken into account when performing routing, as it may be better to traverse a domain with a compatible technology than a domain with a completely different one [2].

On the other hand, since different networks operate with different technologies, strategies for interfacing for both data and control messages should also be considered. In general, two different strategies can be used, either performing electro-optical conversion at the interface between domains or keeping data in the optical domain.

These issues are briefly analysed in this paper. Thus, in Section 2, we study the potential of hierarchical PCE for interconnecting heterogeneous optical networks, and then, we discuss interfacing strategies between heterogeneous domains (in particular, based on optical burst switching and wavelength routing) in Section 3.

2. HIERARCHICAL PCE FOR INTERCONNECTING HETEROGENEOUS OPTICAL NETWORKS

The hierarchical PCE strategy for multi-domain networks consists in employing a PCE per domain (children PCEs), which are in charge of computing intra-domain routes, together with a global PCE (parent PCE), which is in charge of computing inter-domain routes [1]. While each child PCE has only visibility of its own domain, the parent PCE is agnostic to that information, but it knows the set of domains composing the whole network and the set of links connecting those domains. Children PCEs are responsible alone of computing intra-domain routes.

However, in order to compute inter-domain end-to-end routes, the parent PCE works in cooperation with children PCEs, without requiring the exchange of internal topological details.

The hierarchical PCE approach is useful for interconnecting heterogeneous optical networks, since the parent PCE may have knowledge of the technology used in each domain, and such knowledge can be exploited in path computation, for instance by adjusting link weights according to the complexity of the network interface between heterogeneous domains. However, a difficulty for this approach comes from the lack of availability of PCE architectures for other optical networks than WRONs. Thus, in this section we first show an example of how hierarchical PCE can help addressing the interconnection of heterogeneous domains, and then discuss the potential of PCE for optical networks based on different switching paradigms.

2.1 Simulation analysis of interconnected heterogeneous domains

We have performed a simulation study to show how hierarchical PCE can be used to compute both intra and inter-domain routes in a heterogeneous multi-domain scenario while taking into account the compatibility of technologies used in adjacent domains. The topology used in the study is shown in Figure 1a. It is based on the MUPBED topology [3] and consists of six different domains, using two different technologies (say A and B). Inter-domain links are equipped with wavelength converters in their endpoints, and support up to 11 wavelength channels. Moreover, for path computation purposes, inter-domain links which connect domains with different technologies have a higher cost than those connecting homogeneous domains (β times higher). In this way, β is used to model the preference to traverse a domain with a compatible technology rather than a domain with a completely different one. Regarding the load, we have assumed that 80% of connection requests are addressed towards nodes of the same domain, while the remaining 20% of the requests are addressed towards nodes of other domains. Shortest path routing is used within each domain. On the other hand, for inter-domain connections, the parent PCE computes all possible routes between the source and destination domain, and collects from the children PCEs the costs of traversing the different domains as required. Based on that information, the route with lowest total cost is selected by the parent PCE.

Results have been obtained as a function of the β parameter. As β is increased, for low network loads (where the network will operate) the mean number of changes of technology for inter-domain routes decreases, as shown in Figure 1b, with no significant impact on blocking probability (Figure 1c). However, increasing β leads to using longer routes in order to bypass heterogeneous domains, which translates in slightly longer routes (the number of hops increases 2.6% approximately for low loads, around 0.1, in this scenario), and thus typically in higher propagation time.




Figure 1: (a) MUPBED topology (only including low capacity links [3]);

(b) Mean number of changes of technology for inter-domain routes; (c) Mean global blocking probability.

2.2 PCE architectures and extensions for optical burst switching-based domains

As previously mentioned, a difficulty for the use of the hierarchical PCE approach comes from the lack of PCE architectures for other optical networks than WRONs, so here we discuss its use for two additional types of switching paradigms: wavelength routing-optical burst switching (WR-OBS), and OBS.

The WR-OBS architecture combines the processing and buffering capabilities of electronics at the network edge with the advantage of optical transport in the core. However, in contrast with traditional OBS approaches, it uses acknowledged establishment of dynamic lightpaths for the transmission of those bursts, thereby facilitating its implementation with existing technology [4]. In other words, the WR-OBS is a highly dynamic circuit-switched network, where the establishment of lightpaths is triggered by the arrival of data at edge routers rather than by user requests or the network manager.

We have adapted the PCE architecture so that it can be employed for this type of networks [5]. The PCE architecture, shown in Figure 2a, contains typical PCE elements like a *request* and a *response dispatcher*, a *path computing processor* (where different algorithms can be loaded), and a *traffic engineering database* (TED). However, additional features are implemented. Due to the high dynamism of the WR-OBS, the PCE should not rely on an Interior Gateway Protocol like OSPF-TE to keep track of the state of a network, as the information could be outdated and therefore the PCE could assign again to new requests wavelengths that are in the process

of being reserved. Therefore, the PCE needs to keep track itself of wavelength reservations. Hence, the *wavelength reservation manager* is in charge of updating the TED based on the results of path computations and on the notifications received about the end of utilization of wavelength paths. An additional modification is the inclusion of a module for re-scheduling path requests. When a combination of both route and available wavelengths cannot be found by the *path computing processor*, rather than blocking the request, it can be sent to a *path computing retry queue*, with the hope that in a short period of time resources will become available (since the WR-OBS is highly dynamic). This technique has been shown to decrease the blocking probability at the expense of increasing end-to-end delay [6]. Besides these modifications to the PCE architecture, the Path Computation Element Protocol (PCEP) [1] has been complemented with a number of extensions, which include mechanisms to indicate the desired reservation time of a connection and to cancel a connection when resources are no longer required [5]. Finally, we have implemented an emulated setup, and measurements on key performance metrics have shown very good agreement with theoretical models (e.g., Figure 2b) [5].

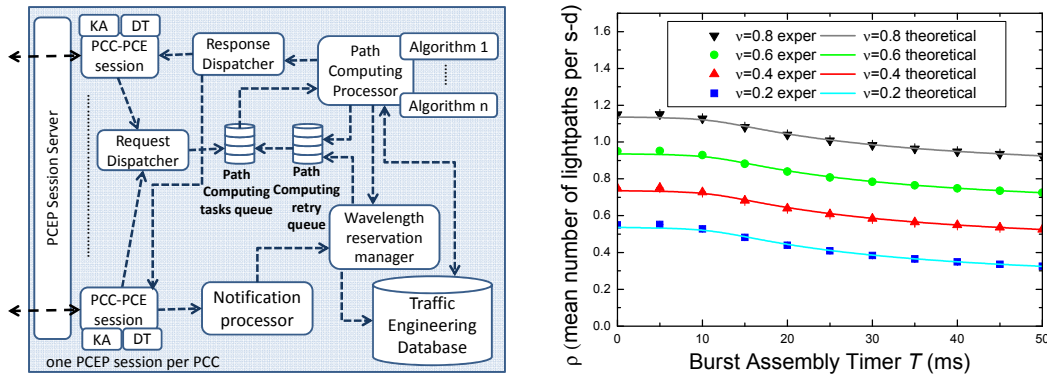


Figure 2: (a) PCE design for highly dynamic domains, e.g., based on WR-OBS (b) Experimental and theoretical mean number of lightpaths per source-destination pair for different traffic loads, v , and timers [5].

The use of PCE in a pure OBS network becomes more complex due the extremely high dynamism of optical burst switching. However, PCE could still be used to periodically compute routes that could be provided to source nodes where bursts are built. In this way, source routing could be used in the transport of those bursts, or at least the most suitable outgoing frontier node could be indicated depending on the destination node, for instance, taking into account the number of changes of technology required (as pointed out in Section 2.1) and the availability of the interfacing strategies that will be now discussed in Section 3.

3. INTERFACING STRATEGIES FOR OPTICAL BURST SWITCHING AND WAVELENGTH ROUTING DOMAINS

In multi-domain scenarios, the border of two domains can be either one edge node from any of the domains or, as shown in Figure 3, the border can be a link between two edge nodes (one from each domain). The first option implies that the control of the interconnection is given to one of the domains, which is not a realistic option. Therefore, the second option has been chosen for the study.




Figure 3. Data interface between domains.

For this study, we have also chosen OBS for the first domain and WRON for the second. In that case, the UNI (User-Network Interface) provided by OBS is electric and the one provided by WRON is optical (lightpaths). For the NNI (Network-Network Interface) we can consider:

- Interface with electrical conversion: this allows the use of buffers and technologies of already existing data networks. Therefore, this alternative will not be reviewed.
- All-optical interface: this option represents the most interesting scenario. Wavelength conversion and fiber delay lines (FDL) allow more flexibility in the interconnection, so we also assume this scenario.

NNI1 is an edge node in the OBS domain, so it is responsible of triggering the establishment of the necessary lightpaths to transport the bursts over the next WRON domain. Therefore, when NNI1 receives a burst, it will act differently depending if the destination node in the second domain has an established lightpath or not.

The first time that NNI1 receives a burst request (a Burst Control Packet, BCP) with destination a certain output node in the second domain, NNI1 signals NNI2 to request a new lightpath. This lightpath will be used to transport all BCPs and bursts from the first domain to the same output node of the second domain, independently of the source node in the first domain. If there is at least one lightpath established with the output node of the second domain, but all of them are busy transmitting bursts, NNI1 is responsible of deciding if a new lightpath is needed considering the estimated traffic in the near future. In some cases, it will be enough to use FDLs to hold up those peaks with more traffic volume. When a lightpath is not used any more, NNI1 is responsible of removing it. It will imply a new signalling to NNI2.

As the lightpath establishment over the WRON domain can take some time (even several seconds), there are some alternatives in how this delay is accommodated from the point of view of the OBS domain. The time offset for the BCP can be increased, at least for the first bursts with destination a new output node of the second domain. This would imply an extra delay not assumable for several applications. Another possibility would be to accept losing bursts in the process of lightpath establishment; however, that does not seem a good solution. One more possibility would be to use techniques of sending empty bursts as soon as a new burstifier (with bursts to a new destination) is created. This would allow to request the lightpath establishment before than the first bursts with data get out from the burstifier. However, since the sizes of bursts are usually small, this scenario would not be realistic as the size of the bursts sent should be much bigger than usual. For that reason the most realistic scenario consists in using FDLs to store the first incoming bursts to allow the lightpath establishment.

In order to study the strategy to transport OBS bursts through a WRON network, we have created OBS incoming bursts by aggregating self-similar traffic generated using the Perlin Noise algorithm [7]. The WRON network employs the 14-node NSFNet topology and is equipped with 13 wavelengths per link and 13 transmitters/receivers per node (at 10 Gb/s). The WRON receives not only requests for lightpath establishment (and release) from the OBS network but also originated in its own WRON nodes (in this case according to a Poisson Process). Results have been obtained as a function of the traffic load generated in the WRON, and also as a function of the incoming data rate of OBS traffic (1, 10 and 60 Gb/s) and average size of bursts (1, 1.75 and 3 MB). As Figure 4 shows, the blocking probability of the requests is mainly dependent on the underlying load of the WRON rather than on the characteristics of the incoming OBS traffic.

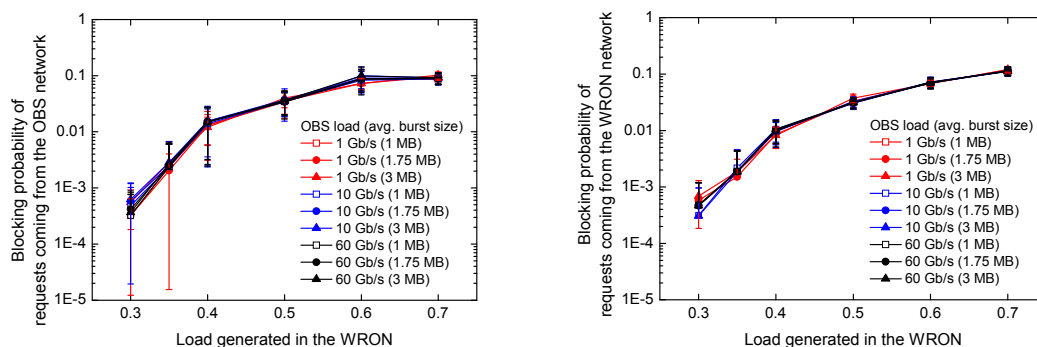


Figure 4. Blocking probability of requests coming from (a) the OBS network (b) the WRON.

ACKNOWLEDGEMENTS

This work has been supported by the Spanish Ministry of Science and Innovation (TEC2010-21178-C02-01, TEC2010-21178-C02-02 and TEC2010-12250-E) and Public University of Navarre through PIF grant. We are also grateful to Diego Martínez Azcona and David Ortega Lora for their help with some programming tasks.

REFERENCES

- [1] F. Paolucci, F. Cugini, A. Giorgetti, N. Sambo, and P. Castoldi: A survey on the path computation element (PCE) architecture, *IEEE Communications Surveys & Tutorials*, vol. 15, no 4, pp. 1819-1841, 2013.
- [2] M. Chamania and A. Jukan: A survey of inter-domain peering and provisioning solutions for the next generation optical networks, *IEEE Communications Surveys & Tutorials*, vol. 11, no 1, pp. 33-51, 2009.
- [3] G. Maier, C. Busca, and A. Pattavina: Multi-domain routing techniques with topology aggregation in ASON networks, in *Proc. ONDM 2008*.
- [4] M. Düser and P. Bayvel: Analysis of a dynamically wavelength-routed optical burst switched network architecture, *Journal of Lightwave Technology*, vol. 20, no. 4, pp. 574-585, April 2002.
- [5] O. González de Dios *et al.*: Experimental demonstration of a PCE for wavelength-routed optical burst-switched (WR-OBS) networks, in *Proc. ONDM 2013*, pp. 269-274.
- [6] E. Kozlovski, M. Düser, I. de Miguel, and P. Bayvel: Analysis of burst scheduling for dynamic wavelength assignment in optical burst-switched networks, in *Proc. IEEE LEOS 2001*, vol. 1, pp. 161-162.
- [7] I. Prieto, M. Izal, D. Morato, and E. Magaña: Traffic generator using Perlin Noise, in *Proc. IEEE GLOBECOM 2012*, pp 1847-1852.