Benefits of Optical Packet Switching for Router By-pass in Metro Networks

Csaba KISS KALLÓ[2], Víctor LOPEZ[1], John DUNNE[2], Óscar GONZALEZ DE DIOS[1], Mark BASHAM[2], Juan FERNANDEZ-PALACIOS[1]
[1]*Telefónica I+D, c/ Don Ramón de la Cruz, 82-84. 28006, Madrid, Spain*
Tel: +34 913373923, Email: jpfpg@tid.es
[2]*Intune Networks, 9C Beckett Way, Park West, Dublin 12, Ireland*
Tel: +44 (0) 1872 279915, Email: mark.basham@intunenetworks.com

Abstract: This paper builds on previous work [1] on IP off-loading over multi-granular photonic switching technologies to provide an in-depth techno-economic analysis for a real metro scenario. The main contribution of the paper is that it demonstrates a 42% capital cost reduction in favour of the proposed IP-offloading architecture, when compared against currently used, typical all-IP architectures. This is achieved by performing packet transport, aggregation, switching and grooming in the optical layer within an IP off-loading architecture using the Optical Packet Switch and Transport (OPST) technology. Such results show that multi-granular optical technologies are a strong candidate for solving the bottleneck problem caused by video streaming, cloud services and mobility in metropolitan area networks. Finally, the paper also discusses other, cumulative benefit aspects of multi-granular technologies, such as video and next-generation cloud services.

Keywords: IP off-loading, future proof networks, optical packet switching, router by-pass, economic modelling.

1. Introduction

According to recent traffic forecasts [2], by 2015 the global data consumption will cross the Zettabyte (10^{21} bytes) threshold. The biggest chunk of this large amount of data is some form of video content. The most popular video service over the Internet is YouTube. It is estimated that 48 hours of YouTube video is uploaded per minute and that the number of views per day is 3 billion.

This huge amount of data highly impacts service operator networks, so Content Distribution Networks (CDN) are used for storing and distributing content more effectively [3]. The idea is either to use a dynamic DNS (Domain Name Server) that points to the CDN server assigned to the geographical region of the user or to perform a HTTP redirection from a central site. Needless to say, the closer the content is to the user the better, because the video traverses less network elements, with less risk of packet loss and therefore reduced quality of experience for the video audience.

There is an ongoing debate in the industry around small distributed versus large centralized caches. The former can be located closer to the user, thus decreasing the chances for bottleneck formation. The latter instead concentrates a larger number of users, thus the popular, cached video is closer to the overall user population. Zink et al. estimated that at the campus network level the probability of finding a specific video in a local cache (aka *hit rate*) is 30% [4]. On the other hand, Cha et al. performed the same estimation considering the whole population of Korea [5]. Three different caching strategies were proposed and the most efficient one produced a 98% hit rate, with a storage size of less than

300,000 videos. This is because the cache population is large and the chance to find a previously downloaded video increases accordingly.

The results of [4, 5] motivate the optimal utilization of efficient network architectures in the metro network. The CDN reduces the traffic from other networks, but as the cache must cover a high number of users the traffic load in the metro network is not reduced. Furthermore, cloud services and mobility [6, 7] are putting increasing traffic load on the metro network.

This paper aims at extending the IP-offloading architecture work of [1] with a techno-economic study on a real network. The IP-offloading architecture based on multi-granular photonic switching technologies is explained in more detail and its economic feasibility is quantified through a thorough economic analysis.

2. Architectural solutions for metro-core networks

Two metro architectures are modelled and evaluated. First, an all-IP architecture that is typically used by network operators is modelled. Second, a model for a sub-wavelength packet forwarding technology called Optical Packet Switch and Transport (OPST) is also created and compared to the all-IP solution.

2.1 – All-IP reference metropolitan network

Currently, carrier networks are based on IP routing from the access nodes to the IP interconnection (aka Internet Exchange). In this scenario the Broadband Remote Access Server (BRAS) functionality is in the IP access router. The BRAS node is responsible for subscriber management and is also used for IP quality of service (QoS) management. In order for access routers to connect to other regions or domains the traffic is switched through one or more transit routers connected in a hierarchical architecture (Figure 1). The number of levels in the hierarchy depends on the size of the network. Transit routers are interconnected through long reach optical transport links. A subset of the transit routers is also connected to core routers that provide connectivity to Internet Exchange routers located outside the operator network. This architecture is used typically as a reference for the aggregation part of a carrier network.

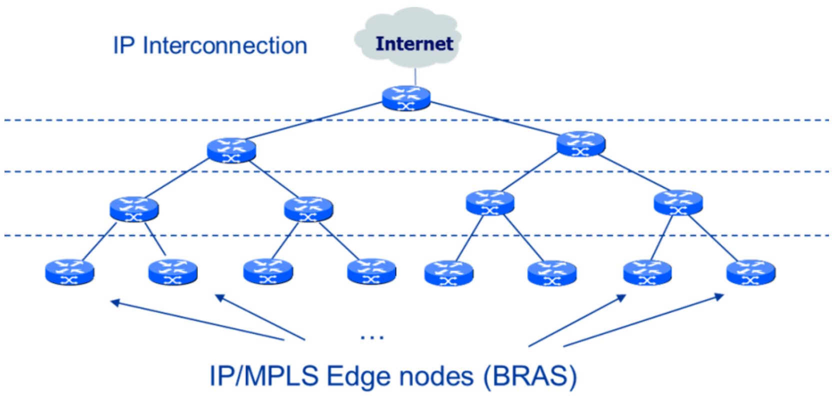

Figure 1: All-IP network hierarchy with IP access routers at the leaf level, transit routers at the middle levels and core routers at the root.

2.2 – Optical Packet Switch and Transport (OPST) metro architecture

Optical Packet Switch and Transport (OPST) [9] is an optical sub-wavelength packet switching technology developed by Intune Networks. OPST fully integrates the packet forwarding of layer 2 with the sub-lambda transport of layer 1 and physical fibre of layer 0 into a single layer (Figure 2). The OPST nodes contain 4 or 8 port cards, each equipped with an ultra-fast tuneable laser capable of tuning between wavelengths at nanosecond speeds. Each port card is assigned a wavelength that is used as its unique identifier. A port card can only drop its assigned wavelength. A node can drop as many wavelengths as the number of port cards installed.

Basic operation involves reading incoming packet addresses, queuing them according to destination and class of service and grouping them into bursts dedicated to the wavelength assigned to the destination port. A scheduler is used for serving all of the queues in a fair manner based on the availability of resources around the distributed switch. The source node laser is then rapidly tuned to the destination wavelength and the packet is sent out on the optical ring. The optical ring – that operates as a switch fabric – can be 300km long while remaining in "plug-and-play" operational mode requiring minimum operator intervention for set-up. Longer ring lengths are possible with additional amplification. The destination OPST port receives the burst, decomposes it into packets and forwards them on at the appropriate client port. The data path is managed through a single OPST control and management layer for Layer 1 and Layer 2, which simplifies network operation and, hence, enables significant operational expenditure (OpEx) reductions.

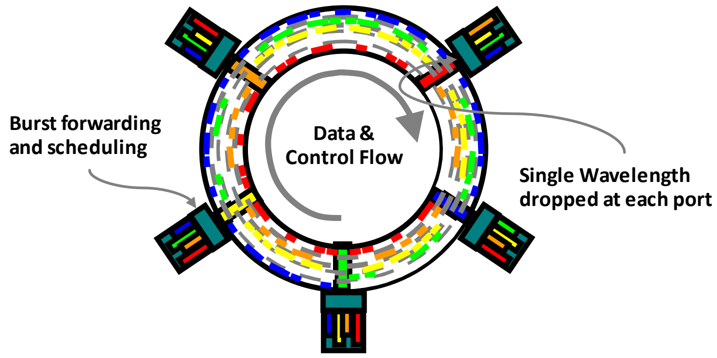

Figure 2: Illustration of a wavelength routed system – Intune OPST

The proposed network architecture combines OPST in the metro network with distribution of the BRAS functionality out to the Optical Line Terminal (OLT) (Figure 3). OLTs are used as the access head-end device in the assumed passive optical network (PON) in the access segment. Next in this paper OLTs with integrated BRAS functionality are referred as IP-OLTs.

The IP-OLT controls packet processing and classification, IP routing and MPLS tunnel creation. Client-facing OPST nodes receive MPLS tunnels from the IP-OLTs and transmit the traffic to the OPST concentrator (i.e., the metro head-end OPST node). Traffic is groomed through the OPST network so that packets with common destination are presented on separate OPST concentrator ports. This removes the need for transit routers in the metro head-end, as the only functionality that is needed is a simple transport pipe. A ROADM-based photonic mesh is used for inter-connecting the metro head-end nodes (i.e., inter-regional connectivity) as well as for connecting them to the core routers (i.e., international connectivity). This ROADM mesh forms a Wavelength Switched Optical Network (WSON) [8].

The control plane information exchange is handled by the control IP router. The role of this router is to avoid the explosion of the number of adjacencies in the control plane. Further, the control IP router facilitates the segmentation of the IP domains in this plane.

2.3 – Connectivity

There are several connection types used in the network architecture presented in the previous subsection. These connections are providing local (i.e., intra-regional), inter-regional and international connectivity and are described below.

a) *IP-OLT to OPST node* and *Interconnection IP router to ROADM*: is based on MPLS tunnels.
b) *OPST node to OPST node*: is performed by tuning the transmitter OPST node to the receiver wavelength of the destination OPST node.
c) *OPST node to OPST concentrator*: depending on the external destination (i.e. inter-regional or international) – the OPST node tunes to the corresponding receiver wavelength of the OPST concentrator inside its region.
d) *OPST concentrators to OPST concentrators* and *OPST concentrators to Interconnection IP router (through the WSON)*: the OPST concentrators and the Interconnection router are connected to ROADMs. These optical nodes are interconnected by means of routed wavelengths. These wavelengths are filled through grooming of a set of sub-wavelengths (towards the same destination) performed through the OPST distributed switch.

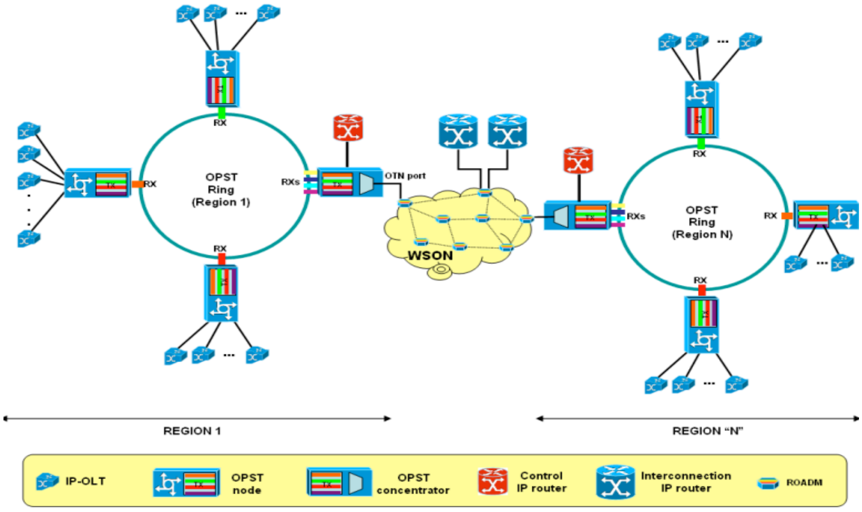

Figure 3: OPST based network architecture

2.4 – Traffic Assumptions

The network consists of 20 access nodes connected by optical spans of up to 140km (Figure 5). The access traffic requirements are for the Madrid Metro Regional Network for the 5 year period between 2010 and 2015. Based on internal reports, 100Gbps traffic is assumed for 2010, increasing by 330% to 330Gbps in 2013 and to 1Tbps by 2015. All traffic is presented to OPST equipment through 1GE and 10GE ports. Full protection is required for 100% of the traffic. A sample of the traffic growth used for a subset of the access nodes is shown in Figure 4.

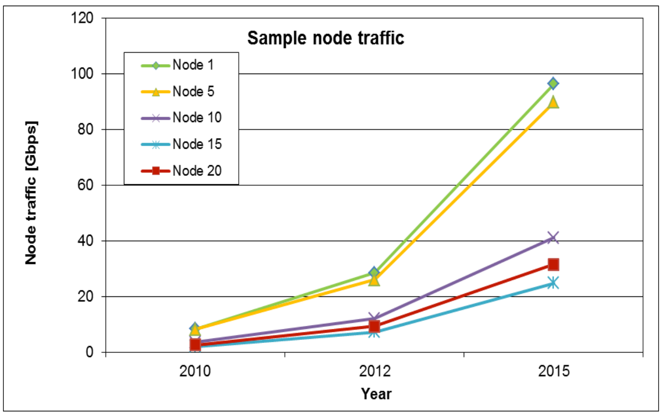

Figure 4: Sample traffic growth used in the cost model for IP-offloading

3. Results

3.1 – Scenario description

Two metropolitan architectures are defined, based on all-IP and OPST technologies.

First, in the *all-IP metropolitan network scenario* 30 OLTs are connected to each IP access node (Figure 5). As previously mentioned, the BRAS functionality is in the IP access router. Twenty access nodes and a pair of transit routers (nodes 25 and 26) are used for collecting the traffic from the access network. Each transit router is connected to one international Interconnection IP router (nodes 21 and 22). Both transit routers are connected to a pair of national interconnection IP routers (nodes 23 and 24).

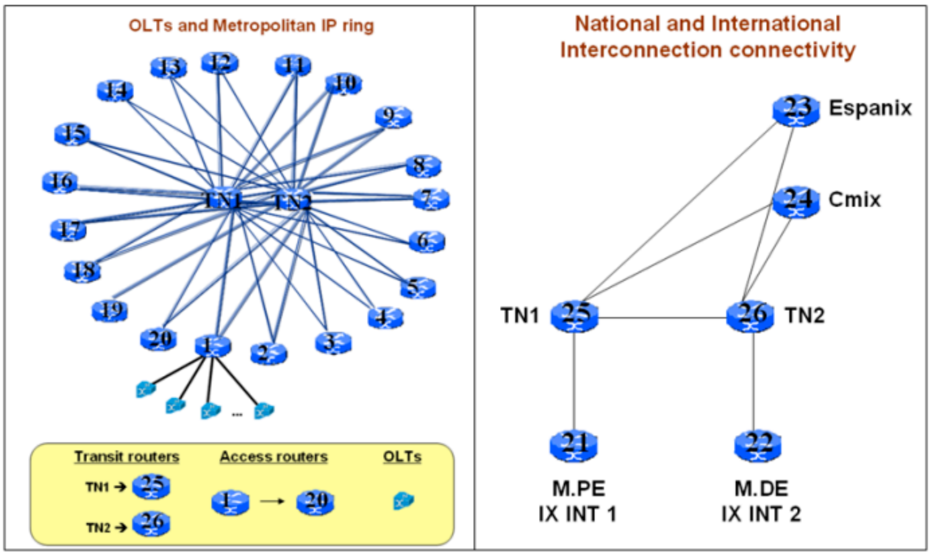

Figure 5: All-IP metropolitan network for the case study

Second, in the *OPST based metro network scenario* (Figure 3) BRAS-enabled OLTs (i.e., IP-OLTs) are connected to OPST nodes, as the OPST network replaces the IP access and IP transit routers. In this scenario OPST takes over the metro aggregation, grooming and switching functionality from the IP routers. The aggregated traffic is also optically groomed, so that packets with the same destination (i.e., national or international exchange routers) are handed off from dedicated ports to the WSON core network. This solution eliminates the need for routing in the transit sites like 25 and 26.

Models were built for both of the above scenarios, and the cost of each solution was derived based on typical component pricing. Monetary units (M.U.) used in the paper should be interpreted as 1MU being the price of a 1GBE SX port.

3.2 – Modelling results

The first and most important result is shown in Figure 5, presenting a CapEx comparison of the two solutions analysed. The figure presents a CapEx reduction on the favour of the OPST solution of 42%, 32% and 35% for 2010, 2012 and 2015, respectively. It can be noted that the first-year CapEx reduction is higher than that of the other two years. This is due to the low utilisation of the chassis slots in the all-IP solution. The OPST solution uses small iVX200 fan-in switches featuring 1G client ports that interface with the IP-OLT and connect to the iVX8000 10GE ports on the uplink. These switches aggregate small 1GE flows into 10GE ones, while also extending the reach of the OPST ring. The iVX200 and iVX8000 are network elements implementing the OPST technology described in [9, 10].

In 2012 the amount of traffic increases significantly (to 330Gbps from 100Gbps in 2010) resulting in a more efficient utilisation of the deployed equipment in the all-IP solution. The equipment cost reflects the traffic growth much closer in this year, thus the CapEx reduction enabled by OPST drops back to 32%. In 2015 both CapEx values increase significantly to reflect that the total traffic in the network is now 1,000Gbps. However, the gap between an IP vs OPST solution increases as the OPST solution enables 35% of CapEx savings for this last year of the study.

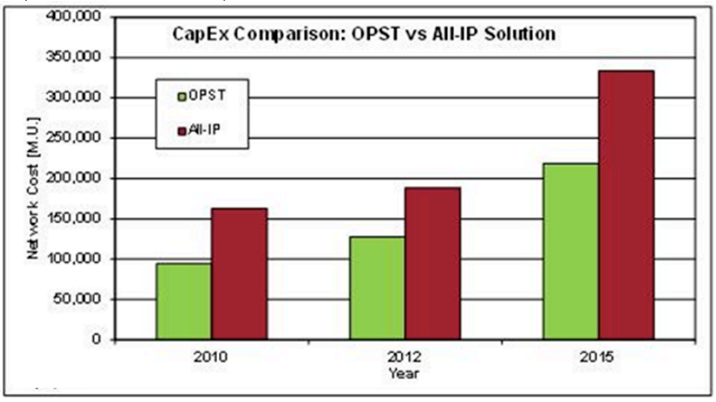

Figure 5: CapEx Comparison: OPST vs All-IP solution

Figure 6 shows the cost impact of integrating the BRAS functionality into an OLT device, transforming it into a true IP access node. The range of 10% to 100% increment is used on the OLT cost for the added BRAS functionality. However, according to current market information the actual cost for this functionality integration is very close to the lower end of the range. In any case, Figure 6 shows that even with the unlikely 100% cost increment 8% to 20% cost savings are enabled. At the lower end (i.e., 10%) of the range the obtained cost savings reach as high as 42%.

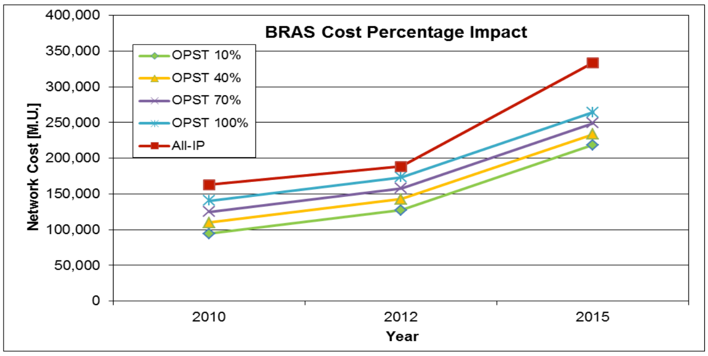

Figure 6: BRAS functionality cost impact when integrated into OLT

Figure 7 shows the CapEx breakdown per component type for both architectures. New OPST cards added to the network also increase the capacity of the distributed switch fabric, as new, dedicated wavelengths are used for each OPST card. This enables the OPST fabric size to linearly scale in line with traffic growth, smoothing out the cash flow requirements and aligning them to revenue growth. Also, the all-IP solution requires more investment upfront to cover the cost of the complete chassis fabric.

3.3 – Discussion: Benefits of an OPST network

The CapEx savings enabled through IP off-loading reflect one aspect only of the benefits enabled by the OPST technology. The calculated 42% of CapEx reduction (shown on the left bar of Figure 9) is due to the capabilities of the OPST technology to aggregate traffic from multiple metro access sites into a single wavelength of a metro ring and transport, switch and groom that traffic while it travels through the distributed OPST switch. Such functions are currently carried out at multiple layers (Layer 1, 2 and 3) and multiple levels in the hub and spoke architecture hierarchy, which means that the current approach is complex and expensive to scale. By using OPST equipment in the metro space almost all of the transit routers are eliminated from the network and further OpEx benefits are obtained through the simplified network operation and scaling.

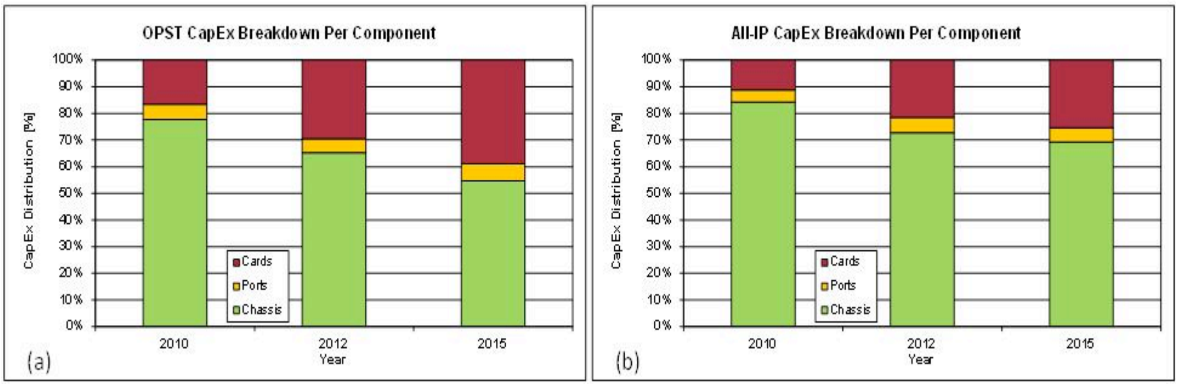

Figure 7: CapEx breakdown per component type. a) OPST; b) All-IP

Beyond metro aggregation there are other areas that benefit significantly from an OPST network enabling IP off-loading. One of the compelling next-generation services to mention is the Virtual Personal Computer (VPC) service [7]. The VPC service is built around low cost terminals (aka thin clients) with reduced functionality – I/O and network connectivity – with all the processing and storage functions being placed on a server located in the network cloud. The VPC service is an attractive one as it enables 51% cost reduction in the first year of operation compared to traditional PCs [11]. However, it also puts stringent requirements on the telecommunications infrastructure as it relies heavily on high-bandwidth, low-latency and low-jitter connections to the data centre. As shown in [12], even small increments of additional latency can make a huge impact on revenue generated by a service. Furthermore, a VPC is a mobile device, which means that user data has to be available from anywhere in the network at a consistent performance. Therefore, the underlying networking technology used for providing this service must handle very cost-effectively the quick delivery of user data to frequently changing points of connection. The VPC service was modelled in depth in [7] and the main conclusion of the study is that an OPST-based solution can deliver the VPC service at 20% of the cost of the current network architecture approach (see right bar in Figure 9). This is another important aspect of OPST-enabled benefits.

Finally, an important area of modern networking is video services. IPTV (Internet Television) and VoD (Video on Demand) services generate more than 40% of the traffic today in carrier networks and this percentage is increasing rapidly. Carriers do not have a proven monetisation solution yet for this video traffic, while their costs for delivering over-the-top (OTT) content through their networks is high. Therefore, finding cost-effective solutions for video service is a major concern. The single optical hop based mesh connectivity provided by OPST in a metro ring enables video streaming servers to be placed close to the user for high performance and away from the storage for a higher statistical multiplexing gain between storage disks and server. Furthermore, the drop and continue multicast of OPST also reduces the required capacity in a metro ring, as on the metro head-end downlink ports the content appears in a single copy, while in current solutions as many copies are required as access nodes connected to the same head-end node. According to early estimates, these OPST features enable CapEx savings in excess of 65% (see middle bar in Figure 9).

The CapEx benefits of OPST described in this section are additive and represent a small sample of the full space of applications that can benefit from such multi-granular optical packet technologies. The quantification of further benefits is subject for future work.

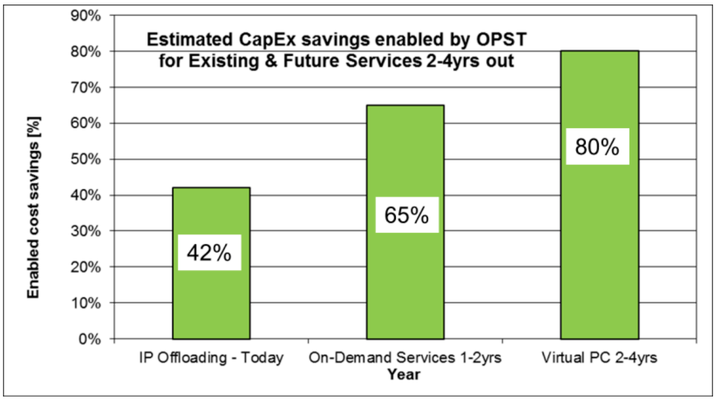

Figure 9: Additive OPST benefit quantifications

4. Conclusions

Current all-IP architecture costs depend significantly on traffic growth: the higher the traffic, the higher the network costs. Consequently, any cost increase is reflected on the service provider's margins. This paper demonstrates that IP off-loading over multi-granular photonic switching technologies enables up to 42% CapEx savings. Such savings are achieved through the removal of most transit routers from the network as their routing, aggregation and grooming functions can be fulfilled in the multi-granular photonic network. Moreover, the paper discusses other benefits enabled by the proposed architecture, specifically, significant cost reductions for video and cloud services.

An all-IP solution requires a high initial investment because the common equipment dominates the CapEx, whereas the CapEx of the multi-granular OPST solution follows the actual traffic demand in the network significantly closer. Consequently, OPST proposes a future-proof network infrastructure investment solution as it yields greater benefits as more video, mobile and next-generation cloud services are added to the network.

Acknowledgements

The MAINS research project funded by the EC's FP7 Programme under agreement nº 247706.

References

[1] J. P. Fernández-Palacios et al., "IP Off-loading over Multi-granular Photonic Switching Technologies", European Conference and Exhibition on Optical Communication (ECOC), Torino 2010.
[2] Cisco White Paper, Cisco Visual Networking Index: Forecast and Methodology, 2010–2015, June 2011.
[3] A. Rosas et al., Evolutionary projects in Content Delivery [online], Future Internet Assembly, 2010.
[4] M. Zink et. al, "Characteristics of YouTube network traffic at a campus network - Measurements, models, and implications", *in Computer Networks* 53(4), 501-514, 2009.
[5] M. Cha et. al, "Analyzing the Video Popularity Characteristics of Large-Scale User Generated Content Systems", i*n ACM/IEEE Transactions on Networking,* 5(17), 2009.
[6] J. Fullaondo et. al, Optical Friendly HiSpeed File Transfer Protocol For Enabling Next Generation Nomadic Virtual PC Services, *ECOC 2011*, Sep, 2011.
[7] Cs. Kiss Kalló et. al., "Cost Reduction of 80% in Next-Generation Virtual PC Service Economics Using an OPST Sub-Wavelength Metro Network", In NOC 2011.
[8] Y. Lee et. al., "Framework for GMPLS and PCE Control of Wavelength Switched Optical Networks (WSON)", IETF Draft, Feb. 2011.
[9] Intune Networks, "Optical Packet Switch and Transport. A Technical Introduction", 2009.
http://www.intunenetworks.com/home/shape-up/core_innovation/opst_technical_introduction/
[10] Intune Networks, "*Verisma iVX8000 Product Data Sheet*", http://www.intunenetworks.com, 2010.
[11] IGEL Technology, "Slimming Down with Slim Computers", 2007.
[12] J. Hamilton, "Perspectives: The Cost of Latency,"
http://perspectives.mvdirona.com/2009/10/31/TheCostOfLatency, 2009.