Optical Packet Switch and Transport: A New Metro Platform to Reduce Costs and Power by 50% to 75% While Simultaneously Increasing Deterministic Performance Levels

John Dunne*, Tom Farrell, Jim Shields
Intune Networks Limited, 9b Beckett Way, Parkwest, Dublin, Ireland
*Tel: (353) 1620470, e-mail: john.dunne@intunenetworks.com**

ABSTRACT

In this paper we present a new optical networking platform called Optical Packet Switch and Transport or OPST. We present some techno-economic modelling results that show how OPST can reduce the Total Cost of Ownership (TCO) of a Metro Network by up to 50%, it can reduce the power consumption by up to 75%, and space requirements are reduced by up to 50%. OPST outperforms existing equipment by having the capability of combining connection oriented and connectionless packet flows over the same infrastructure. An overview of OPST will also be provided, describing the key technology blocks that enable the OPST platform.

Keywords: Optical, Packet, Transport, Switch, Burst, Asynchronous

1. INTRODUCTION

The birth of Skype, Bebo and social networking after 2005, represents a tipping point creating a huge surge in online traffic. This in turn has forced carriers to increase the performance of their networks to overprovision static systems to handle unpredictability. However, this approach is causing a major business problem as the cost of increasing the performance is eroding the margins of the network operators. As a result, carriers do not have the free cash flow to invest in new services and generate new business with improved margins.

Today s networks are not capable of providing the dynamic bandwidth necessary to meet customers demands and guarantee the quality of service at the same time. What is needed is a new approach to networking that would allow for dynamic bandwidth provisioning, without the need for more fibre. Another major concern is the power consumption of existing technologies, particular larger centralised IP Routers. Telecoms Operators are taking their carbon footprint and environmental impact more seriously now than ever, setting aggressive targets to reduce power consumption and carbon footprint in their networks over the coming years [1].

2. OPST: OPTICAL PACKET SWITCH & TRANSPORT OVERVIEW

Optical Packet Switch and Transport is a new networking platform that truly collapses layers 0 to 2 under the same control plane. The control plane runs internally inside a ring network, transforming the entire ring into a distributed switch that operates as a single new network element. The collapsing of layers 0 to 2 is achieved by using ultra-fast tunable laser transmitters on the line side of each externally facing client port. The tunable transmitters act as both transmitters and switches simultaneously, which collapses the layers under one control system.

The ring uses a wavelength routing scheme to address packet flows. Each OPST port has a fixed wavelength filter, which acts as the address. This is implemented using a wavelength selective switch typically used in ROADMs. The filter is set once at the installation of the system. Each tunable transmitter can therefore simultaneously switch its packet flow to a destination as well as transmit to that destination by tuning to the target wavelength. When the transmitters are used on Optical Burst mode, virtual wavelength paths can be set up and pulled down in response to incoming packet flow requirements. The result is an ability to merge packet flows from different sources optically, so that they arrive multiplexed in time at the destination.

Wavelength routed networks have been reported by many researchers for the past 15 years [2-5]. However the majority of these types of systems are built as transport vehicles only and not as distributed switches. Also, many of them use ring-wide synchronisation techniques with some sort of token passing to enable access to the fibre without collisions.

In contrast, the OPST system uses a new Optical Media Access Control system (OMAC), which employs a Carrier Sense Media Access with Collision Avoidance (CSMA-CA). This is an asynchronous access system that avoids the need for ring-wide synchronisation. By removing the dependence on ring-wide synchronisation, OPST becomes more practical to implement and operate because it can adapt to a wide variety of fibre span lengths and fibre types. It also avoids the pitfalls of chromatic dispersion on synchronisation across many wavelengths. Typical synchronisation techniques use an image of what is actually happening on a separate control channel around the ring. If this image is corrupted, the system breaks down, and the correlation between the ring image the actual ring performance is not inherently tied together.

Digital Object Identifier: 10.4108/ICST.BROADNETS2009.7830
http://dx.doi.org/10.4108/ICST.BROADNETS2009.7830

Figure1 shows the main building blocks of the OPST system and these building blocks are described individually below. In the diagram a ring with 6 OPST ports and 6 client ports is shown. It should be noted that these ports can be placed arbitrarily around the ring, or packaged together at any one location into a chassis system in order to share power, cooling and other base infrastructure features.

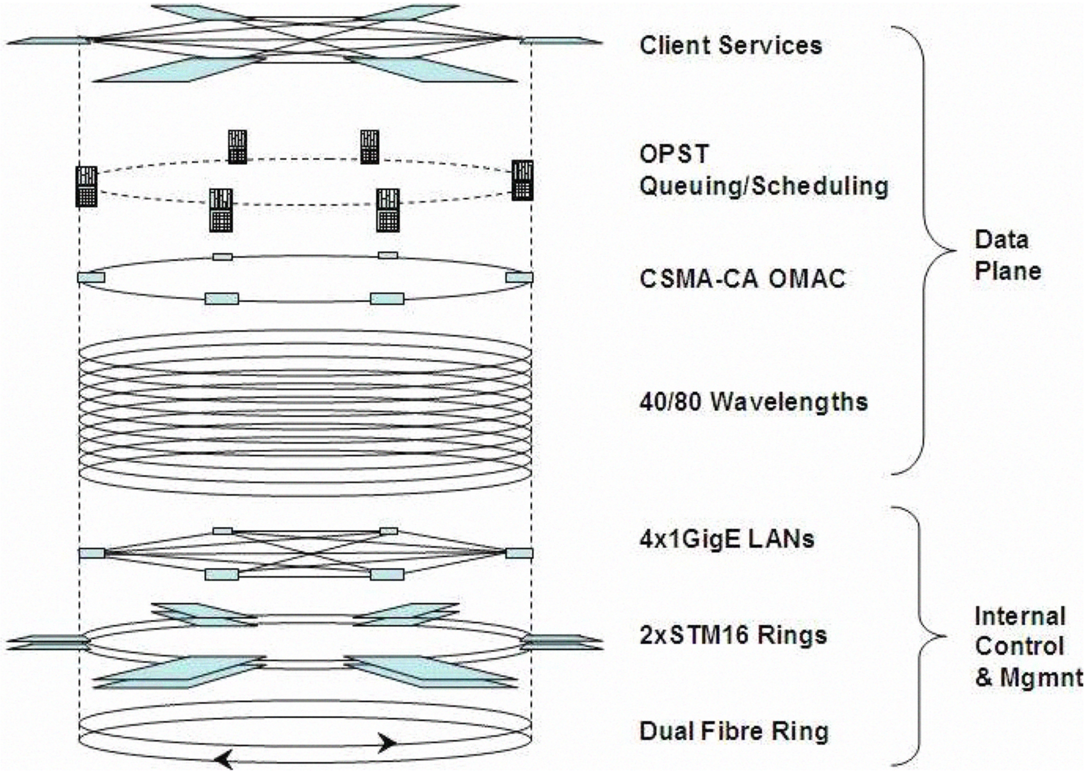

Figure 1: The main building blocks of the OPST system

2.1 Dual-Fibre Ring

The fundamental building block of the system is a dual-fibre optic ring. All signals travel unidirectionally on each ring but in opposite directions to the other ring. The fibre rings are simultaneously the transport media for the system and also each ring is the switching fabric for a distributed switch. The use of dual switching fabrics simultaneously improves the performance of the distributed switch, while providing the high reliability required for carrier grade networks in the transport and connection functions.

2.2 2xSTM16 Control Channels & 4x1GigE LANs

The internal control system and the management systems are built on a separate wavelength dedicated to this purpose. This wavelength is used to daisy-chain a set of STM16 links. Each STM frame contains 2x1GiGE LAN connections, the first for control, the second for management. Each fibre carriers such an internal network in opposite directions for resiliency in the case of fibre breaks or node failure.

2.3 40/80 Wavelengths and the CSMA-CA OMAC

The dual fibres each have 40 or 80 wavelengths used for the data paths. Note that these wavelengths are constantly appearing and disappearing to form virtual paths. None of the wavelengths are permanently on . The lasers and transmitters are used in burst mode all the time. The OMAC uses an optical channel monitor built into every port that acts as a carrier sensor. The scheduling system uses the carrier sense to avoid collisions on the ring, thus providing media access to the enormous 400Gbps-800Gbps of available bandwidth from any port. In effect, it acts like a large set of parallel shared media that are asynchronously accessed.

2.4 OPST Queuing & Scheduling

The OPST layer is a distributed packet switch, whose unit of currency is an OPST frame. This is a proprietary frame format that closely resembles an SDH frame format. Incoming client packets are encapsulated as OPST frames, then queued in a Virtual Output Queue (VOQ) on a destination and class of service basis. The schedulers

Digital Object Identifier: 10.4108/ICST.BROADNETS2009.7830
http://dx.doi.org/10.4108/ICST.BROADNETS2009.7830

are given boundary conditions in the form of credits by the control system. However, from then on, each scheduler operates individually at each port. The behaviour of each scheduler is completely autonomous and ascrynhronous to every other scheduler on the ring.

The scheduler composes a burst by selecting parts of the various OPST frames across the classes of service destined for switching and transmission. The OMAC described above is used to add the burst onto the ring. The effect is that no packets are dropped at the OPST layer once they have been queued. Also note that client packet boundaries do not align to optical burst boundaries, making the system completely agnostic to the type of client traffic being sent across the ring.

2.5 Client Services

The OPST platform can switch and transport any client services, for example both Ethernet and SDH/SONET services can exist simultaneously. On the same ring network, an Ethernet NIC (Network Interface Card) can be placed on one set of OPST ports, and a SDH/SONET NIC can be placed on another set of OPST ports. This is possible due to the asynchronous nature of the underlying switch fabric, alongside the independent behaviour of each wavelength in the system. A single OPST ring can therefore contain a set of Virtual Bridges alongside a set of SDH/SONET cross connects.

The system has another unique attribute by allowing both dimensioned and undimensioned packet flows to exist simultaneously. The dimensioned flows all have guaranteed bandwidth, latency and jitter performance characteristics. Whereas the undimensioned flows can adaptively use up the remaining bandwidth. This ensures that carriers can offer a range of terms and conditions around their Service Level Agreements (SLA), while operating all services over the same infrastructure.

2.6 Summary of OPST Performance Benefits

OPST offers significant simplification of network operations as both the optical and OPST layers are hidden freeing operator resources to focus on rapid introduction of client services to generate new revenue streams. The operator has a single infrastructure to manage, but can build virtual bridges on this infrastructure and resell the physical network, as a collection of virtual networks, to other operators and service providers. Within all of these features is the ability to guarantee the bandwidth, latency and jitter characteristics of dimensioned packet flows in services such as ELine, ELAN and ETree as defined by the Metro Ethernet Forum (MEF). These services can be sold under SLA for a range of classes of service. Finally, the infrastructure can simultaneously offer an adaptive and efficient use of the remaining undimensioned bandwidth for best efforts and internet browsing traffic. This bandwidth can be sold at lower rates but remain profitable as it shares the same infrastructure as other services.

3. MODELLING RESULTS

The cost-performance attributes of the OPST platform were derived using real metro network topologies and traffic profiles by working with several network operators in the European region. We present here a generic example of how these models were constructed and the nature of the results achieved. This example illustrates the principles behind the results, which derive from the following 3 attributes of the system: (1) each tunable laser can address any other node, thereby ensuring the system has the least number of transponders necessary to create a fully meshed network, (2) the OPST system can merge packet flows in the optical layer, ensuring a ring-wide use of statistical gain, which minimises the number of ports required and (3) each metro solution that uses OPST has no additional layer 2 switches, as the switch comes for free as part of the optical ring network. These 3 attributes appear in every network solution built using OPST ensuring a consistently lower CAPEX and OPEX compared to other alternatives. Simultaneously, the performance is enhanced because the simplicity of a single control system over layers 0-2 enables deterministic behaviour of the combined transport and switching function in the metro area network.

3.1 Techno-Economic Modelling Results

Firstly, the techno-economic models were built using metro area network topologies and traffic profiles and by building up complete metro solutions using 3 types of equipment: (1) layer 3 over WDM (2) layer 2 over WDM and (3) OPST. The example below shows a typical Metro Collector network, which is aggregating and transporting traffic up and down through the collector network for multiple services, connecting the IP Core Platforms to the IP Edge devices. Traffic profiles for 2 years with growth from year to year are shown. To illustrate the reduction in transponders that OPST offers, a count of the required number of transponders is shown.

Digital Object Identifier: 10.4108/ICST.BROADNETS2009.7830
http://dx.doi.org/10.4108/ICST.BROADNETS2009.7830

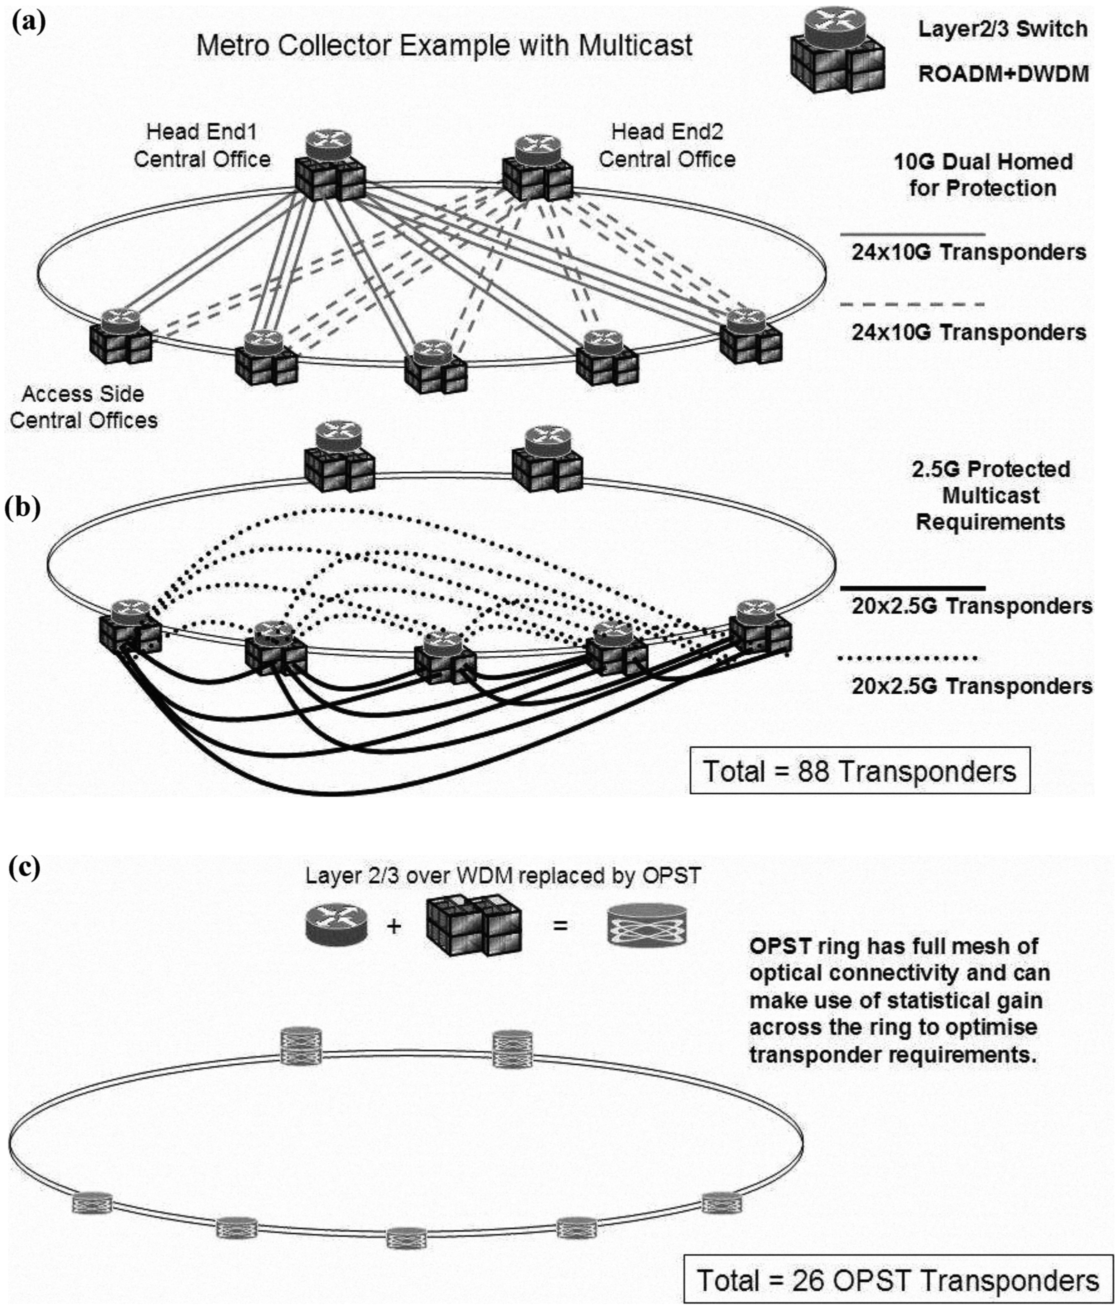

Figure 2: Metro Collector network using 3 types of equipment: (a) layer 2/3 over WDM (b) layer 2/3 with Multicast and (c) OPST.

Each metro solution was then costed by using estimates of the pricing of available commercial equipment and the target pricing for the OPST platform. The total capital equipment expenditure (CAPEX) was calculated for each type of solution. In this particular example shown in Fig. 2, the result was that the cost of OPST was 38% lower than the Layer 2 over WDM solution and 58% lower than the Layer 3 over WDM solution. In fact, over a wide number of models, the average savings were in the range 50-75%, depending on particular traffic profiles and topologies.

A similar operational expenditure (OPEX) model was calculated taking into account space, power consumption, cooling and other operational factors such as truck rolls and manpower to install, operate and maintain. Once again, the results showed an average of ~50% reduction in the OPEX for the OPST solution compared to the Layer 2 and Layer 3 over WDM solutions modelled.

3.2 Power Consumption Modelling Results

Using similar Metro network models described above, power consumption calculations were performed to illustrate the total reduction in power consumption across a complete network solution, when OPST is deployed. Typical power consumption figures for commercially available equipment were used in comparative modelled solutions. For example, the Cisco CRS1 consumes 14.3kW per rack, Cisco 7609 consumes 4.9kW, Cisco 15454 ROADM consumes 1kW, Force10 E1200 consumes 7.2kW, Extreme 10808 consumes 2.7kW and Ciena 4200 consumes 800W.

Taking the average for a 40-port 10GigE line-rate access-side switch as 2.7kW, taking the average for a 32-port 10GigE line-rate head-end switch as 4.2kW and finally, taking the average for a ROADM at 800W and ignoring regeneration features, the comparative power consumptions can be calculated for the metro region. For the Layer 2 over WDM solution, the power consumption sums to 37.9kW, whereas for the OPST solution, the power consumption sums to 11.8kW. There is a 68.9% reduction in power consumption for OPST. In fact, power consumption results consistently fell in the 50-75% reduction range for a variety of network topologies and traffic growth profiles. The power consumption reduction is due to the lower amount of equipment being deployed in the OPST system, where the switches and transport equipment have been truly merged into a single platform.

4. CONCLUSIONS

A new metro platform called Optical Packet Switch and Transport (OPST) has been presented with associated techno-economic modelling results. The results have been developed by modelling real metro networks and using real traffic and topology data provided by carriers. An example of a typical metro collector network is used to illustrate how a 50-75% savings in both CAPEX and OPEX results in the use of the OPST platform versus today s Layer 2 and Layer 3 over WDM solutions. The OPST system core design is also explained to illustrate how it can significantly improve the deterministic performance of dimensioned traffic, while simultaneously improving the performance of undimensioned best efforts traffic. OPST is agnostic to the type of services that are delivered over the system and so in conclusion, OPST represents a breakthrough platform solution for the future performance and economics of Metro Area Networks.

ACKNOWLEDGEMENTS

OPST has been developed by and is wholly owned by Intune Networks and we would like to acknowledge the many years of work that the entire team at Intune Networks have devoted to the development of this new platform.

REFERENCES

[1]Http://www.btplc.com/Societyandenvironment/Environmentandclimatechange/Environmentandclimatechange.htm

[2] RINGO: a WDM Ring Optical Packet Network Demonstrator, R.Gaudino, A.Carena, V.Ferrero, A.Pozzi, V.DeFeo, P.Gigante, F.Neri and P.Poggiolini, ECOC 2001.

[3] Experimental Demonstration of an Access Point for HORNET – a Packet-over-WDM Mulitple-Access MAN, D.Wonglumsom, I.Whiate, K.Shrikhande, M.Rogge, S.Gemelos, F.An, Y.Fukashiro, M.Avenarius, L.Kaxovsky, Journal of Lightwave Technology, Vol.18, No.12, December 2000.

[4] Novel Node Configuration for DWDM Photonic Access Ring Using CMLS, H.Takesue, F.Yamamoto, T.Sugie, IEEE Photonics Technology Letters, Vol.12, No.12, December 2000.

[5]All-Optical WDM Multi-Rings with Differentiated QoS, M.Marsan, A.Bianco, E.Leonardi, A.Morabito, F.Neri, IEEE Communications Magazine, February 1999.

Digital Object Identifier: 10.4108/ICST.BROADNETS2009.7830
http://dx.doi.org/10.4108/ICST.BROADNETS2009.7830