Dynamic CANON: A Scalable Multidomain Core Network

Alexandros Stavdas, Theofanis G. Orphanoudakis, Helen C. Leligou, Konstantinos Kanonakis, Chris Matrakidis, Andreas Drakos, and John D. Angelopoulos, University of Peloponnese

Andrew Lord, British Telecommunications

ABSTRACT

The explosion of current demand has brought the contemporary multidomain core network paradigm to its limit. In the quest for new approaches that exploit recent developments in optical technology, a novel network architecture that obviates most of the expensive and loss-prone centralized all-optical switches is described in this work. It is based on clustered architecture for nodes in optical networks and features a reconciliation between dynamic resource allocation and guaranteed end-to-end network performance in a multidomain network. This article enhances the distributed, collision-free slot aggregation inside domains of clustered core nodes with dynamic switching of slots/frames between the domains. Thus, it can support dynamic sub-wavelength allocations between network domains, using standard burst-switching techniques. This extends the high efficiency and multiplexing gain into the inter-domain network even under highly bursty traffic. It features both low-cost optical add/drop edge nodes exploiting WDM transmission and agile and modular centralized electro-optical switches that are presented in conjunction with the overall network architecture. Its performance exhibits very low burst loss probability traded for a higher but tolerable and bounded delay.

INTRODUCTION

The current migration of access networks from digital subscriber line (xDSL) technologies toward various fiber to the-x (FTTx) solutions will expose the inner networks not only to a continuous increase in traffic but also to higher traffic fluctuations. This will bring the contemporary metro and core network architecture to its limit. The main limitation of this approach is the excessive amount of layer 2/3 (L2/L3) processing and switching of such multi-protocol, multi-bitrate traffic, which can only rely on an excessive deployment of expensive electronic hardware, suffering from high power consumption, management complexity, and large office real estate. The problem is further exacerbated by the fact that existing networks are multidomain, meaning that a large-scale network is partitioned into different administrative domains deploying equipment from different vendors, whereas at their (multi-granular) interfaces, inter-operability is required to ensure effective end-to-end service delivery. This network paradigm cannot scale cost-effectively.

To address this bottleneck, architectures with fewer layer transitions and reduced layer 2/3 processing leading to flatter, that is, less hierarchical, network structures are sought. Flat networks allow for enormous savings in operating expenditures (OPEX) and capital expenditures (CAPEX), and it is the only network architecture that can deliver long-term economically scalable growth. Along these lines, a two-tier core network architecture is proposed where the rigid distinctions between metro, outer core, and inner core become blurred significantly, alleviating the processing burden, when traffic crosses core nodes or when it traverses different domains.

Recently, a considerable number of new approaches for constructing a dynamic core network were proposed around various optical burst switching/optical packet switching (OBS/OPS) solutions. In the OBS/OPS schemes, the lack of effective optical buffering presents a hard dilemma between high loss (when adopting one-way reservation) and high delay (when using a circuit-switch such as two-way reservation). Specifically, OBS solutions with little or no buffering at the nodes have been extensively investigated, for example [1–3], and shown to suffer from heavy burst loss even at low utilization levels due to their ambitious on-the-fly act. The problems arise from conflicts due to the temporal properties of data bursts entering the OBS nodes, which the method can alter only poorly due to lack of buffering. As a result, losses grow exponentially with each node the traffic must cross.

In an alternative OPS-based solution described in [4], a metropolitan area network (MAN) consisted of a number of optical rings where traffic was stored in buffers located in the

 0163-6804/08/$25.00

periphery nodes awaiting the proper slot and wavelength to release the payload. As soon as the slots were launched in the optical layer, they traveled toward their destination with no further buffering. This is a cost-effective and distributed solution that can provide relief from loss, but the two-way reservation used in the medium access control (MAC) protocol that achieves the lossless operation does not scale well beyond metro dimensions due to the prohibitive delay.

An architecture addressing the previously described issues and presenting a scalable solution without additional excessive L2/L3 processing when crossing domain boundaries is presented in this article. It is based on clustered architecture for nodes in optical networks (CANON), extending the initial concept presented in [5] through an enhanced — in terms of scalability — method to interconnect a large number of node clusters even at long distances, offering high end-to-end performance and cost-effective basic building blocks. The remainder of this article is organized as follows: an overview of the CANON network architecture is presented; then node architectures to implement the proposed solution are described; and the mode of operation is analyzed. Finally, the performance is compared to that of current approaches for dynamic optical networking based on OBS.

The CANON Network Architecture

CANON offers an architecture for constructing a dynamically reconfigurable slotted optical core network exploiting both wavelength division multiple access (WDMA) and time division multiple access (TDMA) techniques. The initial concept of CANON is discussed in [5], and a detailed description and evaluation of a protocol for aggregating traffic from ring clusters through pre-provisioned wavelengths in the inter-cluster network is presented in [6]. The scope of CANON is two-fold. The first is to provide a coordinated mechanism for interconnecting large numbers of core nodes of a two-tier core network that are formed either from a number of MANs (like those in [4]) over a national network or from a large number of nodes over an extended geographical area like the Pan-European, National Science Foundation (NSF) backbone networks, and so on. An important aspect is that in both cases, the nodes could also belong to different administrative domains. The second is to reconcile end-to-end guaranteed slot/packet loss performance with statistical multiplexing gains so that the overall network resources are utilized efficiently while end-to-end QoS constraints are respected.

Under the CANON concept, a large-scale network, for example, the one that is schematically shown in Fig. 1, is partitioned into a number of geographically limited areas, taking into account various criteria like administrative domains, topological characteristics, traffic patterns, legacy infrastructure, and so on. An important consideration is that each of these clusters is comprised of a group of nodes in geographical

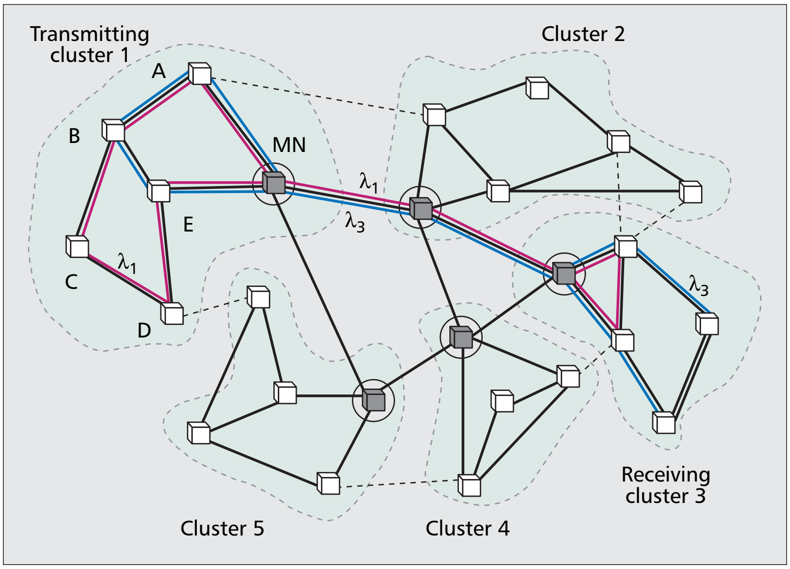

Figure 1. *CANON: A network architecture exploiting node clustering.*

proximity shown in highlighted areas in Fig. 1. The clusters can coincide with administrative domains, but there could be many cases where two or more clusters belong to the same administrative domain. Therefore, in the most general case, the partitioning into specific clusters can be either a static or a dynamic process. In the current work, only the static case is considered.

As shown in Fig. 1, the interconnection between clusters is made through a node that is called a master node (MN) as opposed to regular nodes (RNs) that operate in a "slave mode" under the coordination of the MN. The MN has a dual-role: to act as the gateway between the particular cluster and other clusters and to coordinate the transmission of all other nodes of its own cluster domain. Thus, the problem of controlling a core network, with a large number of nodes, is transformed into the simpler one of coordinating the *inter-cluster* and *intra-cluster* arbitration mechanisms. The coordination of these two mechanisms is a central issue in CANON and the key to ensuring efficient transportation in a multidomain environment.

In the most general case, the intra-cluster topology is further decomposed to a fixed number of partially overlapping, elementary topologies like *ring* or *bus* so that a number of nodes share the same WDM channels. For example, in Fig. 1, one can observe that the partial-meshed network of Cluster-1 can be decomposed to two rings such as A, B, E, MN and A, B, C, D, E, MN. In CANON, the ring/bus topology is viewed as a network building block so it is assumed that a cluster consists of a number of logically independent but physically, partially co-located rings.

Thus, for the intra-cluster operation, the MN decides which WDM channel is to be assigned to which ring, and all nodes within this ring share the same channel. In common with most ring architectures, the shortest path routing principle is *not* followed within a cluster; that is, node A is transmitting to node MN via a longer path, for example, A, B, C, D, E, MN or A, B, E, MN. To simplify the concept in the rest of this work, it is

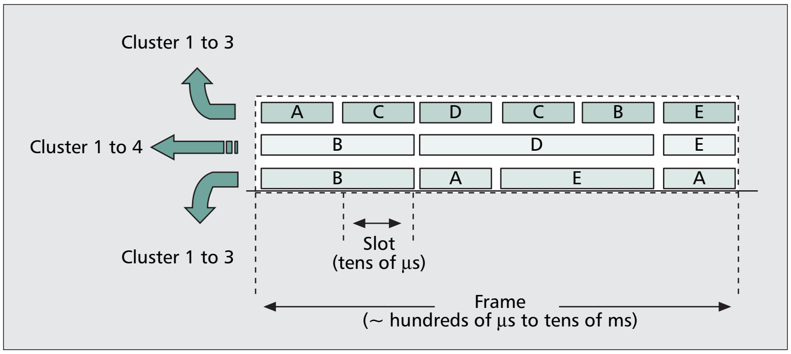

■ **Figure 2.** *A schematic representation of the relation between slots and frame.*

assumed that each cluster is formed by means of a single ring, which is adequate for presenting system operation. In reality, however, reasons of resilience dictate the existence of redundancies. Although there is no room to cover this subject here, it should be mentioned that dual MNs, each controlling one direction of counter-rotating traffic, exist in each ring, in accordance with classical ring protection methods. The inter-cluster network interconnects both MNs of each cluster in a mesh network with enough redundancy as deemed appropriate by the operator (not unlike present synchronous digital hierarchy (SDH) systems).

For efficient end-to-end transportation, two bandwidth granularities are considered in CANON: the *slot* and the *frame*, which are inextricably linked, as shown below. For transportation in the intra-cluster segment, a number of WDM channels are used to forward traffic from the RNs to the MN. Moreover, a separate wavelength channel is used for conveying all control information. The association between the control and data channels is elaborated in [5]. The control channel is processed in every RN, unlike the data channels, which are transported transparently. To ensure a lossless intra-cluster operation, a two-way reservation mechanism between the RNs and the MN is established by means of the MAC protocol, such as in [6], as follows: bursts arriving in the electronic buffers located in RNs are aggregated into fixed size synchronous *slots* (typically 10 µs to 100 µs). As described in [5], each RN reports to the MN the arriving traffic per destination cluster in terms of number of slots via the control channel. Through the same channel, the MN announces to the RNs the time and wavelength for each slot to enter the ring.

Likewise, in the intra-cluster segment we define as frame (typically 250 µs to 50 µs), the minimum time duration over which a wavelength channel is assigned for interconnecting two particular clusters. The duration of a frame is typically two to three orders of magnitude longer than the duration of a slot, and it is transported under synchronous operation. Therefore, the coordination between intra-cluster and inter-cluster control mechanisms is a two-step process requiring the:

- Efficient grooming of slots into frames
- Lossless transportation of frames from the source to the destination cluster via a number of MNs

In the receiving cluster, the mode of operation is reversed so the slot is extracted in the final destination RN.

Regarding the first step, the coordinated transmission from RNs to MN, apart from avoiding collisions, enables traffic aggregation in a distributed manner: the MN decides the allocations so the slots from all RNs in the ring aiming for the same destination-cluster are marshaled into the same pre-specified wavelength(s). In this way, every cluster is transformed into a large distributed all-optical switch, where each RN acts like an input buffered port controlled by the reservation-based controlled process; statistical multiplexing gains are realized. This operation is shown in Fig. 2 assuming that wavelengths 1 (red) and 3 (blue) are to serve the interconnection between cluster-1 and cluster-3, and the second (green) serves the interconnection between cluster-1 and cluster-4 of Fig. 1. The letter in each slot designates the node that hypothetically contributed that slot. In this way in macroscopic scale, large data bursts toward a single destination cluster are formed on the fly by simply instructing nodes to transmit at the appropriate point of time.

To allow ample time in both intra- and inter-cluster operations, the scheduling is made for the next frame. The scheduler implemented at the MN grants RN requests only when enough traffic to form a large frame is collectively reported or when a time threshold is exceeded. This practice is similar to the various OBS aggregation policies where a burst is generated when the awaiting traffic exceeds a certain limit, or a time limit is exceeded. The allocations are decided and issued periodically with a period equal to the cluster round trip time (e.g., 5 ms for clusters of 1000 km).

Regarding the second step, namely, the efficient inter-cluster transportation, many options covering a wide range of technologies can be considered classified into the following three large categories:

- Semi-static
- Fast automatic lightpath switching
- OBS-like

The first category was covered in [6], where it was shown that this approach exhibits excellent loss and tolerable delay performance for core networks of reasonable size (e.g., five to eight clusters of six to eight nodes) due to the efficient aggregation in the intra-cluster operation. It does not, however, scale beyond such sizes, but this issue will be revisited below.

In the second category of intra-cluster solutions, the wavelengths are dynamically assigned between the clusters. In this case, because the scheduling is always made for the next frame, as soon as the MN has an indication of the total amount of traffic that must be transported from the source cluster to a particular destination cluster, the originating and the recipient MNs proceed to reserve the corresponding wavelengths for a duration equal to integer multiples of the frame duration, which for efficiency reasons must be larger than the round-trip time between the two MNs because the set-up and tear-down times cannot be less than the round trip time.

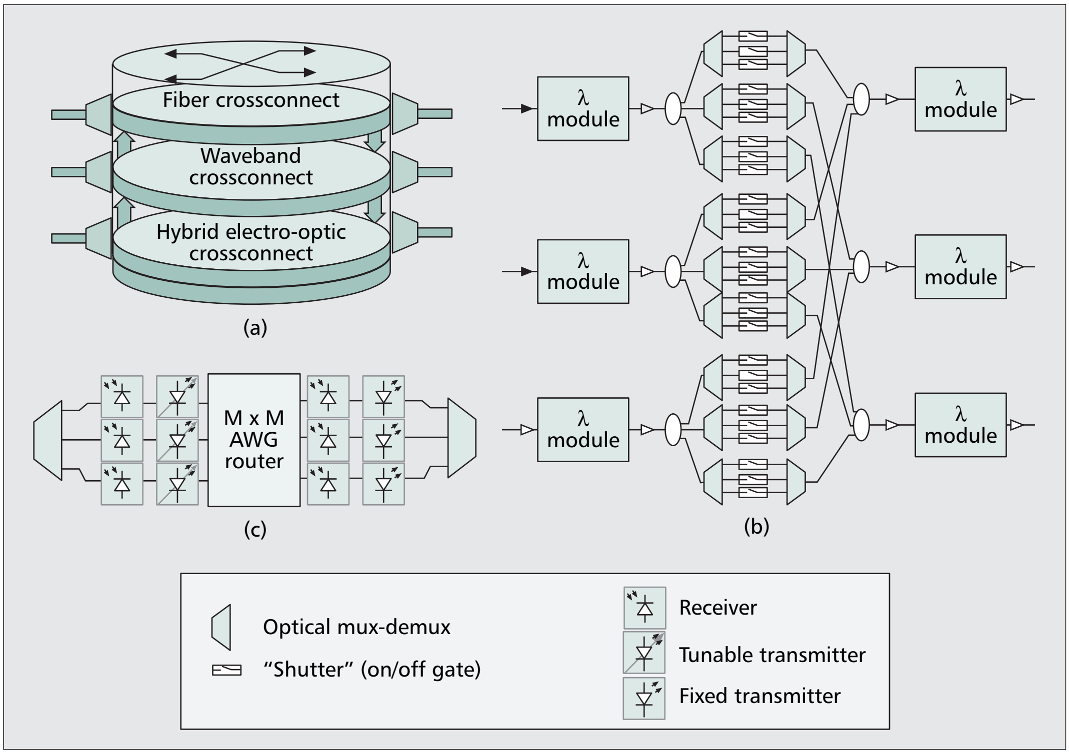

A significant percentage of the existing network infrastructure is based on NG-SONET technology in a ring topology. This allows exploiting the well-known advantages of rings in terms of dimensioning, management, multicast, protection, restoration, and so on. Therefore, this is the starting point of the evolution toward the CANON concept.

■ **Figure 3.** *a) An overview of a CANON master node; b) some of the critical building blocks: the λ-S-λ E-O WXC; c) the λ-module.*

The third case is where the traffic pattern becomes very dynamic so that lightpath switching, no matter how fast, cannot be efficient enough. Under such circumstances, only one-way, OBS-like solutions are much more effective because they can utilize sub-wavelength granularities (i.e., chunks of data of duration much smaller than the round trip time) and achieve multiplexing gain under bursty traffic. This case is investigated in this article, and performance results further highlighting the value of CANON are provided later.

The CANON Node Architectures

A significant percentage of the existing network infrastructure is based on next-generation synchronous optical network (NG-SONET) technology in a ring topology. This allows exploiting the well-known advantages of rings in terms of dimensioning, management, multicast, protection, restoration, and so on. Therefore, this is the starting point of the evolution towards the CANON concept. The first migration step, as shown in [8], is the process of transforming a G.709 frame to an optical slot that can be used in the framework of CANON. This provides the framework and the interface between CANON and legacy systems handling lower switching granularities.

As mentioned earlier, two node types, the RN (used in the intra-cluster segments) and the MN, are employed. The network concept is based almost entirely on transmission with limited, if any, switching. In fact, the main dynamic element is a MAC controlled transmitter and possibly, a reconfigurable add-drop multiplexer (ROADM). The simplest way to exploit such a ring/bus topology is to upgrade NG-SONET add/drop multiplexers (ADMs) to rings with RNs capable to add/drop slots over a number of fibers. In [7], the benefits and the strategy for upgrading an existing infrastructure based on NG-SONET or resilient packet ring (RPR) to an all-optical MAN is shown for two distinctive paths. One is entirely based on transmission, called dual-bus optical ring network (DBORN), for which the optical node is composed of transceivers, a couple of 3-db couplers, and an erbium-doped fiber amplifier (EDFA). The second is based on a λ-selector exploiting fast technologies. Both solutions are cost-effective, modular, and scalable; allowing for a pay-as-you-grow approach so the intra-cluster capacity can grow to very high capacities, deploying these elements only in a gradual manner.

On the other hand, the inter-cluster segment consists of MNs that are coordinating and controlling the entire operation within the clusters and between the different domains. Beyond that, the MNs should cross connect not only the entire capacity of the cluster, but also forward the transit traffic stemming from other clusters toward their destination cluster as shown in Fig. 1. Thus, the MNs are very high-capacity nodes.

To effectively complete the cross-connecting capacity task, the MNs should:

- Scale smoothly to capacities at least an order of magnitude higher than what is available today. For this reason, a highly modular architecture allowing for a flexible and incremental growth in the number of fibers and wavelength channels is required.
- Minimize the bit/byte/packet-level L2/L3 processing and buffering because these contribute the most to scalability restrictions and power dissipation.

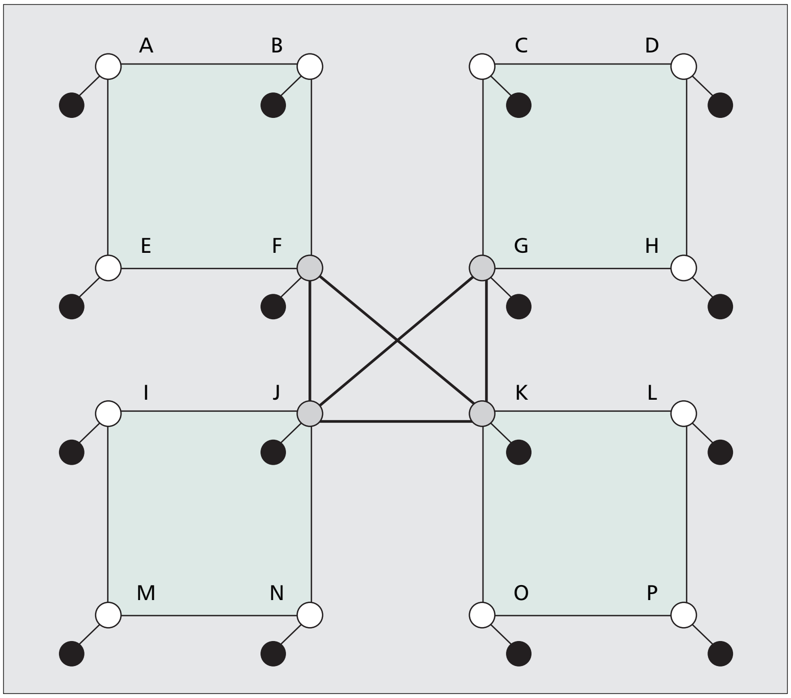

■ **Figure 4.** *Simulated network topology with 16 nodes (which can form four clusters with four MNs).*

- Be able to effectively switch client signals at various rates, ranging from 10 Gb/s up to 100 Gb/s and successfully integrate fiber, waveband, wavelength, and electrical sub-wavelength switching.

Thus, the proposed MN node should be multi-layer as schematically shown in Fig. 3a The three layers are identified as the fiber crossconnect (FXC), the waveband crossconnect (WBXC), and the electro-optic crossconnect (E-O WXC).

The advantage of deploying multi-layer/multi-granular switches in minimizing the number of switching elements was shown in a number of publications ([9, 10]), whereas in [10, 11], the corresponding performance enhancements are quantified. For CANON, these features are of immense importance because the MN must switch traffic at the coarsest possible granularity that still addresses the required functionality. This avoids dedicating expensive lower switching granularity hardware to handle transit traffic, which is the largest part of the traffic through the node. In CANON, the FXC and WBXC are optically transparent switching layers, whereas the E-O WXC is optically opaque. An important aspect is that all three layers avoid any processing of the information content. Here, we propose specific architectures for the WBXC and the E-O WXC, and we explain their mode of operation when these are used as building blocks of an MN in CANON.

The proposed WBXC architecture is the wavelength selective switch (WSS) shown in Fig. 1 of [12]. Although it has been proposed for wavelength switching in [12], the architecture is suitable for any switching granularity. The most attractive feature of the architecture is its modularity for adding either a new fiber and/or a new waveband. As shown in Fig. 3a, the WBXC provides the required connectivity between the FXC and the WXC layers. In the latter case, optimized routing between the WBXC and the WXC can be achieved based on the analysis of [10].

Regarding the WXC as shown in Fig. 3, due to the CANON concept, the slots are well-groomed in frames and wavelength channels, so no wavelength conversion is required for the outgoing WDM channels. However, the following considerations should be taken into account:

- Since the shortest path routing principle is violated, transmission impairments are accumulated, and the WXC layer of the MN should provide 3R regeneration for the outbound channels but without invoking any L2/L3 processing.
- The WXC should cross connect the local outgoing wavelengths with the transit traffic that is handled at the wavelength level. Thus, the WXC must provide wavelength conversion to a number of wavelength channels, as well as to switch individual wavelengths to groom them into wavebands.
- The offering of local add/drop to/from the WBXC and to tributary electronic switches.

The proposed solution addresses all these issues, and a schematic illustration is shown in Fig. 3b in a configuration named λ-S-λ [12], with some important differences as explained below. The λ-S-λ is the next generation of WSS because it allows for wavelength conversion, a feature that is not offered today. The letter "S" stands for space switch and it designates the typical λ-blocker architecture whose schematic layout is given in Fig. 1 of [12]. The letter "λ" designates the λ-*module*, a subsystem allowing for wavelength conversion. The λ-S-λ configuration emerges directly from a λ-blocker by adding blocks, that is, two λ-modules for the λ-S-λ to the already existing structure giving birth to the new functionality.

A schematic illustration of the λ-module is shown in Fig. 3c. The λ-module is a subsystem allowing for wavelength interchange using all optical means. In contrast to [12], the entire block is based on mature electro-optic technology to facilitate rapid deployment: The incoming wavelength is detected, 3R regenerated, and it is buffered to synchronize the outbound from the cluster traffic with the transit traffic. As mentioned, at this stage bit-by-bit processing of the information content is entirely obviated. At the appropriate time, the train-slot is directed to a tunable transmitter, followed by a passive optical router (e.g., arrayed waveguide grating [AWG]), a fixed receiver, and finally a fixed transmitter. With this set up, the incoming wavelengths are interchanged to avoid blocking in the λ-blocker. A second module is required at the output, to ensure unrestricted wavelength-to-wavelength conversion. Also, because the λ-blocker is switching-technology agnostic, different switching family elements with respect to their reconfiguration time can be simultaneously used. The λ-S-λ architecture is rearrangably non-blocking, and it allows for broadcasting or multicasting. It is also both fiber-modular and wavelength-modular. These features ensure that the MN can scale from low to high capacities in a cost-effective way.

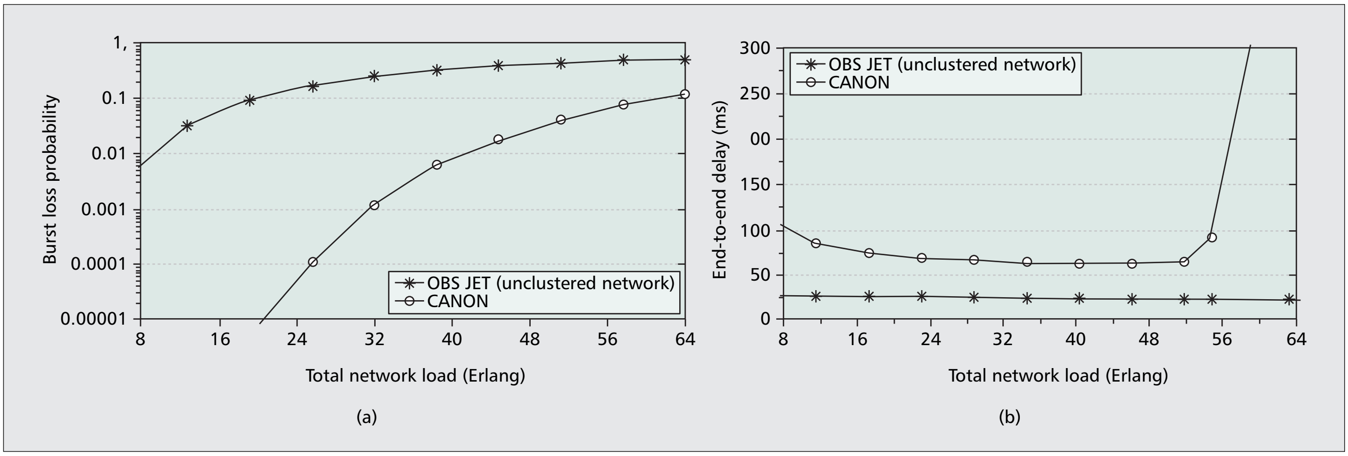

Figure 5. *a) Burst loss probability; b) end-to-end delay for unclustered vs. clustered 16-node core network.*

Performance Evaluation of the CANON Architecture

Extending the work in [6] — where CANON employed pre-provisioned lightpaths in the inter-cluster network — in this section we investigate the additional performance benefits of a dynamic, scalable inter-cluster network, which provides additional multiplexing gains outside the cluster as well. Of course, it is not possible to completely eliminate losses due to burst collisions, as was the case with the logical full mesh assumed in [6], where losses only occurred due to buffer overflow. To evaluate the new dynamic CANON, a generic symmetrical network topology of 16 core nodes as illustrated in Fig. 4 is used under the following assumptions:

- All neighboring nodes have the same distance between them (symmetrically spaced).
- Traffic demand is uniformly balanced among all network nodes, and the generated traffic is destined to all other nodes with uniform probability.
- Each core node receives ingress traffic following a Poisson process through an edge node, where burst assembly is performed.
- Four wavelengths are used by edge nodes, as well as between core nodes in the periphery, whereas sixteen wavelengths are provided in the core links interconnecting nodes F, G, J, and K, because they must be appropriately dimensioned to carry the largest percentage of the network load.

The previous topology represents the case of an optical core similar to existing network backbones like the 16-node Pan-European network or the 14-node NSF backbone. Therefore, link distances were similarly selected, taking into account the average distance among neighboring nodes and the largest span and network periphery of NSF. Figure 5a shows the burst loss probability, and Fig. 5b shows the end-to-end delay versus network load for two cases. In the first case, dynamic resource allocation is achieved by means of a typical OBS one-way reservation scheme (the just-enough-time [JET] signaling scheme proposed in [1]) operating between all 16 core nodes of the network shown in Fig. 4. In the second case, the CANON approach is used in the same network of Fig. 4, now grouping every four neighboring nodes into a cluster to create an equivalent CANON topology for the purpose of achieving a meaningful performance comparison.

An inspection of Fig. 5a and Fig. 5b reveals the basic trade-off between reduced loss and higher delay (due to reservations) that CANON offers over OBS. However, whereas the loss improvement is very significant, the delay penalty is not disruptive for the proper network operation as are the excessive losses, because it remains at acceptable levels for all services (by avoiding the end-to-end reservation). It is worth noting that the low delay of OBS, particularly at high loads, should be treated with caution because it represents a very small percentage of bursts that manage to arrive at their destination, which also explains why delay does not grow exponentially (because lost bursts are not measured!). This also explains the slightly decreasing delay at high loads due to the higher probability of loosing bursts with longer offsets (due to the more contention points) that reduces the average delay.

Conclusions

In this work, the CANON architecture was presented and evaluated as a viable solution to multidomain dynamic core networks able to reconcile guaranteed loss performance of optical slots/packets with statistical multiplexing gains. CANON provides an integrated solution toward efficient, on-the-fly aggregation of traffic in edge nodes and optically performing burst switching in the inter-cluster network. Thus, by hierarchically aggregating traffic inside the clusters through the lossless two-way reservation scheme, one way, OBS-like control for frame/burst switching inside the inter-cluster network can be employed when dynamic fine granularity (i.e., sub-wavelength) resource reservation is required. The overall reduction of contention points (because collisions inside the cluster of origin were resolved preemptively by the control protocol) enables the OBS switching in the inter-cluster network (chosen to avoid further reservation delays) to become far more efficient than in a flat optical core network topology, resulting in excellent burst loss matched with tolerable delay as quantified by the presented simulation results.

ACKNOWLEDGMENTS

This work has been partially funded by the EU project IST-NOBEL.

REFERENCES

[1] C. Qiao and M. Yoo, "Optical Burst Switching (OBS) — A New Paradigm for an Optical Internet," *J. High-Speed Networks*, no. 8, 1999, pp. 69–84.

[2] I. I. Baldine *et al.*, "Jumpstart: A Just-In-Time Signaling Architecture for WDM Burst-Switched Networks," *IEEE Commun. Mag.*, vol. 40, no. 2, 2002, pp. 82–89.

[3] J. Y. Wei and R. I. McFarland, Jr., "Just-In-Time Signaling for WDM Optical Burst Switching Networks," *IEEE/OSA J. Lightwave Tech.*, vol. 18, no. 12, 2000, pp. 2019–37.

[4] L. Dittman *et al.*, "The European IST Project DAVID: A Viable Approach toward Optical Packet Switching," *IEEE JSAC*, vol. 21, no. 7, 2003, pp. 1026–40.

[5] A. Stavdas *et al.*, "A Novel Scheme for Performing Statistical Multiplexing in the Optical Layer," *OSA J. Opt. Net.*, vol. 4, no. 5, 2005, pp. 237–47.

[6] J. D. Angelopoulos *et al.*, "An Optical Network Architecture with Distributed Switching Inside Node Clusters Features Improved Loss, Efficiency, and Cost," *IEEE J. Lightwave Tech.*, vol. 25, no. 5, 2007, pp. 1138–46.

[7] C. Develder *et al.*, "Benchmarking and Viability Assessment of Optical Packet Switching for Metro Networks," *IEEE J. Lightwave Tech.*, special issue on Metro and Access Networks, vol. 22 , no. 11, 2004, pp. 2435–51.

[8] Alcatel 1674 Lambda Gate, tech. description, 3AL 93106 AAAA Ed. 05.

[9] A. Stavdas: "Architectural Solutions Toward a 1,000 Channel Ultra-Wideband WDM Network," *SPIE Opt. Networks*, vol. 2, no. 1, 2001, pp. 51–60.

[10] M. Li and B. Ramamurthy. "Dynamic Waveband Switching in WDM Mesh Networks Based on a Generic Auxiliary Graph Model," *Photonic Network Commun.*, vol. 10, no. 3, 2005, pp. 309–31.

[11] C. T. Politi *et al.*, "Single-Layer Multigranular OXCs Architecture with Conversion Capability and Enhanced Flexibility," *OSA J. Opt. Net.*, vol. 5, 2006, pp. 1002–12.

[12] A. Stavdas, C. Matrakidis, and C. T. Politi, "Migration of Broadcast-and-Select Optical Crossconnects from Semi-Static to Dynamic Reconfiguration and Their Physical Layer Modeling," *Optics Commun.*, vol. 280, no. 1, 2007, pp. 49–57.

BIOGRAPHIES

ALEXANDROS STAVDAS [M'00] received a B.Sc. degree in physics from the University of Athens, Greece, an M.Sc. degree in optoelectronics and laser devices from Heriot-Watt University/St. Andrews University, Scotland, and a Ph.D. degree from University College London, United Kingdom, in the field of wavelength-routed wavelength-division multiplexed networks. He is an associate professor of optical networking in the Department of Telecommunications Science and Technology at the University of Peloponnese, Tripolis, Greece.

THEOFANIS G. ORPHANOUDAKIS [S'93, M'98] received his Dipl.Ing. degree in electrical and computer engineering in 1995 and his Ph.D. in telecommunications in 1998, both from the National Technical University of Athens (NTUA). He is currently with the University of Peloponnese as a visiting assistant professor. His research interests include design and implementation of high-speed networking systems, performance evaluation of MAC protocols, traffic engineering, and network processor architectures. Prior to joining the university, he was with Ellemedia Technologies as a technical manager.

ANDREW LORD graduated from Oxford University in physics in 1985. Then he joined the optical communications team at British Telecom (BT), assisting BT procurement of subsea systems, TAT12, 13, 14, and subsequently, BT's UK core network. He now leads an optical technology team designing BT's future optical core network.

HELEN C. LELIGOU received Dipl.Ing. and Ph.D. degrees, both in electrical and computer engineering, from NTUA in 1995 and 2002, respectively. Since 2007 she has been a lecturer at the Technological Educational Institute of Halkida. Her research interests are in the area of access control mechanisms in broadband networks including hybrid fiber-coaxial (HFC), passive optical network (PON), WDM, metro, and core networks. She has participated in several EU-funded research projects in these areas.

KONSTANTINOS KANONAKIS received his Dipl.Ing. degree in electrical and computer engineering in 2004, and his Ph.D. in telecommunications in 2007, both from NTUA. His main research interests are in the area of traffic engineering, and architectures and control protocols for optical networks and their implementation in hardware. He is participating in the European Union projects MUSE and NOBEL.

CHRIS MATRAKIDIS received a degree in electrical engineering from the University of Patras, Greece, and a Ph.D. degree in digital communications from University College London. He is currently lecturing with the Department of Telecommunications Science and Technology of the University of Peloponnese, Tripolis, Greece. From 1995 to 1999 he was a research fellow at the Telecommunications Research Group of UCL.

ANDREAS DRAKOS [S'04, M'08] received a B.Sc. degree in telecommunications in 2006 from the Department of Telecommunications Science and Technology of the University of Peloponnese, where he has been pursuing his Ph.D. since 2007. His main research interests are in the area of simulation of optical networks, with emphasis on the performance evaluation of WDM networks and switching node architectures.

JOHN D. ANGELOPOULOS [S'92, M'95] received a degree in electrical engineering from NTUA in 1973, an M.Sc. degree in automatic control from Nottingham University in 1977, and a Ph.D. degree in telecommunications from NTUA in 1992. He is currently a professor in the Automation Engineering Department of the Technological Institute of Piraeus. His research interests cover most aspects of high-speed communications with a special focus on shared medium access protocols and implementations.