Scheme for performing statistical multiplexing in the optical layer

Alexandros Stavdas

Department of Telecommunications Science and Technology,
University of Peloponnese, Tripolis, Greece

astavdas@uop.gr

Helen-Catherine Leligou, Konstantinos Kanonakis,
Charalambos Linardakis, and John D. Angelopoulos

School of Electrical and Computer Engineering,
National Technical University of Athens, Athens, Greece

{nelly, kokan, chlin, jangel}@telecom.ntua.gr

The proliferation of data-centric networks has intensified the research efforts for development of architectures and technologies supporting fast reconfigurable optical networks accommodating any temporal and (or) spatial traffic fluctuations without WDM channel overprovisioning. The novel core network architecture presented in this paper is based on partitioning a meshed network into a number of clusters. Within each cluster the architecture exploits metropolitan area concepts (MACs) to achieve statistical multiplexing gain by transparent transportation of fixed duration slots. These short slots, which are controlled by means of a MAC protocol, are combined to form larger frames that are then exchanged between the clusters. Specific nodes that serve as gateways between clusters undertake the coordination of frame switching and end-to-end routing. Thus, traffic within a particular cluster enjoys reservation-based no-loss transport, while for the communication between clusters either robust two-way reservation schemes or even optical burst switching (OBS) approaches are possible. The end result is an architecture that supports statistical multiplexing within the optical layer with good performance and reasonable complexity. © 2005 Optical Society of America

OCIS codes: 060.4250, 060.4510.

1. Introduction

The telecommunications systems that have been installed so far, which exploit the continuous advances in optical technologies for SDH/SONET as well as wavelength-division multiplexing (WDM) transmission, have been quite successful in accommodating demand for low- and high-bandwidth services. However, traffic statistics indicate that the volume of data traffic is still increasing rapidly and much faster than that of voice traffic. Given the degree of burstiness that is already witnessed in the metro part and the continuing broadband evolution in the access part (via xDSL or FTTx), it is expected that further proliferation of data services will make bursty traffic profiles much more prominent in all network sections including the core. Under such traffic conditions, operators cannot afford static overprovisioning of low-utilization WDM channels anymore, leaving as the only other viable alternative the deployment of dynamically reconfigurable networks exploiting statistical multiplexing principles.

To tackle this problem, revolutionary solutions have been proposed for both the core and the metro part of the network. The proposed solutions exploit the strength of both optics and electronics, allowing transport and switching of data chunks of various sizes transparently, while targeting a fully dynamic network control. The optical burst switching (OBS) architecture [1] attempts to adapt the IP paradigm to optical core networks to achieve high multiplexing gain and high performance in terms of delay. It allows a burst to be sent without any prior indication that can be successfully switched by a nearly bufferless core network designed to operate in the way best suited to optical technology. However, its performance suffers from significant burst losses, requiring retransmissions via TCP and causing increased delays and unpredictable loading along the full path (including the access part). Two-way reservation schemes have also been proposed [2, 3] to enhance the performance of the OBS architecture in terms of losses; however, this increases the observed delay values, since each burst waits for a positive acknowledgement that will travel to the destination successfully.

On the other hand, dynamic networking schemes for Tbit/s capacity metropolitan area networks (MANs) have been proposed in which the payloads are switched and routed transparently end to end while the control information is transmitted using a separate (common) channel and processed in every node without using high-speed electronics. The control mechanism is based on a reservation metropolitan area concept (MAC) protocol, ensuring collision-free operation without compromising performance in terms of end-to-end delay. This approach has been pursued within the framework of the project Data And Voice Integration over DWDM (DAVID) [4], which exploits WDM technology in ring topologies to design and implement flexible and dynamic MAN systems. Since the geographical area covered from a MAN is less than 200 km, reservation delays are not prohibitive, thus achieving both high efficiency and low latency, while still avoiding any buffering in the optical domain as well as any optical–electrical conversion of payload data.

In this paper, a novel approach that builds on the advantages and sidesteps the drawbacks of the already proposed schemes is adopted for the design of a dynamically reconfigurable core network offering efficient utilization of the available resources by means of statistical multiplexing. Thus, high performance in terms of losses and delay is reconciled with low implementation cost. The latter depends on the required switching speeds and on the employed control schemes in terms of synchronization needs and algorithmic complexity. The paper is organized as follows: in Section 2 the related work is discussed, while in Section 3 the new architecture, including its control aspects, is presented. Assessment of the architecture is included in Section 4, and conclusions are drawn in Section 5.

2. Methods for Achieving Statistical Multiplexing in Core and MAN

2.A. OBS and Two-Way Reservation Architectures

The concept of OBS aims at allowing access to optical bandwidth in DWDM networks at fractions of the optical line rate to improve bandwidth utilization efficiency. Each node transmits in a dedicated wavelength a control message prior to each data burst that informs the traversing nodes about the required optical configuration for switching of the payload burst, the burst transmission time, and its length. The control wavelength is the only one converted to the electrical domain and processed in every node. Thus, statistical multiplexing is achieved, which along with the all-optical transmission of data represents the most important advantage of the concept.

However, the basic drawback of the concept is the high loss ratios that are observed, since the probability of reaching the destination dramatically decreases with the number of nodes that need to be traversed. This greatly affects the observed delay, which is equal to the burst propagation time, augmented by the round-trip time as many times as there

are unsuccessful attempts. To this delay, the terminal processing time must also be added for schemes that do not attempt retransmissions from the OBS edge node but rely on TCP for retransmissions. Although several reservation and contention resolution schemes (Refs. [5–7]) have been proposed to enhance OBS performance, the improvements are considered rather marginal. Another disadvantage is the high implementation cost, which stems from the need for fine synchronization between control and burst and for fast node reconfiguration to follow the variable size of the burst. Hence, the use of expensive components, such as the semiconductor optical amplifier, is dictated.

Two-way reservation schemes have been proposed to avoid losses: the transmission occurs only if a positive response has been received back from the destination node. Hence, no burst is lost. However, the access delay experienced by the injected bursts is greater than the round-trip time, which is not at all negligible for core networks [3]. Thus, the lower transit delay limit for two-way reservation schemes is equal to the round-trip propagation time plus the burst one-way propagation time augmented by the scheduling time, reaching the order of hundreds of milliseconds in total for networks of several thousand kilometers as is necessary to cover, for example, the continental United States or the EU.

The two-way reservation concept becomes promising only when the time scale of requests and transmissions is of the order of tens of milliseconds, assuming that node-to-node throughput will allow for building bursts of millisecond duration. However, this is not very likely to happen for small networks that have tolerable round-trip times; i.e., to assemble a burst of millisecond duration toward one node, a node needs much more time, causing IP packets to wait for a long time for a burst to be assembled, although once assembled the burst will travel with almost round-trip time delay. Schemes that allow pipelining of the burst assembly phase and the reservation phase have been proposed, but they sacrifice efficiency to delay.

2.B. Metro Concepts

The architecture adopted in the IST project DAVID for the MAN consists of several WDM rings interconnected at a central point, the hub, as shown in Fig. 1. The traffic is converted from the electrical to the optical domain at the source metro node, transmitted by means of fixed size slots and converted back to electrical only at the destination metro node. Virtually lossless end-to-end packet–slot transmission of the payload without introducing optical buffers is possible because of the reservation-based control scheme, which, in turn, is possible only at metro distances. The MAC protocol controlling access to the slotted rings employs a dedicated control wavelength with fields for identifying the wavelength and the time slot in which each fixed size payload should be transmitted, ensuring collision-free operation. The matching of the information of the control slots to the corresponding equal-size data slots is shown in Fig. 2, where just three data wavelengths are assumed. Each control slot contains information for the data that will arrive on all the supported wavelengths in the following slot time. The adoption of a fixed-size slot alleviates the control scheme in the hub switch matrix. Since the MAN is effectively a one-hop scheme, the synchronization requirements can be easily coordinated from the hub because only coarse synchronization, on slot level, is needed. So the overall scheme can be likened to a small OBS with only one hop that can achieve both lossless and bufferless operation by strict and accurate closed-loop control of the reservation-based MAC protocol, an approach made possible only because of the small metro dimensions.

The payloads injected from the ring nodes are switched in the hub without undergoing any optical–electrical conversion because of the "permutations" mechanism, presented in Refs. [4, 8]: input–output ring combinations are decided at the hub and announced in the control slot transmitted on the rings so that all the ring nodes insert a payload destined for the announced ring. (For example, for the first slot shown in Fig. 1, all nodes of ring 1

will insert traffic destined to ring 1, all nodes of ring 2 will insert traffic destined to ring 3, etc.) When the respective slots arrive back at the hub one round-trip time later, they are switched to the already decided output; hence the payload will reach the appropriate ring without undergoing any electrical–optical conversion at the hub. Virtual queuing per destination ring is performed in every node, and IP packets are aggregated–segmented to fill the system fixed-size slot. The input–output permutations are decided either dynamically (based on explicit reservations) or statically.
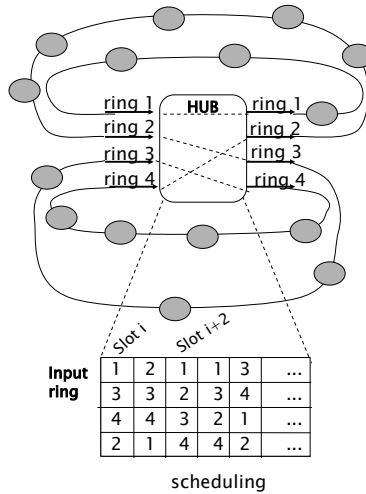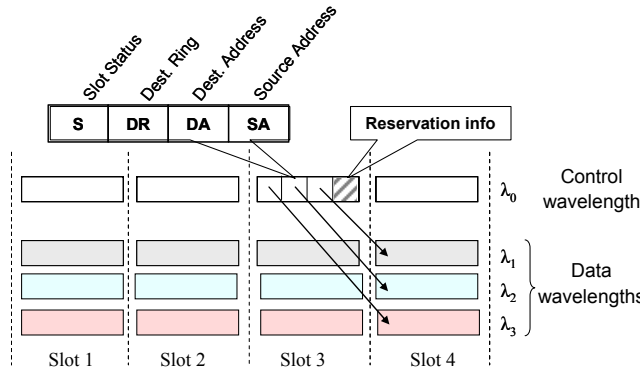

Fig. 1. DAVID metro network architecture.

Fig. 2. Control slot information and its relation to the data slots in DAVID architecture.

The operation of the MAC protocol in each single ring is very simple and based on so-called empty slot access. A node with a slot to send examines the information carried in the control slot (Fig. 2), seeking an empty slot (indicated in the slot status field) destined for the appropriate ring (indicated in the destination ring field). When a node succeeds, it updates the slot status, the source, and the destination node address fields in the control slot region that correspond to the correct wavelength. (Priority is provided to transit traffic.) Thus the payloads travel a safe trip to their destination in the optical domain throughout the system.

To guide the hub, calculate an input–output permutation distribution that matches the

fluctuations of the offered load. The reservation information field includes one field per supported ring. Each node updates the received information, adding the value of its own queue lengths. With the collective reservations a node can grab a slot destined for the right destination ring passing in front of the node that was created by previous reservations by other nodes without waiting for the round-trip time of its own reservation, which in turn may be used by another node. Very small waste is observed since reservations and actual arrivals are equal in the long term. An average transit delay of the order of just one round-trip time has been measured by use of computer simulations, because of the collective character of the reservations. For example, for a metro ring of 80 km, delays of the order of 0.4 ms are experienced. (Successful implementation of the MAC controllers both for the ring node and the hub is presented in Ref. [9].)

The whole system manifests an optimized interplay between centralized control (at the hub) and local decisions in the MAN nodes offering statistical multiplexing gains. It must be pointed out that the system relies on a unidirectional ring topology for which all the resources (fiber, wavelengths, time slots) are shared over the same path, allowing all nodes to optically add–drop packets–slots over the entire WDM channel spectrum. In addition, the system requires fast optical technology in the form of MAC-controlled burst-mode transmitters in the ring nodes and SOA-based (semiconductor optical amplifier) optical switches in the hub; thus, in essence, it operates as an optical packet router. The local reconfiguration speed of these devices is less than the guard band between packets, making the whole operation feasible.

It is important to stress that the control scheme can operate for any slot duration, i.e., even if optical switches with slower reconfiguration rate are employed in the hub. So, the node and the hub reconfiguration speeds can be decoupled and two switching granularities are explicitly introduced: a fast reconfiguration that is associated with the MAC-controlled transceivers in the ring nodes, allowing them to insert their traffic in short-duration slots, and a slowly reconfigurable switch at the hub handling larger slot aggregates.

In Section 3 a broader architecture is presented that builds on the above architecture and the concept of small slots written by fast devices and forming longer frames by trains of such slots that can be switched by slower devices with much lower cost. The frames can reach durations of the order of tens of milliseconds. In contrast with the above metro system, this architecture can cover much longer distances where it is no longer possible to use a fully reservation-based approach for delay-sensitive traffic. Hence, the single hub is not sufficient anymore. However, the basic idea of reservation-based control by the empty slot protocol is used at parts of the network covering shorter distances, while other methods are used for the frames that carry traffic to longer distances.

3. Proposed Architecture

The architecture proposed in this work expands and develops the previously presented concepts to address multihop core networks with meshed connectivity, such as the one illustrated in Fig. 3(a). The nodes employ no optical buffering in the transit paths, and each node can store in electrical buffers only its own local traffic in concert with the OBS principle, where a bufferless optical network is used with electronic buffering in the network periphery only. To overcome the OBS limitations presented in Subsection 2.A, the whole core network (which may have a size measured by its longest round-trip time of almost 100 ms), consisting of a few tens of nodes is partitioned in subnetworks of smaller dimensions based on topological and traffic parameters. Each subset contains a cluster of nodes that are closer to each other or have more traffic exchanged among them or belong to the same administrative domain (e.g., a country). Although such topological or traffic relations will lead to better performance as will become apparent, the clusters can even be defined arbitrarily, as is essentially the case when national borders are the deciding factor. It is pointed out that

this node clustering can be either static or dynamic. In the latter case, the topology may change on a time scale ranging from days down to seconds. In the proposed scheme, one node per cluster serves as a master node. The role of the master node is dual: to act as the gateway to other clusters, thus interconnecting the cluster with at least another cluster, and to coordinate the transmission of all other nodes of its own cluster.
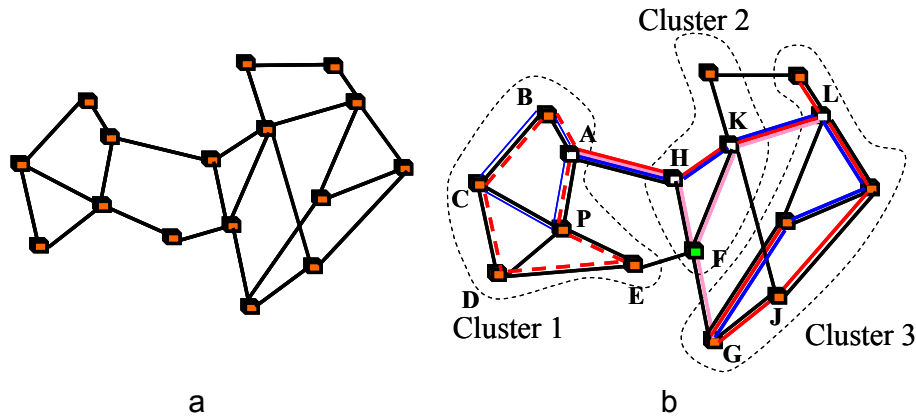

Fig. 3. Network with meshed connectivity (a) and clustering of nodes (b).

Two switching granularities are introduced: the slot and the frame. In the interior of each cluster, both time-division multiplexing and WDM are employed under the supervision and coordination of the master node. All the other nodes within the boundaries of a cluster are allowed to inject traffic in slots of fixed duration. Externally, the master node negotiates a number of wavelengths for the duration of a frame for providing interconnection between its cluster and another cluster. It is imperative to point out that within a cluster the shortest path routing is not followed. Instead, the interior topology of a cluster is decomposed to a fixed number of partially overlapping, ring networks so that a number of nodes share the same WDM channels, as will be explained in more detail in Subsection 3.A. The rationale for coordinating the internal and external scheduling separately but simultaneously is to avoid the drawback of OBS, which is a quite significant loss ratio. The reason for this is the contentions of optical packets for the same exit port under quite limited buffering resources, which leads to collisions. Since several nodes need to be crossed in a continental network, the probability of reaching the destination falls steeply with the number of nodes crossed. Obviously great performance improvement can be reached by reducing the number of nodes where contention is unavoidable, and this is the rationale for the architecture presented in this paper.

In contrast with classical OBS, under the proposed scheme collisions inside the cluster are avoided by reservation-based closed-loop control as in DAVID rings (made possible by the smaller cluster size). As in OBS, however, each wavelength is still shared among a number of nodes offering statistical multiplexing gains in the optical layer. In addition, because of time multiplexing the whole traffic from different nodes of a particular cluster destined for another given cluster forms longer transport units, i.e., frames, enforcing aggregation of traffic towards a destination cluster. Thus, the clustering is a divide-and-conquer approach.

Any payload slot has to wait in the electronic buffer for the outcome of internal and external scheduling. As soon as the go-ahead message is received, the slot is released in the optical transport network and travels to the final destination without any optical buffering and no collisions.

For intercluster (i.e., among clusters) transport, each master node negotiates with other master nodes an end-to-end path dedicating a whole wavelength (or a wavelength band) for the duration of a frame. An example of such a path is shown in Fig. 3(b), where node A, the master node for cluster 1, is interconnected to cluster 3 (master node L) via cluster 2 (master node H, K). For duality reasons, the meshed node interconnection topology of cluster 3 is decomposed into a group of rings in order to distribute the individual slots within a frame to the final destination node.

As mentioned above, within each subnetwork, lossless statistical multiplexing is enforced by means of a reservation-based MAC protocol while the wavelength negotiation between the master nodes could be based on two-way reservation schemes such as fast circuit switching. A reservation approach becomes feasible for communication between clusters (master nodes) because of the traffic aggregations into much fewer paths and longer frames. However, a typical OBS control scheme can also be considered since the number of hops is significantly lower; hence, collision probabilities are dramatically reduced compared with the case in which OBS is applied to all the nodes of the core network.

The adoption of fixed-size transport units (slots, frames) allows for easier control, although it departs from the bit-level synchronization legacy of electronic switching systems. In the proposed scheme coarse synchronization on the slot level is required only for all operations in any transit node. As was demonstrated in Ref. [4], this operation can be achieved in the optical layer by use of relatively inexpensive components. The feasibility of synchronization on the frame level was also dealt with in Ref. [10]. It is pointed out that, although the entire intracluster operation is based on fixed time slots with duration of the order of microseconds, variable-length bursts that are integral multiples of this basic time slot can also be transported as trains of slots. Moreover, since the duration of a frame is several tens of milliseconds, microelectromechanical systems technology can be used to construct the switching matrix, hence reducing their cost while still allowing for statistical multiplexing and all-optical payload transmission at high performance.

3.A. Detailed Description of Intracluster Operations

The partitioning of a meshed network in a group of clusters and the subsequent differentiation of the control mechanisms is an essential step in our proposal. In every node within a cluster, per-destination cluster queuing is performed and all the traffic generated by the nodes of a certain cluster toward another cluster is aggregated, forming large frames for intercluster switching. A crucial aspect of the intercluster coordination and routing is that each cluster should be viewed as virtually decomposed in a number of fixed topology rings. For example in Fig. 3(b), two such rings are illustrated for Cluster 1: the BCDEPA and the BCPA with the same or different master nodes. Note that one node might belong to one or more such rings. The similarity of this scheme to the one presented in Subsection 2.B is apparent.

For demonstration purposes, a four-slot-long frame emanating from Cluster 1 is schematically illustrated in Fig. 4. A letter has been placed within each time slot indicating the node of Fig. 3(b) that inserted this packet during this slot. On the left side of each wavelength, the destination cluster is shown. For example, in λ_1 node C has inserted a burst destined to Cluster 2 with duration equal to three slots, in λ_2 nodes B, E, and P have inserted 1-, 1-, and 2-slot-long bursts destined for Cluster 3 while nodes C and P have inserted 1- and 3-slot-long bursts in λ_3, respectively, destined for Cluster 2. The gaps between the transmissions of different nodes indicate the guard bands between the slots as well as the fact that the nodes operate in the burst mode. Overall, the filling factor of each individual wavelength is higher than could otherwise have been if slots or packets were inserted without coordination. The end result is that traffic toward the same cluster is multiplexed into frames. As in Subsection 2.B, the control information is transported over a dedicated

wavelength channel in slots of equal duration to the data slots, so that a one-to-one relation is established.

To guarantee loss-free operation, the decision on the time slot, the wavelength, and the fiber (ring) on which a packet is transmitted is based on a MAC protocol. One can identify different MAC protocols suitable for this operation. In one approach, the control mechanism adopted in DAVID is modified for the particular situation. The design of the control slot is shown in Fig. 5. The control information includes the destination and source node addresses, the status (empty or occupied), and the destination cluster of the slot for each of the k supported wavelengths, as well as reservation information for all the N possible destination clusters. Based on the slot status and on the destination cluster information, the node decides on the wavelength and time slot to use, giving priority to transit traffic. Each node adds to every QL field, the length of its queue destined to the respective cluster expressed in slots, so that when the reservation field arrives at the master node, the node is informed about the total number of slots required for satisfaction of the relevant traffic. The delay observed even under highly bursty traffic, based on the simulation results reported in Ref. [8], is close to a few round-trip times. For a ring circumference length of 500 km, this will be 5 ms (for two round-trip times) to 7 ms. The offset between the time a node decides that a slot payload directed to a particular destination node is ready for transmission and the actual time where the packet is transmitted is an important networking parameter and will be further analyzed in the next section.
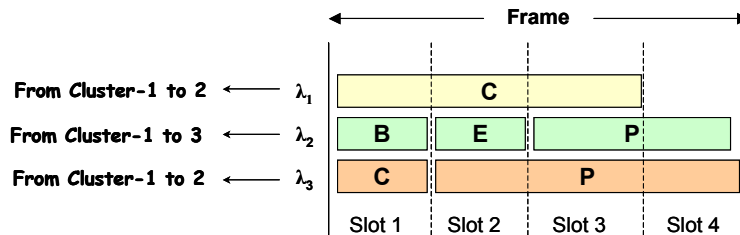

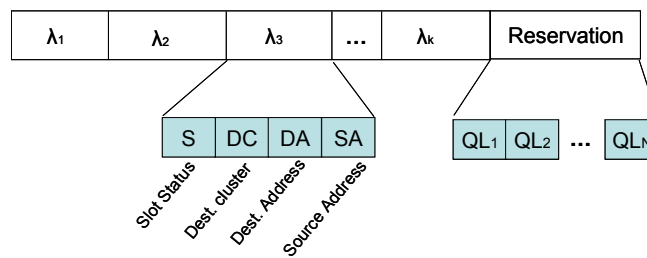Fig. 4. Concept of the frame consisting of multiple slots.

Fig. 5. Control slot format in the case of a ring-topology cluster.

3.B. Detailed Description of Intercluster Operations

The scheduling of interconnections between the master nodes can be decided either in a centralized or in a distributed manner. In the first case, a master node can act as the interconnection scheduler, and all other master nodes place requests to it for {wavelengths, fiber}. This is done as soon as the nodes have adequate call requests, allowing them to estimate the

requested capacity, i.e., the number of wavelength channels needed for gratifying the traffic toward a particular cluster in the next frame. After the relevant negotiation, the interconnection pattern will eventually be scheduled. This procedure is performed on a frame-by-frame basis, and the process is dynamic. Assuming that the maximum distance between master nodes is equal to 4000 km (i.e., propagation round-trip delay around 40 ms) and that the frame duration is 50 ms, it is feasible for a master node to collect all the requests, process them, and distribute the intercluster traffic scheduling for the next frame.

With a distributed control algorithm, it is possible for a master node to select to service a request as transit node while the destination master node decides to reject the request. In contrast, with a centralized control via a single central scheduler, this can be avoided and an optimal solution can be found, leading to higher efficiency. To dynamically allocate the resources adapting to the traffic fluctuations, the explicit requests placed by the master nodes are used to calculate the interconnection pattern for every frame-duration time. Without loss of generality, it is postulated that the number of wavelength channels are the same across all the clusters; otherwise the master nodes are equipped with the corresponding wavelength converters.

4. Control Signaling

The exchange of messages from the time a data burst arrives at a source node is shown in Fig. 6. The source node transmits a request message (at T_{rt}) toward the master node of the cluster. The master node gathers all the received requests from the whole cluster and based on them it is forming the requests for wavelength channel allocation. These requests are then transmitted (as T_{crt}) to the destination master node, possibly through other master nodes. There is an offset time T_1, between the arrival of the request message at the master node and T_{crt} since the master node needs to estimate the requested number of wavelength channels needed to interconnect two clusters. If T_1 is too short, unexpected bursts will be only partially accommodated. On the other hand, if it is too long there will be no time for making the subsequent scheduling tasks and the next frame will be discarded. Therefore, the requests that have arrived also need to be assessed against long-term traffic pattern demands between the two clusters.
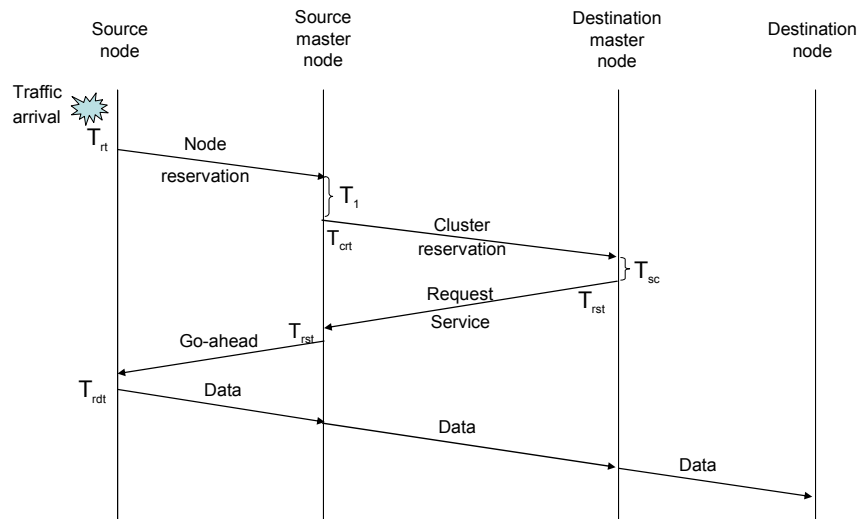

Fig. 6. Signaling and data transmission.

The destination master node based on the requests it has received from all clusters makes a decision after T_{sc} regarding if and how these requests will be serviced. The request service message travels back to the source master node and the go-ahead message is forwarded to the source node that releases the packets (or cancels in case of a failure in setting up a path). The end result is that the total delay equals the reservation time plus the propagation time of the payload, or in more detail, a round-trip time between source node and destination master node, plus the scheduling time of the burst, plus the propagation time between source and destination nodes. This means that the delay is approximately 70 ms (assuming 500 km for the cluster periphery and 4000 km distance between the master nodes) plus the service scheduling time, which depends on the employed logic and the offered load and can be statistically bounded.

5. Conclusions

The attempt to enrich the core networks with scalability and fast configuration led to new architectures using optical burst switching along paths preconfigured just ahead of the burst. However, these random time bursts contending for the same port suffer heavy losses because of the poor buffering capability of optical switches. A novel architecture that still allows for all-optical transmission of the payload but with much-reduced losses and higher efficiency has been presented. The principle it employs is to split the core network nodes into subgroups such that inside each subgroup reservation-based concepts can be used to achieve lossless data transport while bursts transferred among subgroups have much-reduced collision probabilities because of the smaller number of nodes. The adoption of fixed-size frames for the data exchange between master nodes allows the deployment of slowly reconfigurable switching matrices and alleviates the need for fine synchronization, hence reducing the implementation cost. The overall system features high statistical multiplexing gain and nodes implemented directly in the optical layer, allowing better performance in terms of loss and delay.

Acknowledgment

References and Links

[1] C. Qiao and M. Yoo, "Optical burst switching (OBS)—a new paradigm for an optical internet," J. High Speed Netw. **8**, 69–84 (1999).

[2] M. Düser and P. Bayvel, "Analysis of a dynamically wavelength-routed optical burst switched network architecture," J. Lightwave Technol. **20**, 574–585 (2002).

[3] J. Y. Wei and R. I. McFarland, "Just-in-time signaling for WDM optical burst switching networks," IEEE J. Lightwave Technol. **18**, 2019–2037 (2000).

[4] L. Dittmann, C. Develder, D. Chiaroni, F. Neri, F. Callegati, W. Koerber, A. Stavdas, M. Renaud, A. Rafel, J. Solé-Pareta, W. Cerroni, N. Leligou, Lars Dembeck, B. Mortensen, M. Pickavet, N. Le Sauze, M. Mahony, B. Berde, and G. Eilenberger, "The European IST project DAVID: a viable approach towards optical packet switching," IEEE J. Sel. Areas Commun. **21**, 1026–1040 (2003).

[5] K. Dolzer, C. Gauger, J. Späth, and S. Bodamer, "Evaluation of reservation mechanisms in optical burst switching," AEU Int. J. Electron. Commun. **55**, 1–8 (2001).

[6] V. Vokkarane, Q. Zhang, J. Jue, and B. Chen, "Generalized burst assembly and scheduling techniques for QoS support in optical burst-switched networks," in *Proceedings of IEEE Globecom 2002* (IEEE, 2002).

[7] M. Yoo, C. Qiao, and S. Dixit, "Comparative study of contention resolution policies in optical burst-switched WDM networks," in *Terabit Optical Networking: Architecture, Control, and Management Issues*, J. M. Senior, S. Dixit, and C. Qiao, eds., Proc. SPIE **4213**, 124–135 (2000).

[8] Ch. Linardakis, H.-C. Leligou, A. Salis, and J. D. Angelopoulos, "Control of slotted traffic among interconnected WDM rings by explicit reservation," presented at Photonics in Switching 03, Versailles, France, September 3, 2003.

[9] Ch. Linardakis, H. C. Leligou, A. Stavdas, and J. D. Angelopoulos, "Implementation of medium access control for interconnecting slotted rings to form a WDM MAN," J. Opt. Netw. **3**, 826–836 (2004), http://www.osa-jon.org/abstract.cfm?URI=JON-3-11-826.

[10] J. Ramamurtham and J. Turner, "Time sliced optical burst switching," in *Proceedings of IEEE INFOCOM 2003* (IEEE, 2003).