Lightpath bundling and anycast switching (LB+AS): a new paradigm for multilayer optical networks

Pablo Pavon-Marino*, Jose-Luis Izquierdo-Zaragoza
Technical University of Cartagena (Spain)
{pablo.pavon, josel.izquierdo}@upct.es
(*) Corresponding author

Abstract

In IP over WDM multilayer optical networks, IP routers are interconnected by all-optical channels called lightpaths, of typical rates of 10, 40 or, more recently, 100 Gbps. In this context, lightpath bundling (LB) and anycast switching (AS) are control plane and data plane techniques respectively, both of them to be implemented in the IP routers. LB permits grouping a set of lightpaths between two nodes that follow a common route, so that they are perceived by the IP layer as a single virtual link of aggregated capacity. In its turn, AS consists of instructing the router to implement a per-packet-granularity balancing of the traffic among the lightpaths in the bundle, reducing the packet delay and the buffering requirements in the node. This happens transparently to the IP layer that, because of the LB configuration, sees the bundled lightpaths as a single entity. In this paper, we propose the combined application of the LB and AS techniques, as a new paradigm (LB+AS) for optical networks. Applying the LB+AS concept requires seamless changes in the electronic equipment and no changes in the optical infrastructure. We present a case study which shows the significant performance and cost benefits that LB+AS can bring to the network and its inherent scalability. In addition, we discuss other potential advantages of LB+AS, and related open research lines.

Keywords: Multilayer optical networks, lightpath routing, high-performance packet switching

Introduction

Optical multilayer IP over WDM networks are an enabling technology for building the core and metro high-capacity networks in the Internet of the future. In IP over WDM networks, the IP routers are connected by transparent all-optical channels, called lightpaths. A lightpath occupies a wavelength in the traversed fibers and is optically switched at intermediate nodes. For the IP layer, the set of established lightpaths comprises a Virtual Topology (VT) [1], and packet-switched IP traffic is directly routed over the VT. Thus, each lightpath corresponds to a virtual link between two routers: independently on the actual sequence of fibers traversed by the lightpath, the IP layer sees it as a pipe to transmit packets, with typical rates of 10, 40, or 100 Gbps.

The success of multilayer networks is built on their ability to efficiently combine the electronic packet switching for a finer traffic granularity and the optical circuit switching for a coarse granularity. These two modes of switching are conducted by separated equipment in the network. The IP/MPLS routers perform a per-packet processing, buffering and switching of the IP packets, while ROADMs (Reconfigurable Optical Add-Drop Multiplexers) implement the optical add, drop and transparent bypass switching of the lightpaths. Electronic packet switches/routers are costly, bulky and power-hungry (e.g. CISCO CRS-3, 322 Tbps, 80 chassis, approx. 700 kg/chassis, consumes more than 1 MW). Their complexity originates in the need of a fast per-packet processing at very high speeds. Traffic to and from each lightpath is processed by line cards operating at wire-speed (10/40/100 Gbps), with large buffers for storing the packets. Line cards are interconnected by a switch fabric, a complex piece of hardware that permits the delivery of IP packets among the line cards.

The traffic engineers at the IP layer face a trade-off when planning multilayer networks, between lightpath utilization (fraction of the lightpath capacity that is used in average) and network delay performance. This is because high lightpath utilizations can result in an unacceptable degradation of the packet delay/packet losses performance, caused by the traffic burstiness in real networks. This effect has more importance with multimedia traffic, more sensitive to delays, which is becoming dominant in the network. A recent work [2] suggests that the combined application of the *lightpath bundling* (LB) technique in the control plane, and the *anycast switching* (AS) technique in the data plane of the IP routers, can positively bias the utilization versus delay trade-off in the network. In other words, when applied to a multilayer WDM network, this combination permits carrying more traffic with the same delay performance, or carrying the same traffic but significantly decreasing the end-to-end delay and the buffering requirements in the nodes. The lightpath bundling technique consists of bundling together in the network those lightpaths which have a common input and output node and follow the same route (at any wavelength), so that they appear to the upper layer (IP, Ethernet or MPLS) as a single link of aggregated capacity. Anycast switching (AS) means that the output lightpath of a packet can be freely chosen by the router among all the lightpaths in the bundle. This decision is transparent and unrelated to the IP layer, which, in fact, sees all the lightpaths in the output bundle as a unique virtual link, thanks to the LB technique. It is the switching architecture the one that is taking advantage of this new degree of freedom (with respect to unicast switching) to forward the packets to *any* of the output line cards corresponding to the lightpaths bundled. Different switching architectures can implement AS in different forms, and we believe this can be a matter of active research. However, the essence of AS technique is to enforce a *per-packet* granularity balancing of the traffic among the bundled lightpaths, to enhance the statistical multiplexing gain. Once the IP packets reach the output line cards, the lower-layer framing and transmission of the packets progress as if a conventional unicast switching had occurred. In this sense the AS technique is independent of the optical layer technology in the network.

Fig. 1 helps us to illustrate the LB+AS concept. Fig. 1(a) shows an optical multilayer network with four sites. Sites 1, 3 and 4 have an IP router and a ROADM, while site 2 only has a ROADM, and does not process IP packets. Five lightpaths are established at 40 Gbps rate each: three lightpaths from router 1 to router 3, and two lightpaths from router 3 to router 4. Lightpaths 1→3 follow the same physical route 1→2→3, and are bundled together using the LB technique. Thus, they are seen by the IP layer as an aggregated link 1→3 of capacity 120 Gbps (Fig. 1(b)). Similarly, the two lightpaths 3→4 are bundled and seen as a virtual link of 80 Gbps. This configuration is used to carry traffic from node 1 to node 3 and 4, and from node 3 to node 4. In router 1, the AS technique results in an even distribution of the IP packets switched to the three line cards associated to the three outgoing lightpaths bundled 1→3. The same happens in router 3, with the IP traffic targeted to node 4, that is balanced between the two line cards corresponding to the two outgoing lightpaths 3→4 bundled.

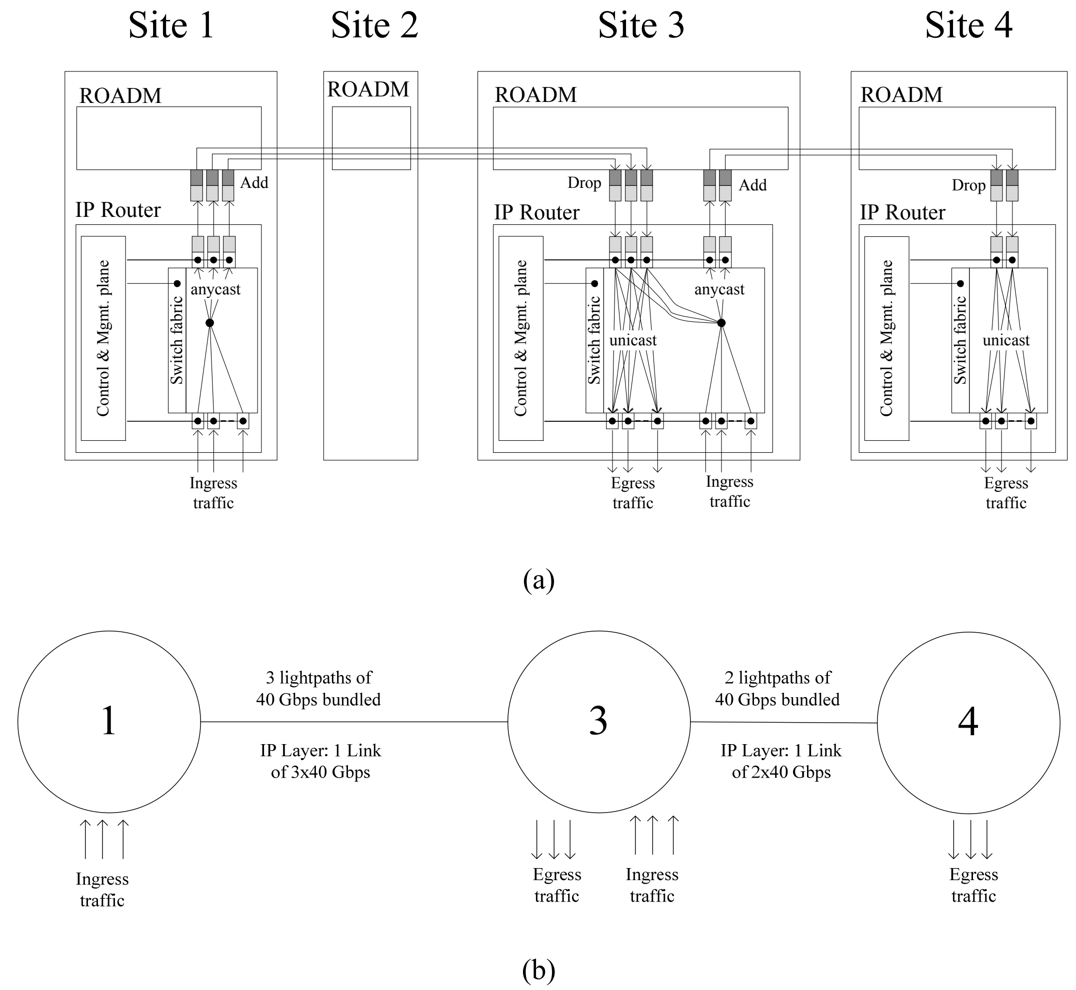

Fig. 1. Example of LB+AS techniques. (a) Equipment scheme. (b) The perception of the IP layer.

The critical point in the LB+AS paradigm is that the multilayer context brings an unprecedented opportunity to massively exploit the benefits of traffic balancing. The reason is that in multilayer networks, the topology observed by the IP layer is *not* the set of fibers in the network (a static infrastructure, that network engineers cannot promptly modify), but the *virtual* topology of lightpaths designed. Consequently, it is possible to plan this virtual topology of lightpaths appropriately, so that the benefits of lightpath bundling are maximized. Moreover, interestingly, the application of LB and AS techniques requires seamless changes in the electronic equipment, and no changes in the optical infrastructure. In particular, LB can be at present implemented in existing Network Management Systems, and in networks based on a GMPLS control plane, by making use of the GMPLS link bundling functionality as described in RFCs 4201 and 4202. However, in practice, the LB configuration is not applied in today's optical networks. In this respect, this work suggests that a new incentive exists for it, if AS is also configured in the data plane of the IP routers.

In this paper, we explore the utilization versus delay trade-off benefits of the LB+AS paradigm. Our case of study is a multilayer network built on top of the well-known Internet2 reference topology [3] fed by self-similar IP traffic. We use the Lightpath Bundling Aware Multilayer Planning Algorithm (LBA-MPA) presented in [2] for planning the network so that LB+AS benefits are exploited. We show that, for the same network infrastructure and traffic, applying the LB+AS techniques can reduce end-to-end network delay in more than one order of magnitude. Correspondingly, if we fix an average end-to-end accumulated queuing delay target to design our network, and if LB+AS techniques are applied, then: (i) a ~30% reduction in the cost and energy consumption of the optical layer can be obtained for the same carried traffic, or (ii) up to a ~50% increase in the carried traffic is possible for the same network cost. These results support the application of the LB+AS paradigm to existing multilayer networks.

The remainder of this article is organized as follows: First, we dedicate a section to show the core idea that propels the performance gain of the LB+AS concept. Then, we describe the Internet2 case study and their results. Afterwards, a section elaborates on the merits and open challenges related to the application of LB+AS scheme to existing networks. Finally, the article conclusions are summarized.

Where the performance gains of LB+AS come from?

The AS extra degree of freedom to select the output lightpath of an arriving packet can be used to implement a very fine per-packet traffic balancing among the bundled lightpaths, resulting in traffic smoothing. This is the core effect powering the performance benefits of the AS with respect to conventional unicast switching. To illustrate this, we focus on a simple scenario consisting of two IP routers connected by a set of b_D lightpaths bundled together (b_D stands for *bundling degree*). The traffic traversing the bundle is given by the aggregation of 100 ingress ON-OFF sources operating at a 10 Gbps bit rate. Each source generates during its ON period IP packets with a length given by an Internet-mix trimodal distribution (40, 552 and 1500 bytes, with probabilities 58%, 33% and 9%) [4]. ON periods follow a Pareto distribution of average duration 1 MB, and a shape parameter dimensioned according to [5], so that the superposition of the sources produces a self-similar traffic of Hurst parameter H=0.6. The duration of the OFF period has an exponential distribution, with average duration adjusted to fit the appropriate load of the source. The IP routers have a conventional VOQ (Virtual Output Queueing) switch architecture, with an internal speed-up and a scheduler which permits the emulation of output buffering behavior. Anycast switching is implemented as a JSQ (Join the Shortest Queue) policy, as described in [2]: for each packet, the ingress line card selects the output lightpath in the bundle, corresponding to the virtual output queue with less bytes pending to be transmitted. Note that this decision can be taken independently for each line card.

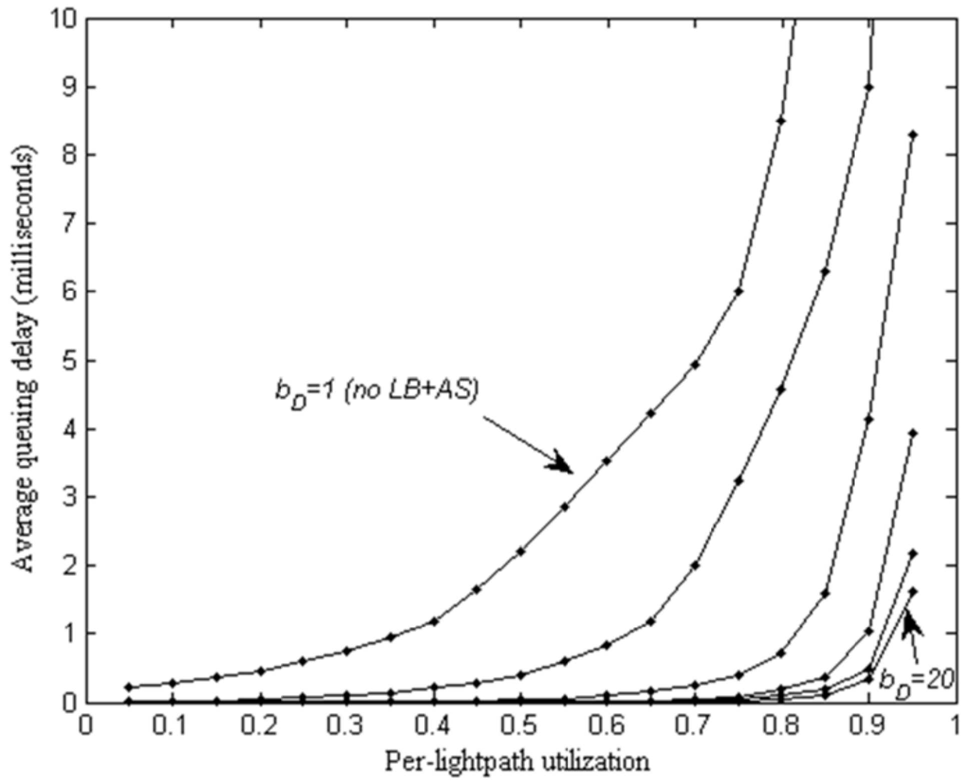

Fig. 2. Average queuing delay performance in a 2-node network, b_D={1, 2, 5, 10, 15, 20}

Fig. 2 illustrates the average packet delay performance (in milliseconds) with respect to the per-lightpath utilization in the system, for different bundling degrees b_D={1, 2, 5, 10, 15, 20}. The case b_D=1 is equivalent to not applying LB+AS to a set of lightpaths. Clearly, as the bundle degree grows the average queuing delay for the same amount of traffic carried drops. For example, if we have 10 lightpaths which can be bundled together, and fix a maximum average queuing delay target of 1 ms, then: (i) the maximum load per lightpath can be 35% if the lightpaths are not bundled, (ii) or 90% if the LB+AS techniques are applied. Alternatively, if we have an utilization of 60% in the link, the average queuing delay drops from 3 ms to a few microseconds if LB+AS techniques are applied.

The per-packet granularity of the balance makes AS different to other traffic balancing alternatives (e.g. Ethernet 802.1AX standard), which define a per-flow granularity balancing. Balancing at a per-flow granularity means that the flow associated with the packet must be identified at wire-speed. This is carried out according to an implementation-dependent (and often proprietary) definition of what a "flow" is *for traffic balancing purposes*. For instance, a switch implementation can generate an 8-bits flow identifier for each packet received, hashing its input and output IP addresses. Then, the packets with the same flow identifier are switched to the same output link using conventional unicast switching. The mapping between flow identifiers and output links is designed to make a uniform balance of the traffic load among the links in a bundle. This is a non-trivial task that requires monitoring the load associated to each flow identifier, and periodically update the flow-to-link map to avoid large deviations in the balance. In comparison to this scheme, AS provides a finer and more perfect balance at a packet granularity, without the need of either wire-speed packet flow identification or flow load monitoring. However, since the AS balance does not rely on flow identification, it can occur that two IP packets of the same multimedia or TCP connection traverse different lightpaths within a bundle, which can cause IP packet out-of-order delivery. Still, the order between IP packets within a connection traversing several bundles can be kept, if each router maintains the order among the packets arriving from the same lightpath bundle targeted to the same output bundle. This packet ordering is preserved if (i) the switching equipment is able to emulate output buffering and (ii) a JSQ policy is applied.

The Internet2 case of study

In this case study we base on the Internet2 network topology and the reference traffic matrix M presented in [3] for this network. We consider 10 Gbps transmitters in all the nodes and two possible wavelength grids of number of wavelengths W={40,80}. In each case, we normalize the traffic matrix following the same procedure described in [2]. As a result, the traffic matrix $M(W)$ associated to network load factor ρ=100%, is the largest matrix, of the form $a \cdot M$, $a>0$, that can be supported by a multilayer network where lightpaths can be loaded at a 100% of average utilization. As shown in [2], this maximum traffic matrix can be calculated optimally solving a linear formulation, and no larger traffic matrix $a'M$, $a'>a$ can be supported without oversubscribing a link. Finally, the traffic matrices for other load factors ρ={20%, 25%, ..., 95%} are calculated as the ρ fraction of the matrix $M(W)$.

The cost of the network is estimated as the number of lightpaths required to carry the traffic. This is a measure of both the CAPEX costs and the energy consumption OPEX costs in the optical layer. On the CAPEX side, the number of lightpaths determines the number of transponders, commonly assumed as a good CAPEX indicator. Regarding the energy efficiency, the number of lightpaths is a precise indicator of the energy consumption at the optical layer, since the energy consumption of a lightpath is, in practice, independent of its utilization. In other words, once the lightpaths are established, they roughly consume the same energy whatever the amount of traffic they actually transmit [6].

Given a number of wavelengths W and a load factor ρ, we use the LBA-MPA [2] algorithm to plan the multilayer network. The LBA-MPA algorithm consists of four consecutive steps. Steps 1 to 3 heuristically search for a tentative VT and IP routing (Step 1), lightpath routing (Step 2) and wavelength assignment (Step 3) for the network, with the constraint that the maximum lightpath utilization is limited to u ($0<u<1$), an internal algorithm parameter. Wavelength conversions are not considered. If a valid solution is found, the algorithm progresses to Step 4. In this step, the algorithm bundles together all the lightpaths with common end nodes and a common route, so that they appear to the IP layer as a single link of aggregated capacity. Then, the IP routing in the network of lightpaths calculated in Step 1 is replaced by a new IP routing in the network of lightpath bundles. The new IP routing is computed with the objective of minimizing the average network delay. This objective intrinsically produces network plans where the bundles of higher degree have a higher utilization: an intended target to exploit the traffic balancing more efficiently.

The LBA-MPA algorithm can be executed in two modes: assuming an underlying LB+AS network, or not. The second mode is a sub-type of the first one, in which the Step 4 of the algorithm is eliminated. Then, lightpaths which could be bundled together are not, or equivalently, all the lightpaths are in a bundle of degree one. For every (W, ρ, LB+AS yes/no) test we execute the LBA-MPA algorithm 16 times, one for each value of internal parameter u={0.20, 0.25,..., 0.95}. For a given u parameter, it may happen that the LBA-MPA algorithm cannot find a valid solution. These cases are rejected. For those cases when a valid solution is found, we conduct an event-driven simulation of the resulting multilayer network. The target of these simulations is evaluating the average end-to-end queuing delay of the IP packets in the designed network. The simulations are built in the OMNeT++ framework [7]. Scripts are used to automate the generation of the OMNeT++ network models from the solutions computed by the LBA-MPA algorithm. In the simulations, each node is assumed to implement a VOQ arbitration able to emulate output buffering behavior, with infinite buffering per line card, and a JSQ scheduling policy as described in the previous section. The statistical properties of the IP traffic generated between each pair of routers are the same as the ones described for the 2-node network example, and thus have self-similar properties.

Five samples are obtained in each simulation point, and confidence intervals are calculated by the *t*-Student method, validating the accuracy of the results. Each sample is an independent simulation that stops when at least 10^8 packets are processed by each node. The simulations were executed in the Ben-Arabi Supercomputing facility [8], located in Murcia (Spain), using more than 1,500 hours of processor time.

Fig. 3 illustrates the benefits in terms of delay performance improvement brought by the LB+AS paradigm. As a measure of the delay performance we use the average end-to-end queuing delay of the packets in the network. This means averaging the accumulated time spent by the packets in the buffers of the traversed IP routers. Note that this is both an indication of the delay (but omitting the transmission and propagation delays), and also an indicator of the amount of buffering required in the routers. For each W={40,80} and each network load factor ρ, we collect all the possible network instances ranging all the tested values of internal parameter u. Then, we dimension the network using the particular u factor (i) for which the network cost is minimum, and (ii) but for which, when LB+AS technique is *not* applied, the average end-to-end queuing delay is below 3 ms. In other words, we take the cheapest network that is able to carry the traffic without applying the LB+AS concepts, satisfying the delay performance target set. Then, we compare the network queuing delay with that of an equivalent planning, with the same cost and traffic, but which applies the LB+AS paradigm.

The main conclusion that can be brought from Fig. 3 is that LB+AS techniques significantly improve the delay performance in the network: reductions obtained in average end-to-end delay are, in general, between one and two orders of magnitude. In addition, it is interesting to see that the higher the traffic is (that is, higher load factor, or higher number of wavelengths), the more intense the delay performance improvement becomes. This is because, when LB+AS is not applied, the number of lightpaths in the network for a given delay performance target grows (approximately) linearly with the traffic volume. In other words, more traffic means network plans with more lightpaths and nearly the same per-lightpath utilization. Then, applying LB+AS to networks carrying higher traffic volumes results in opportunities to bundle together more lightpaths and producing bundles of higher degree. As illustrated by Fig. 2, spreading the IP traffic in bundles of higher degree, while keeping the per-lightpath utilization, betters the packet delay performance.

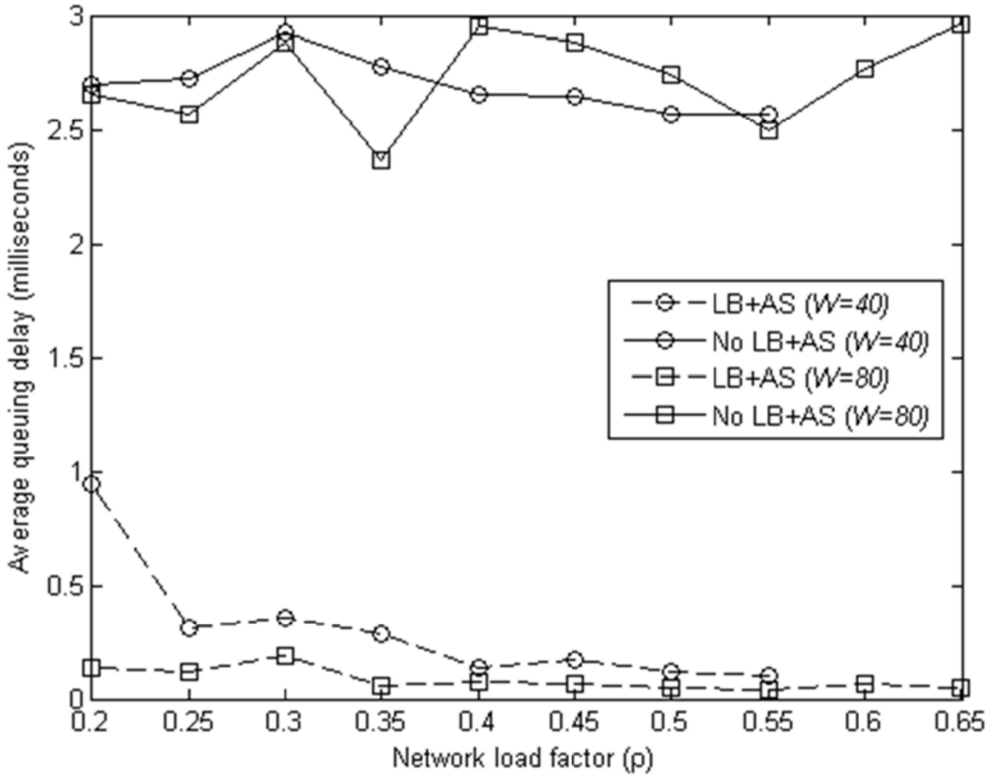

Fig. 3 Average end-to-end queuing delay performance vs carried traffic comparison

Finally, note that horizontal axis in Fig. 3 ranges the load factors up to 55% if *W*=40, and up to 65% if *W*=80. These are the maximum load factors for which it was possible to have no-LB+AS solutions within the network delay target. However, when LB+AS schemes are implemented, it is possible to carry more traffic while still fulfilling the delay performance target set. Put differently, LB+AS can be also seen as a scheme for enhancing the economic efficiency in the network, translating the delay performance benefits into (i) cost reductions or (ii) revenue increases (carrying more traffic). Fig. 4 helps us to illustrate this, by looking at the utilization versus delay trade-off from other point of view. We set again a maximum end-to-end average queuing delay in the network of 3 ms. Then, for each load factor, we select the network infrastructure with lowest cost, which is able to satisfy the delay target in two different cases, (i) if the LB+AS techniques are applied, and (ii) if the LB+AS techniques are not applied.

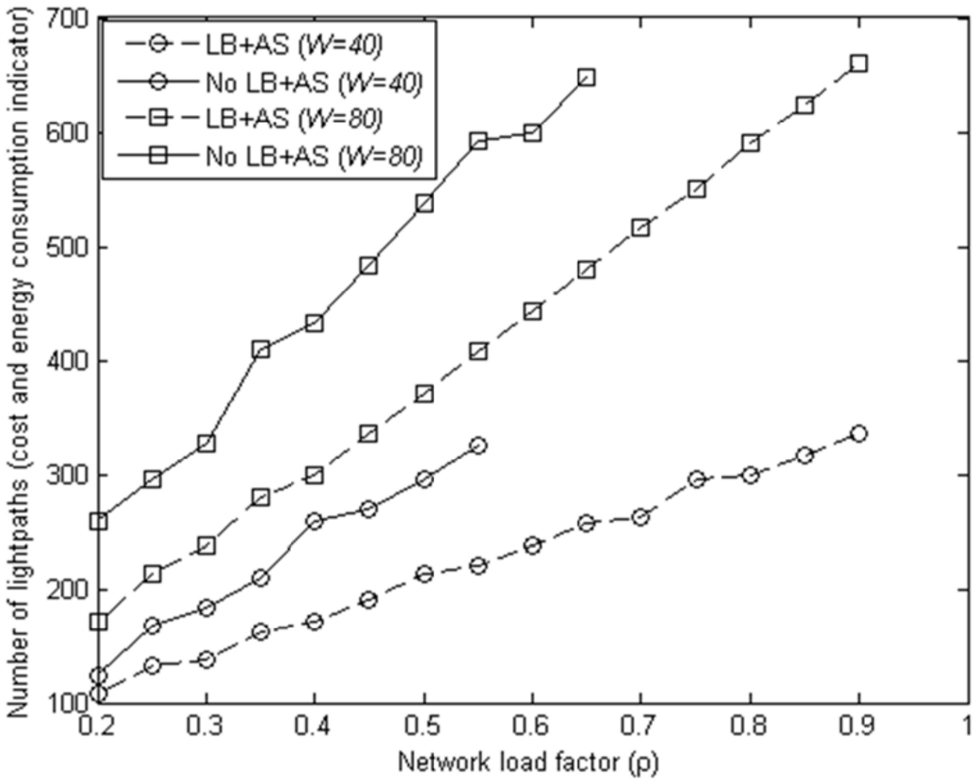

Fig. 4 Cost vs carried traffic comparison

Observing Fig. 4, we note that, for a given number of wavelengths *W*, the cost-vs-traffic lines corresponding to LB+AS are (i) “below” and (ii) “at the right-hand side”, with respect to the lines corresponding to no-LB+AS cases. Being “below” means that applying LB+AS may permit using less lightpaths in the network, to carry the same traffic complying the delay target. Recall that less lightpaths means less CAPEX/OPEX costs at the optical layer, including lower energy consumption. The percentage of reductions in network cost is in most cases between 20% and 30%, being, in general, slightly better for higher loads. On the other side, being “at the right-hand side” means that if we fix the network cost budget, applying LB+AS techniques to the network permits carrying more traffic still complying the delay target: between 20% and 50% more. Likewise, LB+AS combination permits loading the network at factors of 90% fulfilling the delay target, improving the 55% and 65% maximum load factor limits mentioned above for *W*=40 and *W*=80 respectively.

Fig. 3 and Fig. 4 give a hint of the good scalability properties of the LB+AS paradigm, which works better for networks with a higher number of wavelengths per fiber. This is because the benefits of the LB+AS concept strongly depend on the opportunities to bundle together more and more lightpaths. As a result, we can infer that the LB+AS paradigm is inherently scalable: carrying more traffic in the network works in favor of the LB+AS concept, since it means more opportunities to create bundles of higher degree. Table I helps us to perceive this beneficial effect, and also shows some actual numbers of bundle degrees obtained in the Internet2 case. It reports the average degrees of the bundles planned by the LBA-MPA algorithm, for the network infrastructures used in Fig. 4. As can be seen, the average bundle degree grows in an approximately linear form with the traffic carried, reaching 4.5 and 9 lightpaths for a 90% load factor in *W*=40 and *W*=80 cases respectively.

Table I. Average bundling degree. Internet 2 case.

	ρ	0.2	0.25	0.3	0.35	0.4	0.45	0.5	0.55	0.6	0.65	0.7	0.75	0.8	0.85	0.9
W	40	1.5	1.8	1.8	2.1	2.4	2.7	2.8	3	3	3.6	3.6	3.9	4.2	4.3	4.5
	80	2.4	2.8	3.3	3.9	4.2	4.7	5.2	5.7	6.2	6.7	7.2	7.6	8.2	8.7	9.0

The merits and challenges of LB+AS

The results in this paper support the interest in the LB+AS paradigm, and encourage further investigations in related topics. In this section we summarize and review some of the merits of the proposed scheme, together with open challenges:

- **Network Cost/QoS/Throughput triangle**: The principal effect shown in this paper related to applying LB+AS techniques to the network, is that the triangle binding (i) network cost, (ii) network delay performance (QoS), and (iii) amount of carried traffic (throughput), is positively biased. Specifically this means that more traffic can be carried in the same network keeping the same delay performance targets by applying LB+AS. Furthermore, for a given network infrastructure, the same traffic can be carried with better delay performances, etc. In this project, we used the average end-to-end queuing delay metric, and arbitrarily selected a threshold of 3 ms, with the intention of illustrating the potential benefits of LB+AS paradigm. The results promote the investigation on how LB+AS can affect other QoS metrics in the network and other network scenarios.
- **Network scalability**: As illustrated in this paper, the benefits of the LB+AS paradigm improve for higher network traffic, since a higher traffic volume means more lightpaths to set-up and more chances to bundle them together. Consequently, the LB+AS scheme becomes a future-proof option, suited for handling foreseen traffic growths. Furthermore, the incoming multiplexing systems aggregating 160 wavelengths in a fiber (e.g. CISCO ONS 15808) would work in favor of the network efficiency.
- **Gradual upgrade of the transmission equipment**: Increases in the traffic volumes require from the network operators periodic upgrades in the lightpath capacities. However, the increases in the lightpath capacities depend on the technology, and follow coarse steps of x4 and x10. For instance, capacity upgrades available for 10 Gbps lightpaths are now 40 Gbps, and 100 Gbps, with no mid-way alternatives. LB+AS paradigm permits to the network architect a more fine grained upgrade in the optical capacity installed connecting two nodes, with little impact on the IP layer: e.g. it is possible to bundle together two or three 10 Gbps lightpaths, which for the IP layer act as a lightpath of capacity 20-30 Gbps.
- **Suitable for lightpath-on-demand services and traffic dynamicity:** LB+AS paradigm enforces an operation of the network where the IP layer (in general operated by Internet Service Providers -ISPs) can maintain stability (e.g. stability in the IP routes), and adapt to traffic variations by on-demand hiring new lightpaths to the optical network carrier. Technologically, a lightpath-on-demand market is now enabled by new and more agile directionless-colorless ROADMs installed by network carriers. What LB+AS adds to the picture, from the ISP point of view, is the possibility to (i) bundle together the lightpaths hired between two nodes, and (ii) adjust the bundle capacity to

the traffic volumes carried by dynamically setting up/tearing down lightpaths. Then, the LB+AS-aware packet switches would seamlessly balance the traffic among the lightpaths in the bundle at every moment. As a result, the adjustments in the bundle sizes are seen by the IP layer as a capacity update in a virtual link (lightpath bundle), and not as new links that should be advertised to other nodes, and added to the IP routing tables. By doing so, the signaling bottleneck which threats the dynamic optical networks of the future would be greatly mitigated. The preliminary work in [9] suggests that this process would permit the network to efficiently adapt to traffic variations, while strictly maintaining IP routing stability. This and other implications should be studied, with the aim of investigating the LB+AS paradigms as an enabling technique for integrating the future lightpath-on-demand services into the ISP normal operation.

- **Implementation of AS techniques:** AS is a novel technique proposed in [2]. The introduction of AS functionality in commercial routers opens a research line: how existing (or new) switching architectures and schedulers can be tuned to exploit the AS extra degree of freedom in choosing the packet output port in the fabric? In this aspect, the buffering requirements should be re-evaluated, since the observed reduction in the queuing delay brought by the AS technique, is coupled to a reduction in the buffering requirements in the line cards. Finally, enforcing packet order (e.g. by implementing JSQ policies) or permitting a limited packet out-of-sequence, are options that should be studied under a new light: in high-speed multilayer networks, links are shared by a huge number of TCP and multimedia connections. This is a relevant difference, since the number of packets separating two packets of the same connection can be high. Then, small packet misorderings created by the AS could be accepted if they do not affect packets of the same TCP or multimedia connection, and thus do not affect the performance of the upper layers. These aspects are a current line of work in our group.

Conclusions

Lightpath bundling and anycast switching are promising candidates for defining the reference architecture for the backbone and metro network of the Future Internet. Applying the LB+AS paradigm requires no changes in the optical infrastructure of the network. At the electronic layer, LB can be already implemented in existing GMPLS control planes and Network Management Systems, while new mechanisms to implement AS in existing packet switches are being investigated. The preliminary results in this paper suggest notable benefits, and encourage the investigations in the open research lines related to the application of LB+AS paradigm to existing multilayer optical networks.

References

[1] Rajiv Ramaswami, Kumar N. Sivarajan, “Optical networks: a practical perspective”, Morgan Kaufmann Publishers 1998.

[2] Pablo Pavon-Marino, “Lightpath bundling and anycast switching: a good team for multilayer optical networks", 15th Conference on Optical Network Design and Modelling (ONDM 2011), Bologna (Italy), February 2011.

[3] Internet 2 Global Research Network Operations Center, Web Site, [Online], Available: http://www.abilene.iu.edu/.

[4] Agilent technology; Packet over Sonet/SDH (POS) functional testing, product note, 2000.

[5] Ronald G. Addie, Timothy D. Neame, Moshe Zukerman, "Modeling Superposition of Many Sources Generating Self Similar Traffic", in *Proceedings of the 1999 IEEE International Conference on Communicaiones (ICC '99)*, Vancouver (Canada), June 1999, pp. 387-391.

[6] Joseph Chabarek, Joel Sommers, Paul Badford, Cristian Estan, David Tsiang, Steve Wright, "Power Awareness in Network Design and Routing", in *Proceedings of the 27th Conference on Computer Communications (INFOCOM 2008)*, Phoenix (USA), April 2008, pp. 457-465.

[7] András Varga, "The OMNeT++ Discrete Event Simulation System", in *Proceedings of the European Simulation Multiconference (ESM '2001)*, Prague (Czech Republic), June 2001, pp. 319-324.

[8] Ben-Arabi Supercomputing facility, [Online] http://www.cesmu.es.

[9] Pablo Pavon-Marino, Nina Skorin-Kapov, "Stable Routing with Virtual Topology Capacity Adjustment: A novel paradigm for operating optical networks", *Journal of Optical Communications and Networking*, vol. 3, no. 12, pp. 902-911, December 2011.