Transport mechanisms for mobility support in optical slot switching-based next-generation mobile backhaul networks

N. Benzaoui, Y. Pointurier, B. Uscumlic, T. Bonald, Q. Wei, and S. Bigo,

(Invited Paper)

***Abstract*—We propose and evaluate several transport mechanisms to support User Element mobility at the optical layer in a next-generation (e.g., LTE-A) mobile backhaul network implemented using our previously proposed optical slot switching technology. We show with simulations that, with smart introduction of additional electronic processing for mobility traffic only (traffic of user elements that have moved from one cell to another), very strict latency constraints (sub-ms) can be met even for this mobility traffic.**

***Index Terms*—Optical packet switching; network performance; dimensioning, CoS, LTE-A mobile backhaul.**

I. Introduction

With the evolution of the optical transport technologies, energy-efficient solutions for the transport in the radio access part of cellular backhaul networks have become available. In this paper, we consider the use of a fine granularity optical transport solution: optical slot switching (OSS), an energy-efficient [1] (thanks to optical transparency for transit traffic), low-latency switching technology [2], where user packets are encapsulated into fixed duration containers called slots, to interconnect base stations (eNodeBs or eNBs) in an LTE-A mobile backhaul network. Here, optical slot switching is applied at the metro level for backhaul networks with a few (and up to a few dozens) of nodes (eNBs): Fig.1 shows one example of LTE-A mobile backhaul network with eNBs interconnected in a ring structure. We focus on Packet Optical Add/Drop Multiplexer (POADM; see top of Fig.1), an implementation of OSS rings. By shifting the support of UE (User Equipment) mobility (from one cell to another one) to the optical layer, OSS has the potential to dramatically reduce the total handover (HO) duration by 50-70% as shown in [3], assuming (conservatively) a queuing delay of 2 ms in the OSS backhaul. We show in this paper that this latency can in fact be much further reduced to well below 1 ms for real-time traffic. To do so, we propose and evaluate 5 optical transport mechanisms that will effectively enable the mobile OSS backhaul to fulfill the aforementioned handover duration reduction.

N. Benzaoui, Y. Pointurier, B. Uscumlic and S. Bigo are with Bell Labs, Alcatel-Lucent, Nozay, France.

T. Bonald is with Telecom ParisTech and LINCS, Paris, France.

Q. Wei Formally with DOCOMO Euro-Labs, Munich, Germany.

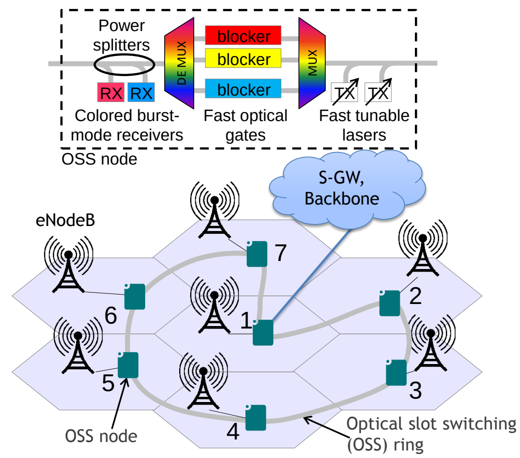

Fig. 1. POADM-based mobile backhaul ring

We first propose 2 electronic forwarding mechanisms, the first one based on "slot drop and forward" transmission, which optimizes latency, while the second is based on intermediate re-encapsulation, to optimize capacity efficiency. To find an optimal trade-off we propose a third, adaptive mechanism. In order to further reduce latency for forwarded traffic we propose a fourth, transparent forwarding mechanism based on all-optical "slot drop and continue" transmission. The mechanisms are presented in Section 2. We evaluate and compare the performance of the 5 proposed mechanisms in Section 3. In Section 4 we propose dimensioning algorithms that minimize the number of network transponders (which drives the network cost) for all 5 mechanisms and compare them in terms of cost/dimensioning. We draw brief conclusions in Section 5.

II. Optical Slot formation and queue management

Before describing the proposed traffic forward mechanisms, we fist explain the basic optical slot formation and how queues are managed.

At an OSS (eNodeB) node, packets arriving from clients (e.g. mobile terminals) are placed in a temporary packet queue. Once enough packets to fill a slot have arrived in the temporary packet queue, this queue is emptied and the

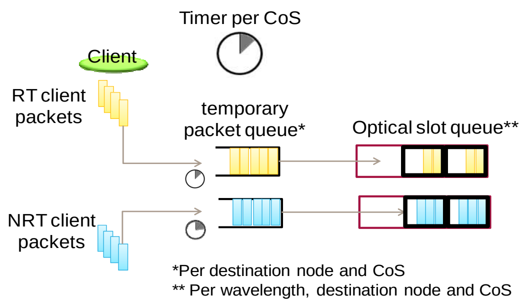

Fig. 2. Optical Slot formation.

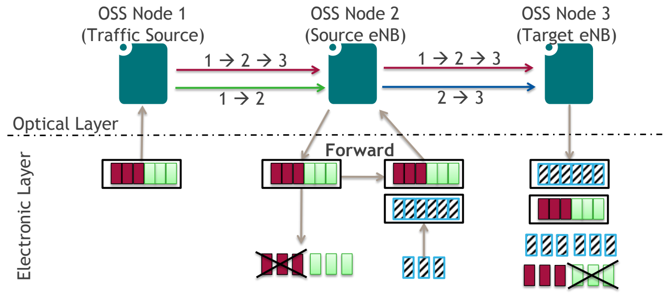

Fig. 3. N1: Electronic forwarding without re-encapsulation

optical slot is formed and placed in the optical slot queue, waiting for its insertion on the channel. To cap the optical slot formation time, we use a timer, which is triggered by the arrival of the first client packet, and which expiration causes the completion of the formation of the optical slot.

In this work we consider that a typical traffic in an LTE-A mobile backhaul is divided into two classes of service: real-time (RT) and non-real-time (NRT). RT traffic is packet loss and latency-sensitive while NRT (or best effort) is more tolerant to delay and packet losses. The classes of service are managed by identifing the temporary packet queues by destination (eNodeB) node and CoS, such that each optical slot contains packets with the same CoS. The utilization of per-CoS queues enables sending RT traffic before NRT traffic. At the slot assembly step we apply one timer per CoS: a short timer for RT traffic (to respect latency restriction) and a larger timer for the NRT traffic (since it is less sensitive to latency).

III. Traffic forward mechanisms

A typical LTE-A handover procedure is depicted in Fig. 3, where a user receives traffic from Node 1 and is moving from a cell (Node 2) to another cell (Node 3), both served by same service gateway [4]. During handover procedure, data is first received by Node 2, which forwards it to Node 3. When the handover is complete, Node 1 sends data directly to Node 3. In an OSS-based LTE-A backhaul, user packets are encapsulated in slots according to their destination eNodeB. Therefore, slots sent from Node 1 to Node 2 contain user packets with destination Node 2 but also packets that have to be forwarded from Node 2 to Node 3. Information about the intermediate and final destinations of client frames is carried by the out-of-band control channel. Based on this information and without further user packet processing, Node 2 knows which slots should be forwarded. Latency during handover is thus essentially impacted by the forwarding mechanism in the optical layer [1].

[1]Thanks to the ring topology, the forwarding mechanisms proposed here also apply to the multicast types of traffic such as IPTV or CoMP [5]. CoMP is a cooperation technique that requires cell cooperation to increase user capacity at cell edges, which should experience less than 1 ms end-to-end latency, i.e., less than 500 µs within one access ring.

The forwarding mechanisms proposed in this paper presents a trade-off between latency and channel utilization for the traffic that has to be forwarded from the source eNodeB to the target eNodeB. In the next section we propose to evaluate different forwarding mechanisms, first using a simple than a realistic scenario, to designate which one brings the best trade-off between latency and dimensionning.

A. CoS-oblivious electronic forwarding mechanisms: N1, N2, M3

As a reference, we first present two naive forwarding mechanisms to support mobility directly at the optical layer

1) Naive forwarding mechanisms: N1, N2:

a) Naive Mechanism N1:Electronic forwarding without intermediate re-encapsulation:

As illustrated in Fig. 3, if an eNodeB (Node 2) receives a slot containing user packets that should be forwarded to another eNodeB (Node 3), the slot is 1) received by Node 2; and 2) forwarded without any modification to Node 3, after 2 optoelectronic conversions. The slot is de-capsulated at Node 3, which discards irrelevant packets whose destination is not Node 3. The advantage of this mechanism is that small channel re-insertion latency is added by Node 2 to mobility traffic (traffic that has to be forwarded from Node 2 to Node 3). This is done by using a higher priority for forwarded traffic at Node 2. Indeed, even if slots are queued at Node 2, the latency is limited thanks to the use of priority scheduling for forwarded traffic. However, since forwarded slots contain user packets which destination is not Node 3, some offered channel capacity between Nodes 2 and 3 is wasted.

b) Naive Mechanism N2: Electronic forwarding with intermediate re-encapsulation:

As illustrated in Fig. 4, and contrary to Mechanism N1, after slots de-capsulation at Node 2, user packets that have to be forwarded to Node 3 are re-encapsulated with user packets sent from Node 2 to Node 3. The advantage of this mechanism is that, since slots sent from Node 2 to the Node 3 contain only user packets with destination Node 3, no offered channel capacity is wasted. The disadvantage is that, due to the re-encapsulation at Node 2, forwarded

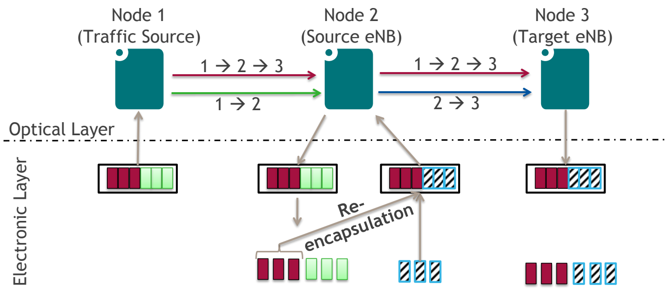

Fig. 4. N2: Electronic forwarding with re-encapsulation

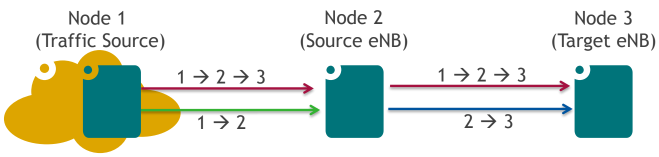

Fig. 5. 3-node Scenario

traffic may observe substantial additional latency, which may adversely affect real-time (RT) traffic.

2) First evaluation of Naive traffic forwarding mechanisms on a simple scenario: We first provide the following simple scenario to evaluate the compromise between latency reduction and throughput saving of the above forwarding mechanisms. As described in Fig. 5 we consider a 3-node network. Node 1 (traffic sender) sends traffic with intensity “A” to Node 2 with destination of Node 2 clients and traffic with intensity “A” that should be forwarded to Node 3. Node 2 sends traffic with intensity “A” to Node 3. There is a single wavelength in the network that is shared by all nodes. It can easily be seen that A =1/2 saturates the link between Node 2 and Node 3. In the following, we normalize the network load to this saturation. The slot duration is D=8 µs and the channel rate is 10 Gb/s. Fig. 6 shows the latency (queuing in temporary and optical slot queues) with respect to the network load for a) the forwarded traffic and b) the inserted traffic at source eNodeB, in dashed curves using electronic forwarding without re-encapsulation and in solid curves electronic forwarding with re-encapsulation.

As explained before, the electronic forwarding without re-encapsulation (N1) keeps the latency low for the forwarded traffic but causes offered channel capacity waste; this is visible on the inserted traffic that can reach only 70% of the load (instead of 100%). In contrary, using the electronic forwarding with re-encapsulation (N2) all the offered channel capacity is used (both forwarded and inserted flows reach 100% of the achievable load) at the expense of a higher latency experienced by the forwarded traffic which get higher using this mechanism than the first one.

Overall, N1 and N2 trade off network capacity (optimized by N2) for end-to-end latency (optimized by N1.) For low-loads the electronic forwarding without re-encapsulation mechanism is suitable, since even if a part

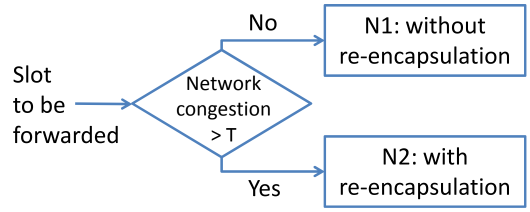

Fig. 7. M3: Adaptive forwarding mechanism.

of the bandwidth is lost it does not matter since there is no need to this resource. For high loads electronic forwarding with re-encapsulation mechanism is suitable since the bandwidth is a precious resource and should be saved even if it is to the detriment of added latency to the forwarded traffic.

3) Adaptive forwarding mechanism: M3: We propose a third, simple and smarter, mechanism M3 to remove the aforementioned latency-capacity trade-off. This mechanism combines N1 and N2 and selects which forwarding mechanism to use at every slot based on the network load. Mechanism M3 is described in Fig. 7. The re-encapsulation process (N2) is used at Node 2 only when the channel load is high (and network is congested) and capacity should not be wasted. When the channel load is low, N1 is used and capacity (which is then abundant) is leveraged to reduce latency. Because the load is likely to vary in a real network scenario, the re-encapsulation decision is taken dynamically, slot by slot. M3 is distributed: each node independently assesses network congestion by counting the number of slots locally queued for insertion; a pre-defined number of queued slots (threshold T) indicates network congestion.

4) First evaluation of adaptive forwarding mechanism on simple scenario: Using the same scenario described in Section III-A2 we evaluate the new adaptive forwarding mechanism and compare it to previous mechanisms N1 and N2.

On Fig. 8 we keep the results of the evaluation of the electronic forwarding with and without re-encapsulation and compare them with the queuing delay obtained using the adaptive forwarding mechanism. Fig. 8 shows that using the adaptive forwarding mechanism we take advantage of each electronic mechanism depending on the network load (seen locally by a node): we take advantage of the electronic forwarding without re-encapsulation mechanism at low-loads and the one with re-encapsulation at high-loads. Hence, with the adaptive forwarding mechanism we insure low latency for the forwarded traffic at low loads and save offered channel capacity at high loads.

B. CoS-aware and adaptive electronic forwarding mechanisms: M4

M4 is depicted in Fig. 9. To further cap latency for very sensitive mobile traffic such as CoMP, we propose to enhance the adaptive mechanism M3 by introducing two levels of network congestion detection, as shown in

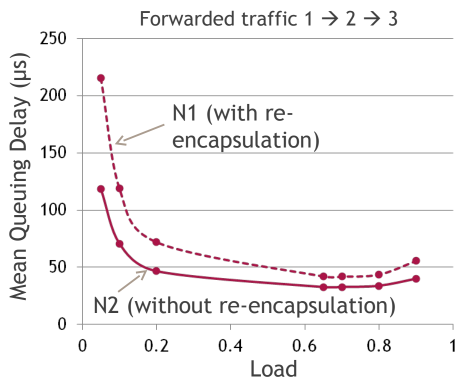

(a) Forwarded traffic (traffic sender → Target eNodeB)

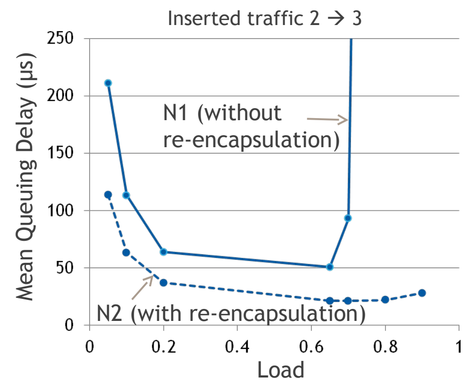

(b) Forwarded traffic (source eNodeB → Target eNodeB)

Fig. 6. Comparison between electronic forwarding with and without re-encapsulation

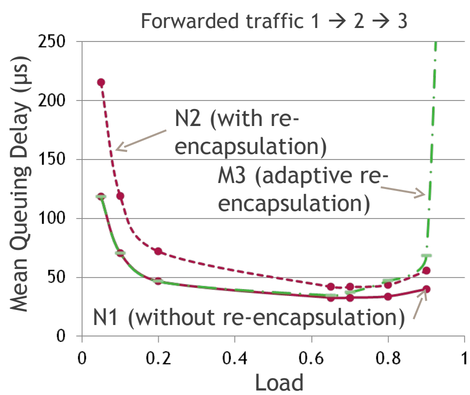

(a) Forwarded traffic

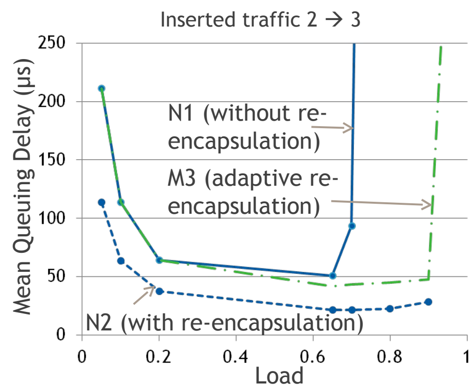

(b) Inserted traffic

Fig. 8. Comparison between electronic forwarding with and without re-encapsulation and adaptive forwarding mechanism

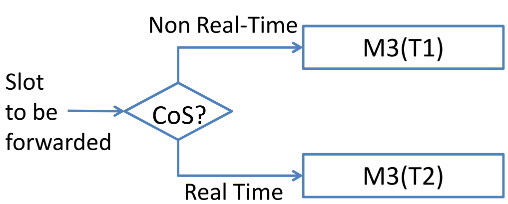

Fig. 9. M4: Cos-aware adaptive forwarding mechanism.

Fig. 9. A first congestion threshold T1 applies to NRT slots, and a second threshold T2>T1 applies to RT slots. This mechanism keeps the latency for forwarded RT traffic as low as possible by saving channel capacity by re-encapsulating NRT forwarded traffic when the network is congested. Comparably to M3, M4 is also distributed.

C. *Transparent forwarding mechanisms: M5*

In order to further reduce latency for forwarded traffic we propose a transparent forwarding mechanism based on all-optical transmission. As illustrated in Fig. 10, a slot that contains user packets for Node 2 and Node 3 is dropped at Node 2 and continues to Node 3 without going through electronic processing at Node 2. The advantage

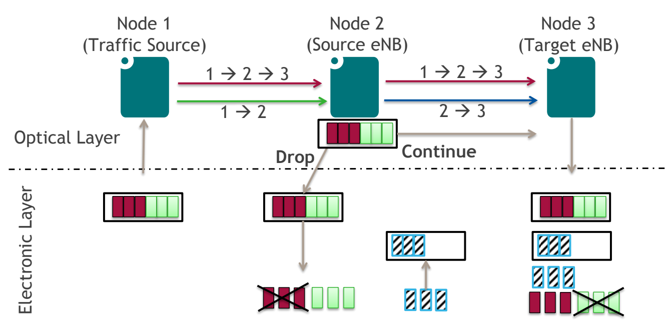

Fig. 10. M5: Transparent forwarding mechanism

of this mechanism is that it adds no latency for the forwarded traffic since it is treated as transparent traffic. Note that, at a given node, transit traffic has the priority over inserted traffic. However, like N1, M4 wastes offered channel capacity, in addition, M4 makes the dimensioning more expensive, as will be seen in Section V.

IV. Performance evaluation

In the following we first evaluate the performance of the electronic forwarding mechanisms on a realistic scenario

TABLE II
HANDOVER TRAFFIC DISTRIBUTION

Traffic sender	Source eNodeB	Target eNodeB	Data rate (Gb/s)
Hub	eNodeB	eNodeB	0.1
Hub	eNodeB	Hub	0.3
Hub	Hub	eNodeB	0.1

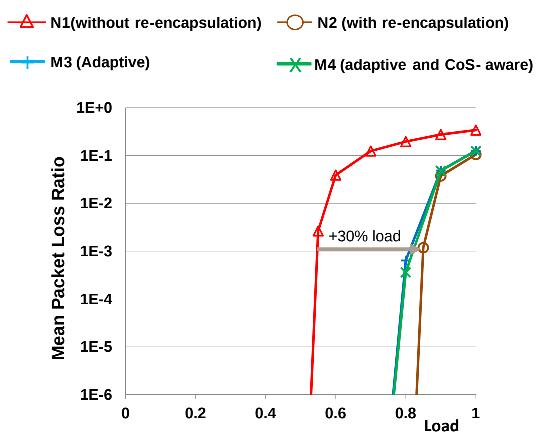

Fig. 12. Mean Packet Loss Ratio for all traffics.

(N1, N2, M3, M4) then compare it to the performance of the transparent forwarding mechanism.

A. Evaluation of electronic forwarding mechanism on realistic scenario

In this section we propose, using a realistic multi-CoS scenario to evaluate the previous electronic forwarding mechanisms: electronic forwarding with and without re-encapsulation, adaptive and CoS-aware adaptive forwarding mechanisms. We simulate an OSS ring with 7 eNodeBs each connected to an OSS node as described in Fig. 1; a selected OSS node (e.g., Node 0) acts as a traffic hub and collects traffic to/from backbone or other rings. Traffic characteristics are shown in Tabs. I. To correctly represent the handover procedure, not only the pair (source, destination) has to be concidred for the handover traffic but the triplet (traffic sender, Source eNodeB, Target eNodeB). The traffic demand value are given in Tab.II.

OSS rings are dimensioned to ensure the support of such traffic demand using the following custom algorithm: we allocate transponders for each node, starting with the most loaded node (allocating resources for unicast traffic first at a given node), and use first fit for wavelength allocation. CoS-management is performed by encapsulating RT and non-real-time (NRT) traffics in separated slots, and giving priority to RT slots at the channel insertion. The slot duration is 8 µs, the M3 threshold T1 = 1 slot, T2 = 40 slots, and timers are set to 450 and 900 µs for RT and NRT traffic, respectively. We assume geographically uniform user mobility.

Fig. 11 shows the latency (queuing and insertion) for all four proposed electronic mechanisms, for RT and NRT traffic forwarded by any intermediate node. Fig. 12 shows the mean packet loss ratio for all traffics. Load 1 (100%) corresponds to the load for which the dimensioning algorithm is run. The electronic forwarding mechanism without intermediate re-encapsulation keeps the latency low for the forwarded traffic but causes offered channel capacity waste. As shown in Figs. 11 and 12, the maximum supported network load is only 50% for a packet loss ratio (PLR) below 10^{-3} (above this limit the network is considered unstable). In contrary, all the offered channel capacity is used (supported load: RT up to 100%, NRT up to 80%) with the forwarding mechanism with intermediate re-encapsulation, at the expense of high latency for the forwarded traffic, which then becomes larger than the 500 µs target for RT traffic even for low loads. This is because the RT forwarded traffic experiences up to 2 RT timer durations: at source and intermediate nodes[2]. Using the adaptive forwarding mechanisms M3 and M4 the network benefits from each electronic forwarding mechanism depending on the network load: low latency and high supported load. With M3, the strict latency limit of 500 µs for RT traffic is met for load up to 55% only. With M4, by differentiating the processing for RT and NRT traffic, we are able to ensure very low latency below 500 µs for the forwarded traffic at low loads and save offered channel capacity at high loads. M4 ensures that RT forwarded traffic meets the latency constraint for load up to 80%, and is generally able to support a total network load of at least 95% with a negligible PLR for RT traffic (all losses are experienced by NRT traffic thanks to CoS-management.) Hence, M4 divides latency by up to factor 2 compared to N2 (Fig 11) and supports 30% additional network load compared with N1 (Fig 12).

B. Performance comparison between Electronic and Transparent forwarding mechanism

Using the same scenario described previously, we evaluate M5 and compare it with the best electronic forwarding mechanism (M4.) Fig. 13 shows that the transparent mechanism M5 presents same performances as the adaptive mechanism M4 for low to medium loads [0.05-0.5]. For RT forwarded traffic M5 achieves a slightly lower latency, since with M5 the forwarded traffic never goes through the electronic layer. Hence M5 ensures that RT forwarded traffic meets the latency constraint for load up to 100%. But this behavior has an impact on transparent and NRT forwarded traffic, since using M5 congestion and losses occur for loads above 70% as shown in Figs. 13 and 14.

V. COST/DIMENSIONNING COMPARISON BETWEEN ELECTRONIC AND TRANSPARENT FORWARDING MECHANISM

In this section we compare electronic (N1, N2, M3 and M4) and transparent (M5) forwarding mechanisms in term of dimensioning and cost.

[2]Using lower timer values would reduce latency but also maximum carried load.

TABLE I
TRAFFIC CHARACTERISTICS [6], [7]

Type	Demand	Pkt size (Bytes)	Pattern	CoS
Data	1.5 Gb/s up, 3 Gb/s down	1500	Centralized	RT, NRT
Voice over IP	8.5 Mb/s up and down	162	Centralized	RT
CoMP data	1 Gb/s	1500	Peer to peer	NRT
CoMP signaling	774 Kb/s	92	Peer to peer	RT
Handover inter-ring	0.1 Gb/s up and down	1500	Centralized	NRT
Handover intra-ring	0.1 Gb/s	1500	Peer to peer	NRT

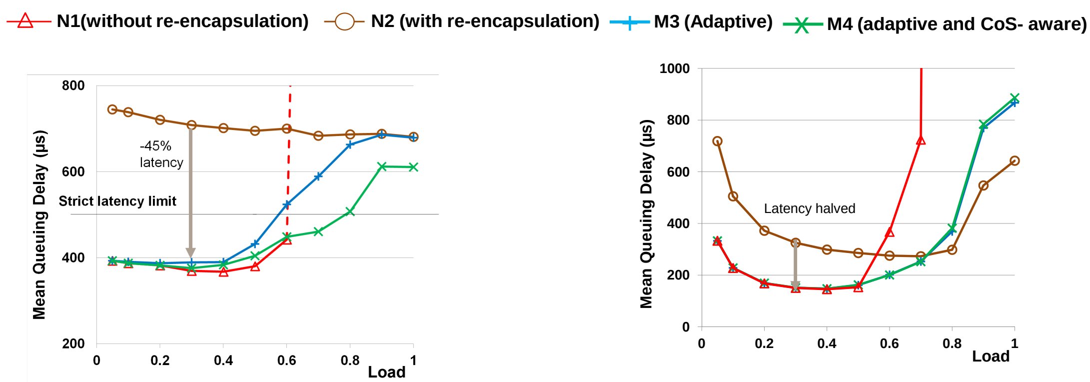

(a) Forwarded Real-Time traffic (b) Forwarded Non Real-Time traffic

Fig. 11. Comparison between electronic forwarding mechanisms N1, N2, M3 and M4

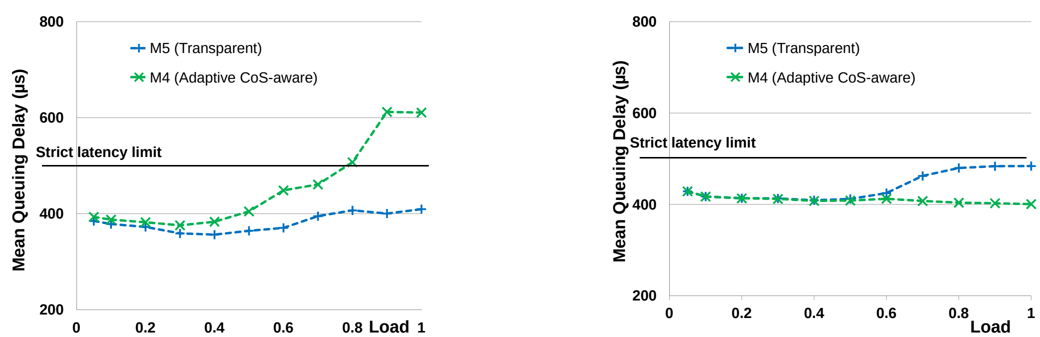

(a) Forwarded RT traffic (b) Transparent RT traffic

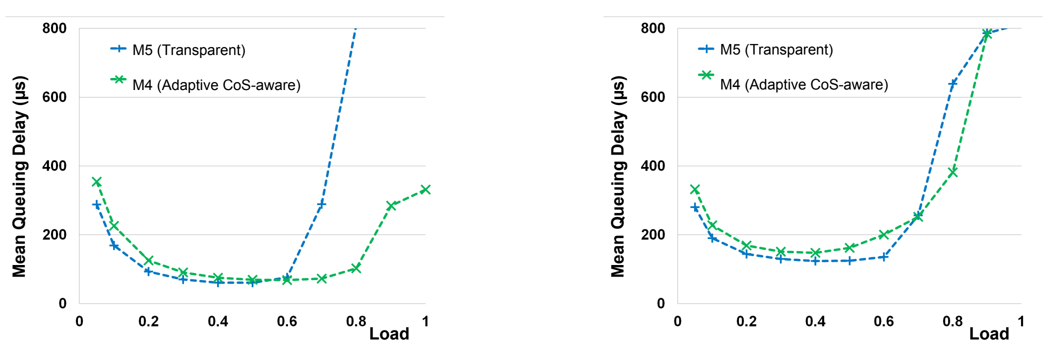

(c) Transparent NRT traffic (d) Forwarded NRT traffic

Fig. 13. Latency comparison between forwarding mechanisms M4 and M5.

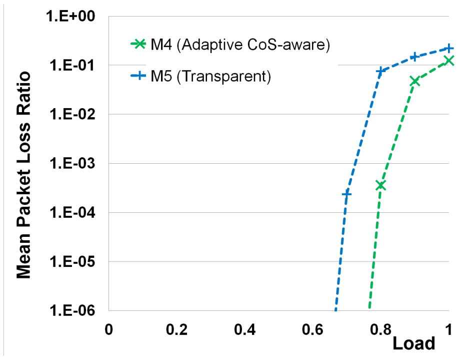

Fig. 14. Mean Packet Loss Ratio comparison between M4 and M5, for all traffic.

A. Dimensioning algorithm

In the network planning problem considered here, the objective function is to minimize the number of employed transponders (TRX). When allocating a TRX to a network node, one shall consider whether it is because a transmitter or a receiver is needed at that node, since a single TRX is composed of a fast-wavelength tunable transmitter and fixed-wavelength receiver. The dimensioning is done for a given traffic matrix, that contains the unicast, multicast and hand-over traffic demands, i.e. all types of traffic considered here.

1) Electronic forwarding Mechanisms N1, N2, M3 and M4:: It is easy to note that allocating the minimum number of transponders is polynomial problem in the case of electronic forwarding for any type of traffic. Indeed, in electronic forwarding, the entire traffic is electronically processed at each intermediate hop, and then, the packets are re-inserted into the optical medium, by using any of the available wavelengths. Consequently, to minimize the number of transponders is equal to minimizing both the number of transmitters and receivers needed at each node. Since the transmitters are fast-tunable, it is sufficient to minimize the number of receivers at each node to obtain the optimal solution. The algorithm that we use is the special case of use of the algorithm called "Minimize Receiver Cost First" - MRCF from [8]. The algorithm enables minimization of the number of needed wavelengths, by using the "first-fit" wavelength assignment. In order to avoid the trading between the receivers and wavelengths, each receiver starts wavelength assignment from the new wavelength.

2) Transparent forwarding Mechanism M5:: For case of the optical forwarding, the transit traffic is not OEO converted (not a subject of Optical-Electronic-Optical conversion), which helps in reducing the energy consumption. Since the intermediate nodes do not have the wavelength converters, the multicast flows need to be received on the same sets of wavelengths by all the nodes belonging to the same multicast group. Thus, there is a trade-off between the number of wavelengths and transponders that shall be allocated in the ring, and this alone would mean that the problem of finding the optimal design is non polynomial, since different OSS nodes can receive at disjoint wavelength sets. However, because of the presence of hand-over, any optical data packet can become the subject of handover forwarding by its downstream neighbor. The easiest way to achieve this is to enforce the sharing of all the wavelengths by all the destinations in the network. In another words, the same set of receivers will be used at each network node. Thanks to this, optical forwarding yields a polynomial design. To minimize the number of TRX for case of optical forwarding, we simply minimize the number of receivers per each node. Since the transmitters are fast-tunable, to final number of TRX is obtained as the maximum of the number of transmitters and receivers needed at each OSS node. Thus, we calculate the number of receivers $N_{rx}(i)$ and transmitters $N_{tx}(i)$ at node i , for a given traffic matrix. The needed minimum number of TRX per node i is then $TRX(i)$. The final number of TRX at each OSS node is $N^F_{trx} = max_i TRX(i)$, which is also the number of wavelengths in the network. The wavelength allocation is not performed, as we suppose that the traffic load is equally shared (load-balanced) over all channels, which is possible thanks to the fact that each node receivers on the same set of wavelengths. Note that routing the forwarding traffic (due to handover) is different in optic and electronic forwarding. Indeed, in electronic forwarding, the forwarded traffic is electronically converted first, and then re-inserted. Additional transponders are needed for the reinsertion of this traffic. For optical forwarding, the packets are forwarded transparently, by keeping the same wavelength. Multicast traffic can also change the wavelength in electronic forwarding scenario, because of the transit queueing of all traffic.

B. Cost comparison results

The cost comparison is performed for our 7-node ring, for different values of traffic load, obtained by scaling demand (Tab. I) by $a > 0$.

In the first scenario Fig. 15, we scale the entire traffic matrix (unicast + multicast + forwarding) by a factor α. We can see that electronic forwarding (N1, N2 ,M3 and M4) mechanism are much cheaper in terms of the number of needed transponders (Fig. 15a). In terms of wavelengths (Fig. 15b), the electronic solution needs over 2 times more wavelengths. It is since optical solution minimizes the number of needed wavelengths, since all wavelengths are shared.

To better understand the impact of the forwarding traffic on the final network cost, we consider the scenario 2 Fig. 16, where only the forwarding traffic is increased (case corresponding to urban scenario, with high density of quickly moving users). We scale the forwarding traffic by a factor β, while the other traffic (unicast + multicast) has the standard value ($\alpha = 1$). The idea behind considering the scaling of forwarding traffic separately is that electronic solution needs additional transmitters for re-insertion of the forwarding traffic, while optic forwarding keeps the forwarded packets on the same wavelengths.

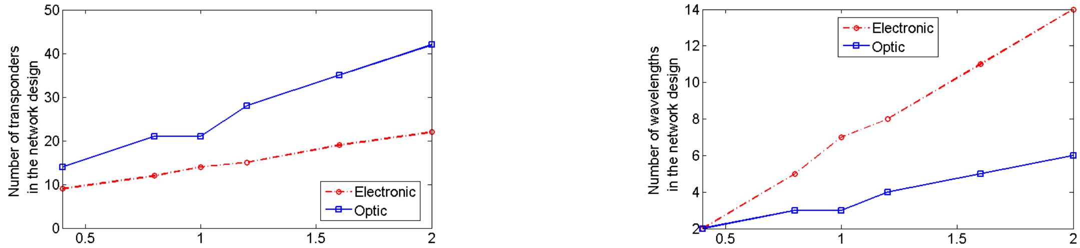

(a) Number of transponders (b) Number of wavelengths

Fig. 15. Scenario 1: 7-node network dimensioning, scaling the typical traffic matrix.

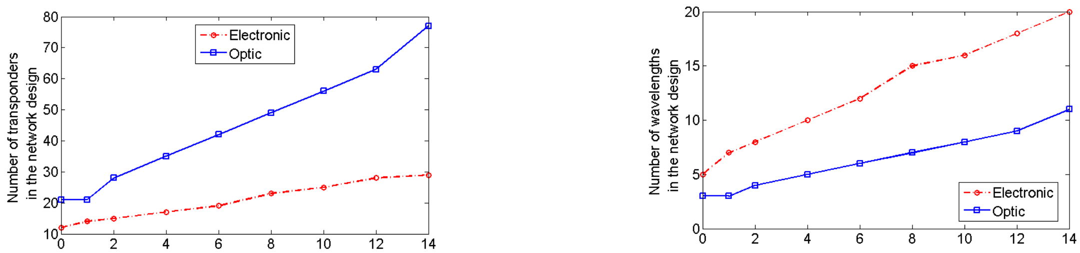

(a) Number of transponders (b) Number of wavelengths

Fig. 16. Scenario 2: 7-node network dimensioning, scaling the forwarding traffic .

However, independently of the values of forwarding traffic, electronic forwarding makes the network much less expensive. Optic needs over 3 times more transponders. Although the number of wavelengths is minimized in optic forwarding, all the nodes are mandatorily equipped with the same number of receivers, which increases the total network cost.

C. Discussion

We have seen in Fig. 13 that transparent forwarding M5 enables the lowest possible latency in an LTE-A mobile backhaul network as traffic subject to handovers bypasses the electrical layer completely. Latency gains, however, are limited to a few slot durations when the network is not congested, which is the time spent in queues at intermediate nodes when the load is reasonably low (i.e., when the network can indeed be operated). Since the slot duration is very small (2 orders of magnitude lower) compared with the latency constraints, the impact of transparent forwarding is also reduced. In addition, any energy gain brought in by all-optical forwarding is more than offset by the number of required TRX in the network, which not only drives energy consumption but also network cost to values higher than with electrical forwarding. Overall, if a source node knows whether its traffic is unicast or multicast, then some channels may be dedicated to unicast traffic and others to multicast traffic. However, in the context of mobile traffic, a source node cannot know beforehand whether a slot will be forwarded or not by an intermediate node towards its final destination, and the optical multicast option has an expensive design as shown previously, in this context, handling mobility traffic using the adaptive forwarding mechanism M4 improves both performance and cost.

VI. Conclusion

We compared several mechanisms to handle mobility traffic directly at the optical layer in an optical slot switching-based mobile backhaul network. We showed that, using an appropriate mechanism with electronic processing of mobility traffic only, the proposed network can provide very low (sub-ms) latency even for real-time traffic. Indeed, the proposed CoS-aware adaptive forwarding mechanism M4 presents the best trade-off between latency and offered channel capacity in optical slot switching-based mobile backhaul networks, where eNodeBs need to forward traffic from one cell to another e.g. upon user element mobility. We ensure a total network supported load of up to 80% while meeting the strict latency constraint of next generation mobile backhaul networks. Low latency is a very important requirement in the mobile backhaul segment but in some other segments such as datacenters, being able to ensure it, especially for forwarded traffic, without over-dimensioning the network is a goal reached by our adaptive forwarding mechanism.

Acknowledgments

This work was partly funded by DGE (France) and BMBF (Germany) through the CELTIC SASER project.

References

[1] Y. Pointurier, B. Uščumlić, I. Cerutti, A. Gravey, and J.-C. Antona, "Dimensioning and energy efficiency of multi-rate metro networks," *Journal of Lightwave Technology*, vol. 30, no. 22, pp. 3552–3564, 2012.

[2] N. Benzaoui, Y. Pointurier, T. Bonald, Q. Wei, and M. Lott, "Optical slot switching latency in mobile backhaul networks," *Lightwave Technology, Journal of*, vol. 33, no. 8, pp. 1491–1499, Apr. 2015.

[3] Q. Wei, M. Lott, V. Mikhail, and W. Kellerer, "Mobility management in optical mobile network," in *WCNC*, Apr. 2014, pp. 3130–3135.

[4] I. Shayea, M. Ismail, and R. Nordin, "Advanced handover techniques in lte- advanced system," in *ICCCE*, Jul. 2012, pp. 74–79.

[5] Q. Wei, J. Bazzi, M. Lott, and Y. Pointurier, "Multicast in mobile backhaul with optical packet ring," in *WiMob*, Oct. 2013, pp. 181–188.

[6] NGMN alliance, "Guidelines for lte backhaul traffic estimation," white paper, https://www.ngmn.org/uploads/media/NGMN_Whitepaper_Guideline_for_LTE_Backhaul_Traffic_Estimation.pdf.

[7] S. Brueck, Z. Lu, J. Giese, and M. Amin, "Centralized scheduling for joint transmission coordinated multi-point in lte-advanced„" *WSA*, pp. 177–184, Feb. 2010.

[8] B. Uscumlic, I. Cerutti, A. Gravey, P. Gravey, D. Barth, M. Morvan, and P. Castoldi, "Optimal dimensioning of the WDM unidirectional ECOFRAME optical packet ring," *Springer Photonic Network Communications*, vol. 22, no. 3, pp. 254–265, 2011.