Impact of transponders and regenerators wake-up time on sleep-mode enabled translucent optical networks

Albert Pagès[1], Massimo Tornatore[2], Jordi Perelló[1], Salvatore Spadaro[1], Annalisa Morea[3]
(1) Universitat Politècnica de Catalunya (UPC), Barcelona, Spain; (2) Politecnico di Milano (PoliMi), Milano, Italy; (3) Alcatel-Lucent Bell Labs, France; e-mail address: albertpages@tsc.upc.edu

Abstract: Sleep-mode enabled transponders and regenerators yield to substantial energy savings; however, their non-negligible wake-up time may degrade the network performance. We show that an appropriate dimensioning of the devices per node can compensate such effect.

OCIS codes: (060.4250) Networks; (060.4510) Optical communications.

1. Introduction

Wavelength Division Multiplexing (WDM)-based optical networks have helped reducing the power consumption of core networks by allowing large amounts of traffic to bypass routers through optical switching [1]. However, additional energy savings in translucent optical networks can be achieved by using energy efficient opto-electronic (OE) devices. In [2], authors presented a novel architecture for both transponders and regenerators capable of putting some of the transponder (TSP) modules either in a low-power consumption mode (*idle* mode), or turning them off (*off* mode) when they do not support traffic; TSP modules can be then turned back in a fully operative (*on*) mode through suitable control plane-based procedures whenever new connections are established. Dynamic management of TSP power states is carried on according to current traffic load and to the requirements in terms of connection set-up time. With such a power management, the authors in [2] demonstrated that up to 56% power savings can be attained with respect to traditional WDM networks, where devices are always powered-on regardless if they are transmitting data or not. Authors extended their work in [3], where different classes of traffic are considered. In such a context, the set-up time requirements of the connections become a critical parameter, as typically high-priority traffic (e.g., real-time traffic) requires short set-up times, while low-priority traffic does not. Considering that the boot-up time (i.e., the time required for the transition of a TSP or a regenerator (REG) from the *off* to *on* mode) can be quite long, available TSPs and REGs in the *off* mode cannot be allocated to support high-priority traffic, because its requirement in terms of set-up time can easily not be matched. Therefore, TSPs and REGs in the *off* mode should only be allocated to low-priority traffic, while TSPs and REGs in *idle* mode should be used to allocate high-priority traffic, since the much lower activation time (time required for the transition from the *idle* to *on* mode) can match more stringent set-up time requirements. Note that, by reserving some of the available devices per node to high-priority traffic and putting them on *idle* state, it is possible to improve the overall readiness of the network, satisfying both high and low-priority connections requirements, and, at the same, reducing the power consumption of the network.

Although there are works that analyzed the blocking probability in sleep-mode enabled optical networks, such as in [4], there are no works, to the best of our knowledge, on the evaluation of the impact of the wake-up time of node equipment on the overall network blocking probability (especially for high-priority traffic) in translucent core networks, considering both TSPs and REGs. Essentially, depending on the duration of the TSP/REG wake-up process, and the arrival rate of the connections, there may be some period of time where devices in *idle* state to be allocated to incoming high-priority connections may still not be ready. Therefore, a proper dimensioning of the number of devices that are put in *idle* mode per node to support high-priority traffic must be done. Then, the impact of the wake-up time on the blocking of the connections must be evaluated. To this end, in this work we evaluate the degradation on the overall network performance in terms of blocking probability due to the wake-up time of the devices in a translucent optical network. The evaluation of the power savings with respect to traditional WDM scenarios is also performed.

2. TSPs and REGs allocation strategy

In this paper, we consider two classes of connections, namely, high-priority and low-priority connections, characterized by two different requirements in terms of connection set-up time. Specifically, we assume that the high-priority connections have stringent set-up time requirements. To manage the establishment of the high-priority connections according to their set-up time target, a possible energy-efficient strategy consists on reserving part of the node resources (TSPs and REGs) for high-priority traffic, setting them in *idle* mode for a prompt establishment of such connections. Assuming that every node is equipped with *N* TSPs and *N'* REGs, a fraction of such resources

---

The work on this paper has been funded by Spanish Science Ministry through the project ELASTIC (TEC2011-27310). The work is part of the Green Touch Consortium.

(*m* and *m'*, respectively) are set in *idle* state and will be exploited by high-priority traffic. The rest of the unused resources (those that are neither transmitting/receiving data nor reserved for high-priority traffic) are left in *off* mode and used for future low-priority incoming connections. In order to maintain the pool of resources reserved for high-priority traffic, every time an *idle* device (either a TSP or REG) is allocated for a new high-priority connection (and thus set to *on* mode), if and only if spare *off* devices exist, a wake-up operation will be triggered and an *off* device will start changing to *idle* mode. Hence, high-priority connections are blocked if no *idle* resources are available. As for low-priority connections, if no *off* devices are available when they arrive, they are also blocked, since the devices in *idle* state are reserved only for high-priority connections. Additionally, when a high or low-priority connection is torn-down, if the number of *idle* devices is lower than *m* (TSPs) or *m'* (REGs), the released device is hibernated to *idle* mode. Otherwise, it is completely shut down (*off* mode).

Regarding the wake-up operation (transition from *off* to *idle*), its duration can dramatically affect the blocking levels on the network, specifically for the high-priority traffic. After an *idle* device is selected to support an incoming high-priority connection, a wake-up operation is also triggered to maintain the pool of reserved resources (a spare *off* device is set to *idle*). But, in case that additional *m*-1 (*m'*-1) high-priority connections have to be established employing that particular node during this wake-up operation, all *idle* devices will become exhausted. So, any further high-priority connection arriving before the ending of the wake-up operation will be blocked due to the lack of *idle* devices. Depending on the wake-up time (T_w) and inter-arrival rate of the connections, the blocking probability may achieve unsustainable levels for network operators.

3. Test scenario and Results

To analyze the impact of T_w, we have executed simulations employing the 16-node European Optical Network (EON) [5] with real physical distances [6]. We are considering a translucent scenario, with nodes equipped with a pool of REGs to perform regeneration when needed. We assume that bidirectional connections arrive at the network following a Poisson distribution, with exponentially distributed inter-arrival times (IATs) and holding times (HTs). Thus, the network load is defined as HT/IAT. All offered traffic is uniformly distributed among all source/destination pairs. Moreover, it is assumed that all connections request for a single wavelength. The bit-rate of the connections is 100 Gb/s, employing a 28 Gbaud PDM-QPSK modulation, for which the transparent reach without need of regeneration is 1200 km [7].

Connections are routed employing a K-Shortest Path (K-SP) strategy to better exploit the network resources, choosing the first available path among the candidates set. Wavelengths are chosen under a First Fit (FF) criterion, and the wavelength continuity constraint must be ensured along regeneration spans. Connections with physical paths longer than the transmission reach are regenerated in the first potential node along the path. REGs are assumed to allow wavelength conversion. For a base scenario with an ideal wake-up time ($T_w = 0$ s.), we have dimensioned the network resources (wavelengths, TSPs and REGs) in order to have an overall blocking probability (considering high and low-priority traffic classes) not higher than 1%, assuming an offered network load of 140 and mean HT = 1 h., with a share of high and low-priority traffic of 30% and 70%, respectively, and K = 6 candidate paths in K-SP. As a result, we have found that all network links have to be equipped with 40 wavelengths, and all network nodes with 30 TSPs and 15 REGs. The number of *idle* devices reserved for high-priority traffic in the basic case (*m* and *m'*) is 2 and 1, respectively.

Taking this base scenario, we analyzed the impact of the value of T_w on the network blocking probability (BP). Since the impact of T_w is tightly related to the traffic dynamicity (how many high-priority connections may arrive during a wake-up operation), we have investigated, for various degrees of dynamicity and values of T_w, how BP values of high-priority traffic evolve. To this end, we define $r = HT_{base}/HT$ as the dynamicity of the traffic, denoting HT_{base} the mean HT in the base scenario (i.e., 1 h.), modifying the mean IAT of the connections in all cases so as to maintain the same network load of 140. Note that for higher *r* values, both the mean HT and IAT of the connections are reduced. Hence, under the same offered load, more dynamic connections are being established. Fig. 1 (left) depicts the evolution of the BP for high priority traffic as a function of *r*, considering multiple values of T_w. The base scenario ($T_w = 0$ s.) is also depicted as a reference.

It can be appreciated that the BP of high-priority connections increases with the dynamicity of the offered traffic. In this sense, it can be seen that low T_w values allow for a high range of dynamicity without significant performance degradations ($BP \leq 1\%$). As the value of T_w starts growing up, lower traffic dynamicity can be supported. For example, if $T_w = 1$ min, the base dynamicity ($r = 1$) is barely supported and if $T_w = 5$ min, even the lowest dynamic traffic cannot be supported without significant network performance degradation. Although not shown in the figure, for all the tested scenarios, the observed BP of low priority traffic remains below 1.25%.

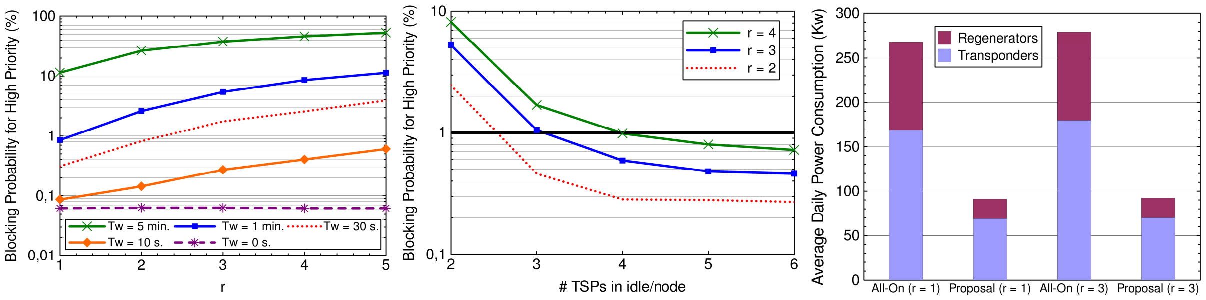

Fig.1- BP as a function of *r* (left); BP as a function of the number of *idle* TSPs (center); Average daily network power consumption for various scenarios (right).

In light of these results, we have investigated the main causes of blocking of high-priority connections. The vast majority of the blocking (70-90%) is due to the lack of *idle* TSPs, whereas the second cause is due to the lack of *idle* REGs, as every connection employs 2 TSPs, while REGs are only used sparsely. Note that these values might change depending on the dimension of the network. Taking the case T_w = 1 min., we have investigated for various values of *r* how many TSPs per node should be set to *idle* in order to reduce the BP of high-priority traffic below 1%. Fig. 1 (center) depicts the obtained results, where the black thick solid line marks the desirable 1% BP value. As shown, although in the base scenario 2 *idle* TSPs per node were reserved for high-priority traffic, an additional *idle* TSP per node should have been reserved to meet the desired BP if $r = 2$, while 2 additional TSPs are required to meet this objective for $r = 3$ and 4. Additionally, note that, as the number of *idle* TSPs per node increases, the reduction in BP starts to stabilize. This is due the fact that, with enough *idle* TSPs, the main component of the blocking starts to be the lack of *idle* REGs. Hence, to further reduce the BP of high-priority connections, more *idle* REGs should have to be reserved per node. Note, however, that as more TSPs are set to *idle*, the BP of low-priority connections starts to increase, since fewer resources can be exploited by them. Hence, to maintain low-priority BP constant, a number of *off* TSPs equal to the additional TSPs set to *idle*, should have to be equipped in the nodes, thus increasing the overall network cost.

Finally, we also have evaluated the average daily power consumption under the daily traffic profile presented in [8] that can be achieved with the proposed energy-aware TSP/REG allocation strategy. To this goal, we have set T_w = 1 min. and $r = 1$ and 3, comparing it to the case where all TSPs and REGs are always active (*on* mode). For the sake of fairness, in the all-on scenario, the same dimensioning used to get the results presented above has been assumed. Moreover, for $r = 3$, two additional *idle* TSPs are reserved per node, as obtained from Fig. 1 (center), increasing also by 2 the number of TSPs set in *off.* Fig 1 (right) depicts the obtained results, separating the power contributions from TSPs and REGs. We assume the same power figures for the devices as in [2]. It can be appreciated that huge savings (around 65%) are attained when the proposed power management strategy is employed. Furthermore, we observe that it is possible to compensate the effect of T_w in more dynamic scenarios (r = 3), with only a marginal increase in the power consumption (around 1.3%) when compared against less dynamic scenarios ($r = 1$). This is due to the fact that the additional resources required to compensate the increased traffic dynamicity only consume a significant amount of power when really needed, while they benefit from the low power *idle* mode when not used.

4. Conclusions

In this work, we evaluated the impact of the wake-up time of transponders and regenerators in a multi-service translucent optical network employing an intelligent power management strategy. We showed that significant degradations on network performance take place as the device wake-up time increases for very dynamic scenarios. We demonstrated that this effect can be suppressed with a small increase of the *idle* devices per node at expenses of a negligible increase (around 1.3%) on the average network power consumption.

References

[1]. Y. Zhang et al., Comm. Surveys and Tutorials, vol. 12, num. 4, pp. 441-458, 2010.
[2]. A. Morea et al., Proc. of ECOC 2011, We.8.K.3.
[3]. A. Morea et al., Proc. of OFC/NFOEC 2012, OM3G.6.
[4]. F. Musumeci et al., Proc. of ICTON 2013, We.A4.5.
[5]. S. De Maesschalck et al., Photonic Net. Comm., vol. 5, num. 3, pp. 203-225, 2003.
[6]. A. Betker et al., MultiTeraNet Report, 2003.
[7]. O. Rival et al., Proc. of OFC/NFOEC 2011, OTuI4.
[8]. A. Morea et al., J. Lightw. Technol., vol. 31, num. 5, pp. 779-785, 2013.