

Demonstration of GMPLS-Controlled Device Power Management for Next Generation Green Optical Networks

Annalisa Morea¹, Jordi Perelló², Fernando Agraz², Salvatore Spadaro²
(1) Alcatel-Lucent Bell Labs France, (2) Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
Author e-mail address: annalisa.morea@alcatel-lucent.com

Abstract: We propose GMPLS protocol enhancements enabling dynamic power state management of optoelectronic network devices for setup-time-aware differentiated connection provisioning. Daily energy savings of 56% are demonstrated in a Pan-European-like network compared to always on scenarios.

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

1. Introduction

To cope with the increase of the required Internet bandwidth, numerous standardization bodies and network researchers have been engaged in improving the energy efficiency of telecommunication networks. Current optical backbone networks work statically, that is connections are set-up at the beginning or during the network life and last for several months or even years. For this reason, deployed devices are conceived and managed under the assumption that all network resources are always on. Such static behavior is inefficient, especially because the traffic transported by the network presents relatively predictable variations such as daily and weekly fluctuations.

Various research works have proposed to adapt the energy consumption of the network to the amount of transported traffic by letting the network devices take two power states, either be on or be in sleep mode. In [1] strategies for managing the link sleep mode has been proposed, while in [2] our proposal concerns the power management of optoelectronic (OE) devices. While allowing devices to switch to sleep mode provides considerable energy savings (40%) at low loaded networks, we consider this strategy unpractical because it jeopardizes the network resiliency by reducing the network connectivity in case of fiber cuts some node indeed become isolated and recovery path can be unavailable. On the other hand, switching OE devices on and off does not compromise the network topology, guarantying a level of resiliency comparable to the one of current networks and has been shown to bring up to 60% energy savings compared to always-on architectures.

Specifically, in [2] we proposed a protocol enhancement for the management of three power states of OE devices. An OE device can be totally or partially powered (*up* and *idle* state, respectively) or switched off (*down* state). The transitions between the different power states are managed by a node as a function of the number of devices at a given power state in it. In this paper we propose GMPLS protocol enhancements for dynamically managing the transponders to be associated to an incoming connection requests according to their service class.

2. Proposed scheme

The proposed scheme towards next-generation green optical networks leverages on extended signaling and routing functionalities of GMPLS to effectively manage the power state of the network transponders and regenerators according to incoming traffic variations. As detailed in [2], three different power states can be implemented in OE devices, namely, *up*, *idle* and *down*. The *up* state describes a fully powered and operational OE device; the *down* state describes an unpowered and unused one; while the intermediate *idle* state represents an unused and semi-powered OE device, where the components requiring thermal stabilization are powered on, allowing a fast wake-up time (few tens of ms). Idle state ensures around 90% energy savings compared to a fully powered OE device [2].

In contrast to the work in [2], which did not differentiate among connection provisioning requests, a prioritized network scenario is addressed in this paper, where gold, silver and best-effort traffic classes coexist. The figure of differentiation among traffic classes is the connection setup time, which is achieved in an energy efficient way by partitioning OE devices to allocate for one or another traffic class. In particular, we target to a maximum connection provisioning time of 50 milliseconds for the gold traffic class, 1 second for the silver one and a few minutes for the best effort traffic [3]. The *idle* to *up* state transition time risks to endanger the gold class connection provisioning time; for this reason we configure part of the network transponders and regenerators to always remain *up*, while silver and best-effort traffic classes are allocated on *idle* and *down* OE devices, respectively. Beside these available *up* OE devices, we ensure in every node the presence of enough transponders and regenerators at the *idle* state to serve silver connections, and the rest of devices are set down state and used by the incoming best-effort requests. As in [2], the power state management relative to an OE is locally performed by the node hosting it.

The choice of an OE device at a given status to the incoming request is orchestrated from the GMPLS control plane building on the RSVP-TE signaling extensions already proposed in [2], which allow the device specification

in the Path message. As depicted in Figure 1 (right), a new ‘S’ bit is added in the RSVP-TE ADMIN STATUS object. In this way, together with the already defined ‘A’ bit, the aforementioned OE device power states are represented. When a gold connection is signaled, both ‘A’ and ‘S’ bits of the ADMIN STATUS object are set to ‘0’, indicating that up state transponders/ regenerators must be used. For silver connections, the ‘A’ bit is set to ‘0’ and the ‘S’ bit is set to ‘1’, meaning that idle transponders/regenerators are needed to allocate them. Finally, best-effort connections are signaled with the ‘A’ bit set to ‘1’ and the ‘S’ bit to ‘0’.

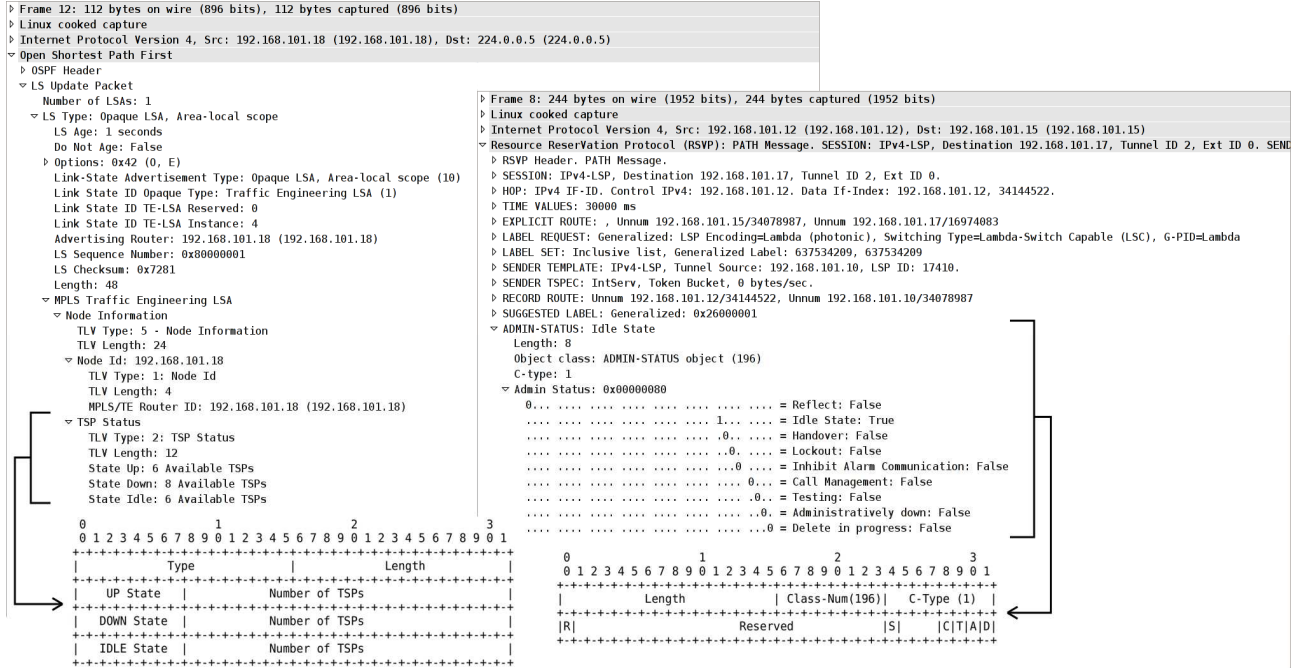


Figure 1. Proposed TSP Status TLV for OSPF-TE (left) and Admin Status TLV usage in RSVP-TE (right).

Note that in [2], connections were signaled over shortest paths with no knowledge of either wavelength availability on links or regenerators at nodes. This approach resulted in high connection blocking probabilities due to the lack of network resource knowledge during the connection signaling phase. Aiming at improving network performance, this work also proposes extensions to the GMPLS OSPF-TE protocol, so that the available regeneration capabilities at any node is disseminated to all the remainder nodes in the network, jointly with the wavelength availability information. As depicted in Figure 1 (left), a new sub-TLV (Type Length Value) named TSP Status is introduced in the OSPF-TE opaque Link State Advertisements (LSAs) containing the number of *up*, *down* and *idle* transponders in a node (an add/drop device, A/D, is represented by a transponder, while a regenerator, REG, by two transponders). This TLV is inserted into a Node Information top level TLV (type 5) [4], as transponder availability is a property of the optical cross-connect (OXC) and not of the link. Besides, a Node Identifier sub-TLV allows the network routing entity, i.e. the Path Computation Element (PCE), to associate the OE devices to their respective node. Indeed, taking the contents of the OSPF-TE opaque LSAs, the PCE can populate its Traffic Engineering Database (TED) with both wavelength and regenerator availability information, which is used to compute the end-to-end routes requested by its client Optical Connection Controllers.

3. Experimental results

Aiming to assess the benefits of the proposed GMPLS-controlled OE device power management scheme, experimental evaluations have been conducted over the ASON/GMPLS CARISMA test-bed, located at the UPC premises in Barcelona. A Pan-European optical network composed of 16 nodes and 23 links has been configured for these purposes, where each link carries 10 bidirectional wavelengths. We consider 100Gb/s Polarization-Multiplexed QPSK modulated channels with 1200km of transparent reach [5].

Regarding the offered traffic characteristics, uniformly distributed connection requests arrive at the network following a Poisson process. Connection Holding Times (HTs) are exponentially distributed with mean 3 hours. Hence, different loads are generated by changing the connection Inter-Arrival Times (IATs) accordingly (i.e., load = HT/IAT). We assumed peak traffic of 80Erlang and to emulate daily traffic variations, we consider the off-peak traffic as 50% the maximum load [6], reached with 10Erlang scaling steps. For the service class distribution, we assume a 20/30/50% gold/silver/best-effort traffic ratio. To allocate this peak traffic, network nodes are equipped with 20 local add/drop transponders and 10 regenerators in total. Resources are partitioned among the different

classes: 6 A/Ds and 3 REGs are reserved for gold class connections; among these devices at least one unused A/D and REG are always up and the remaining unused ones are set idle. Then 6 A/Ds and 3 REGs are for silver class connections, the unused devices are always idle. Finally, 8 A/Ds and 4 REGs are reserved for best effort connections and unused devices are always down. With this device partitioning we ensure 1% gold class blocking probability for the peak load, as shown in Fig.2(a); higher blocking probabilities are observed for the other classes.

We consider a centralized PCE responsible for assigning the most appropriate routes, wavelengths and regenerators for those incoming connection requests. In order to avoid contention of resources among different connections under establishment, a pre-reservation of resources, is implemented in the PCE. Throughout the evaluations, we assume idle-up and down-up state transition times of 20ms and 60s [2], respectively. Under these assumptions, gold class setup times result in average ~ 10 ms. In the opposite extreme, best-effort traffic suffers from very high OE device wake up times, thus resulting in setup times ~ 150 s. Contrarily, using *idle* OE devices to allocate silver connections, in average 60ms setup times can be achieved. We remind that if the current protocol for the connection setup is used, no setup time differentiation will be observed (all connection setup times equal 10ms) because neither device partition nor power management is performed.

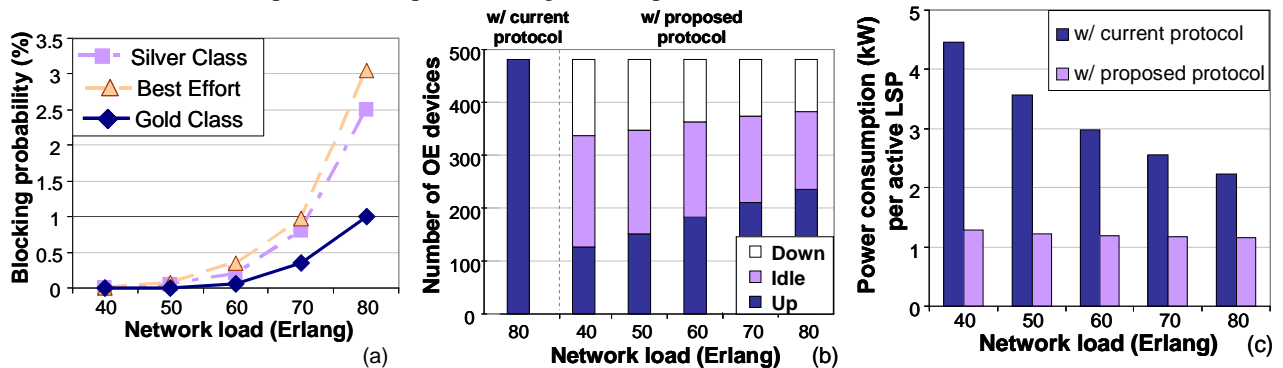


Figure 2. Blocking probability (a); total number of OE devices in *up*, *idle* and *down* states (b); and power consumption per active LSP (c).

We compare now the number of powered OE devices at the different power states when the proposed protocol enhancement and the current protocol (without power management) are used (both ensuring 1% of blocking gold connection at 80Erlang). For the current protocol the number of unused resources in a node is independent of the number of active connections, opposite to the proposed protocol. From Fig.2(b) we observe that the number of down devices varies from 30% to 20% at low and high traffic loads, respectively; and the number of idle devices ranges from 44% to 31%, respectively. If now we estimate the energy advantages of this power management, we quantify in Fig.2(c) the average power consumption per active connection (LSP) in the network with the two policies. Power assumptions are the same as in [2]. Thanks to the introduction of the OE power state management, 70% and 50% power savings are achieved at low and high loads, respectively. We also observe that when the proposed enhancement is implemented, the power consumption per LSP slightly varies with the traffic load (peak vs. off-peak: -10%), while with the current protocol the LSP power consumption is inversely proportional to the traffic load (peak vs. off-peak: -50%). If now we consider daily traffic variations as in [6], the total amount of energy savings is: 56%. We proved that the introduction of device partitioning for connection service class only slightly decreases the network energy efficiency (4% lower than [2]) and is suitable for service class performance requirements.

4. Conclusions

In this paper we proposed an enhancement of current GMPLS signaling and routing protocols enabling the dynamic power management of optoelectronic devices in the optical layer as a function of the request service class. A policy for the management of optoelectronic devices at different power states is introduced such as to guarantee a limited blocking ratio for high priority requests (gold class). Energy savings of 56% are observed when compare the energy requirements for the network managed with the proposed GMPLS protocol enhancement with the one managed with current GMPLS protocol set.

This work has been partially supported by the Spanish Science Ministry through Project ELASTIC (TEC2011-27310) and by the French public authorities in the frame of 100GRIA project.

5. References

- [1] I. Cerruti et al., J. Light. Techn. Vol.29, No. 15, pp. 2292-2298.
- [2] A. Morea, et al., Proc. of ECOC 2011, We.8.K.3, Sept. 2011.
- [3] A. Morea, et al., Bell Labs Techn. J., Vol.14, No.4, pp. 193-211.
- [4] R. Martinez et al., J.of Light. Techn., Vol.28, No.8, 1241-1255.
- [5] O. Rival, et al., Proc. of OFC 2011, OTuI4.
- [6] <http://asert.arbornetworks.com>.