Power Management of Optoelectronic Interfaces for Dynamic Optical Networks

A. Morea[1], S. Spadaro[2], O. Rival[1], J. Perelló[2], F. Agraz[2], D. Verchere[1]
(1) Alcatel-Lucent Bell Labs, France (2) Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
annalisa.morea@alcatel-lucent.com

Abstract: We propose a protocol enhancement to manage the power state of optoelectronic interfaces in automatically reconfigurable optical networks. For dynamic traffic scenarios 67% of energy savings have been estimated for the optoelectronic interfaces.

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

1. Introduction

With the improvement of ICT technologies over the last few years and their price erosion, more and more services (e.g. 3D or ultrahigh definition television) have been developed and have become an indispensable element of everyday activities. An immediate outcome has been the sustained rapid Internet traffic increase (~50% CAGR) with no abating in sight in the near future [1]. To support this traffic growth, research activities the improvement of network energy efficiency has become a growing concern [2]. Indeed, while nowadays ICT consumes ~ 8% of the total electricity over the world, this amount may double over the next 5 years [3]. Hence, power requirement will be a major constraint of next generation networks.

Today's transported traffic presents dynamic evolutions, some of which are predictable, e.g. daily variations [4]. But optical networks are quasi-static, with connections established for months or years. This requires all network elements to be powered up for peak traffic (including enough over-provisioning to absorb unexpected variations) and is thus a major source of energy inefficiency. A first step towards making networks more efficient is the introduction of dynamics, with the possibility of lighting up and tearing down optical connections according to traffic variations. Control plane protocols, such as the Generalized Multi-Protocol Label Switching – GMPLS –, protocol set, enable the aforementioned dynamic connection management. However, to enhance the energy efficiency of the network, such protocols will need to be coupled with an ability to switch on and off the underlying optoelectronic devices (OE), (i.e. add/drop transponders and regenerators) as a function of the network load.

In this paper, we propose a policy for powering on/off OE devices in nodes so that connection set-up times are cushioned against OE device transient effects. The proposed policy enhances GMPLS towards next-generation green optical networks. In the following, the description of the protocol extensions and the list of components present in 100Gb/s OE devices with their relative power consumption are provided. This information will be used in a GMPLS-enabled optical network test-bed to estimate the energy efficiency improvement of the proposed policy.

2. Power management and description of OE devices

In the GMPLS signalling protocol (i.e. RSVP-TE) [5] a bit, A, is committed to administratively set up or down a specific Label Switched Path. While $A=0$ during normal network operation, the network administrator can set $A=1$ in the RSVP-TE Path/Resv messages to request for specific actions to be taken locally at each network node. In this work, we propose to enhance this utility to enable a dynamic power management by defining a set of actions to be performed upon reception of RSVP-TE Path/Resv messages. To do that, we associate three 'power' states to each OE device: *up*, *idle* and *down*. Transitions between these states are described in Fig. 1. *Up* and *down* states describe the OE fully powered on and off, respectively. The *idle* state describes an OE not transmitting any traffic but ready to do so at short notice. We suppose that the transition delay from *idle* to *up* state is comparable to the delays required for the connection set-up (protocol forwards and optical switching) while the power requirement in the *idle* state is a small fraction of the one in the *up* state. The *idle* state is introduced to allow thermal stabilization of critical components (e.g. laser). The time required for the stabilization of such components ('wake-up' transition) is assumed shorter than the demand inter-arrival time.

Therefore we propose to extend the signaling protocol in [5] by introducing a new bit 'S' so that when $A=0$, S discriminates whether the OE devices are *up* or *idle*. Conversely, if $A=1$, S discriminates between *down* and *damaged*, as indicated in Table 1. The management of the power state relative to an OE is performed by the node hosting it. In this study we suppose that in every node there is always at least one transponder and one regenerator *idle*, while all remaining unused devices are *down*. For a high priority connection request, the bits A and S in the related RSVP-TE Path message are set to 0 and 1, respectively, thus requesting for the *idle* resources in order to minimize the connection set-up time. During the signaling phase, the transition of one additional OE device

978-1-55752-932-9/11/$26.00 ©Optical Society of America

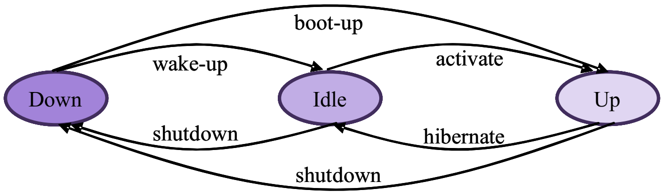

Figure 1: Scheme of the three proposed power states relative to an optoelectronic device and the transitions between them.

Table 1: Protocol description of the different states relative to an OE device.

Proposed power management	A	S
Up state	0	0
Idle state	0	1
Down state	1	0
Damaged	1	1

(equivalent as the one currently allocated) from down to idle is triggered, so as to ensure at least one *idle* OE device for future connection demands. In contrast, for a low priority connection request, the bits *A* and *S* in the related RSVP-TE Path message are set to 0 and 0, respectively. Thus, the boot-up of *down* resources for the new connection request are triggered, which is finally established once the underlying OE devices are stabilized to the *up* state. When a connection is released, the node hosting each OE associated to the connection controls if there is already at least one equivalent device in *idle* state in the node. If so, the freed-up OE is set to *down*, otherwise it is set to *idle*.

In this study the OE devices are: transponders and regenerators (Fig. 2). Transponders receive data from the client side and encode it on a wavelength in the optical domain after suitable framing and addition of a forward error correction (FEC) overhead. Due to the bi-directionality of traffic, transponders are bi-directional and symmetric operations are guaranteed from the WDM optical domain to the client side. Whenever the signal becomes too degraded to reach the destination transparently, regenerators receive the optical signals and re-emit it after FEC.

In this work, we consider 100Gb/s signals with Polarization-Multiplexed QPSK modulation. The transponder unit is connected to the router through 10Gb/s lines (client cards). In Table 2 we list the typical components required for such 100Gb/s transponders (TSP) and regenerators (Reg) [6], their respective power consumption as well as components staying powered-up when the device is *idle*. We assume that the components requiring high stability before being fully operational are the lasers at the emitter and the receiver (~60s [11]). Other devices require from micro- [13] to milli-seconds [7] to be operational; these values are of the same order of magnitude than the time required by the control plane signaling functionality to set-up the connection. Hence we suppose their activation does not delay the setting up process. A complete study would require inclusion of thermal effects within whole racks rather than at the component level. This complex subject is however well outside the scope of this article.

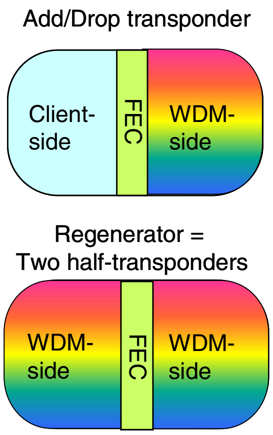

Figure 2: Schematics of an add/drop transponder and a regenerator.

Table 2: Components in 100Gb/s OE devices at *up* and *idle* state and their power consumption.

	Component	Unit	TSP	Reg	Power dissipation (W)	Powered at *idle* state	Ref
Client side	Client-card	10	•		3.5		[7]
	Framer	1	•		25		[8]
FEC	FEC	1	•	• (x2)	7		[9]
E/O modulation	Drivers	4	•	• (x2)	9		[10]
	Laser (on-off)	1	•	• (x2)	6.6	•	[11]
O/E Receiver	Local oscillator	1	•	• (x2)	6.6	•	[11]
	Photodiode + TIA	4	•	• (x2)	0.4	•	[12]
	ADC	4	•	• (x2)	2		[13]
	DSP	1	•	• (x2)	100		[13]
FEC	FEC	1	•	• (x2)	7		[9]
Client side	Deframer	1	•		25		[8]
	Client-card	10	•		3.5		[7]
	Management power		•	•	+20% total power	•	[12]
	Total power		**351W**	**414W**		**18W TSP** **36W Reg**	

3. Network scenario

The value proposition of the energy savings due to the introduction of *idle* and *down* states in dynamic reconfigurable optical networks has been experimentally evaluated in the GMPLS-enabled CARISMA test-bed located at the UPC premises in Barcelona. For this evaluation, a Pan-European network composed of 16 nodes and 23 links has been configured [14], where each link carries 10 bidirectional wavelengths. In order to accommodate the incoming traffic each node is equipped with 20 add/drop transponders and 10 regenerators. No dispersion management is configured in the test-bed; hence, the transparent reach of 100Gb/s signals is set to 1200 km [15].

Regarding the traffic characteristics, uniformly distributed connection requests arrive to the network following a Poisson process with average holding time equal to 3 hours. In this work, the maximum offered load has been 80 Erlang, which results into a blocking probability around 1%. Such an offered load represents the daily peak traffic.

To take into account the daily traffic variations, we progressively scale down the offered traffic at off-peak hours to 50% of the peak traffic [4], decreasing the load by 10 Erlang steps. For simplicity, in this first study we assume that the entire offered traffic is high-priority traffic, so that no direct *down* to *up* state transitions are triggered.

5. Experimental results

In commonly studied dynamic optical networks all resources have to be lit-up to ensure the establishment of any incoming connection. For the studied network and offered loads all OE (335 add/drop transponders and 160 regenerators) in the nodes have to be powered on, yielding a total 184kW power consumption. Thanks to the introduction of our proposed dynamic control of 'power' state, the number of powered-up devices can be greatly reduced. Fig. 3 shows the number of powered (*on* and *idle*) devices when the proposed power state control is implemented. At 80 Erlang the powered OEO devices are in average: 154 transponders and 48 regenerators, for a global energy consumption of 74kW. The ratio of unused but powered OEs is 60% for networks operating with current protocols (only *on* state) and 14% when the proposed extension is implemented (*on* and *idle* states). The capability of controlling the OE power state allows 60% of energy savings, due to the 60% reduction on the number of powered OE and the 14% of OE at *idle* state.

If we now take into account the daily traffic variations, current networks are not able to reduce the number of powered devices accordingly; hence no reduction in energy consumption at low traffic loads is experimented while the percentage of unused and powered devices rises up to 80%. With the proposed protocol enhancement, the number of devices in the *up* state follows linearly the traffic load. The number of *idle* devices is independent of traffic changes and depends only on the number of nodes, causing a 473W floor in the energy consumption at very low networks loads. The power efficiency of the proposed solution is estimated considering the daily traffic variations measured in [4] for the considered network, obtaining an average energy consumption of 60kW against 183kW for current network architectures (67% savings), as it is depicted in Fig. 4. Considering 0.1€/kWh energy cost, these savings equals to 10k€ per transported Tb/s per year.

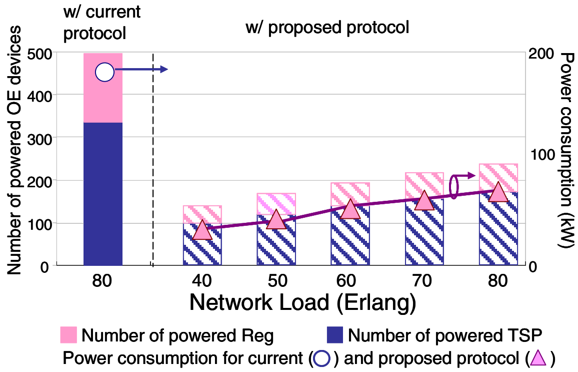

Figure 3: Number of powered devices (left axis) with the current (full) and proposed (hatched) power state management and the power consumption of the two solutions (right axis).

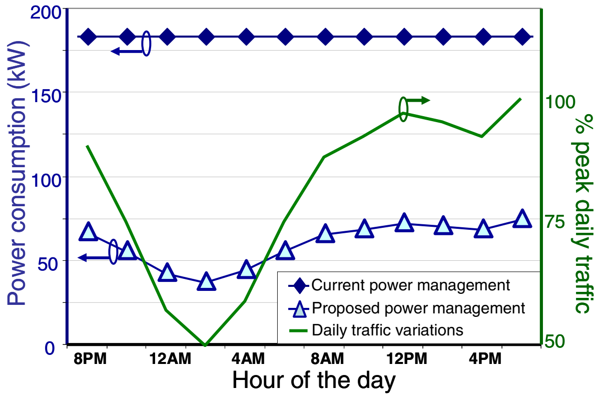

Figure 4: Estimation of power consumption of dynamic networks when daily traffic variations are considered.

6. Conclusions

In this paper we have proposed an enhancement of current GMPLS signaling protocol enabling the dynamic power management of optoelectronic devices in the optical layer. Three power states have been introduced to allow the dynamic on-off switching of optoelectronic devices, including an intermediate *idle* state. We dimensioned a European network with dynamic traffic assumptions. The proposed protocol enhancement enables up to 64% of energy gains. For a network with 100Gb/s connections this amounts to 95kW savings, that is 10k€ per transported Tb/s per year. Future work will focus on the dissemination of power state information through the GMPLS routing protocol (OSPF-TE) to optimize the connections route selection.

This work was partly supported by the French public authorities in the frame of 100GRIA project.

7. References

[1] D.C. Kilper et al, J. Select Topics Quant. Electr., Iss.99, pp.1-10.
[2] www.greentouch.org.
[3] Y. Zang et al., Comm. Surv. and Tut., Vol.12-No.4, pp. 441-458.
[4] http://asert.arbornetworks.com.
[5] http://tools.ietf.org/html/rfc3209.
[6] G. Charlet, Compte Rendus Phy, Vol.9, pp.1013-1030.
[7] OC-192, data sheet available on http://www.jdsu.com.
[8] OIF-CEI-28G-VSR Project – on http://www.ieee802.org.
[9] S. Dave et al., paper JWA14, OFC 2011.
[10] SHF 826H, data sheet available on www.shf.de.
[11] ITLA, data sheet available on http://www.jdsu.com.
[12] R.S. Tucker, J. Select Topics in Quant. Electr., Issue 99, pp. 1-14.
[13] I. Dedic, paper OThT6, OFC 2010.
[14] http://www.ufe.cz/dpt240/cost266/.
[15] O. Rival et al., paper OTuI4, OFC 2011.