Distributed Support of Link Sleep Mode for Energy Efficient GMPLS Networks

Isabella Cerutti, Nicola Sambo, and Piero Castoldi

Scuola Superiore Sant'Anna, Pisa, Italy, {i.cerutti, n.sambo, p.castoldi}@sssup.it

Abstract   *A distributed strategy for selecting the links to be set in sleep mode and re-routing the light-paths is proposed for WDM networks. GMPLS extension for advertising link sleep mode is introduced. Considerable power savings are achieved without affecting the blocking.*

Introduction

In optical WDM networks, a large amount of capacity is deployed, but such capacity lies unused most of the time. Indeed, networks are typically operated at very low levels of utilization, in order to ensure the required level of quality of service and blocking probability. Also, the level of utilization can fluctuate following the periodic changes in the traffic pattern. However, the power dissipated by the optical devices in the network remains almost the same regardless of the network utilization. One example is given by the optical amplifiers on the network links. The power drained by the amplifiers remains the same whether few or a large number of lightpaths traverse the link.

For a scalable and energy-efficient network, it becomes necessary to ensure that the overall power dissipation is proportional to the utilization (or load) while ensuring the required level of performance (e.g., blocking probability). A solution to this issue consists in switching-off, or setting in *sleep mode*, the unused devices[1,2] and minimizing the number of devices in *active mode*. Sleep mode represents a low-power, inactive state from which devices can be suddenly waken-up when required. Although not available yet in most devices for optical networks, support of sleep mode is currently advocated by standardization bodies and governmental programs[3,4].

This paper proposes a distributed strategy for selecting the links, whose devices (such as the amplifiers) need to be set in sleep mode. The lightpaths traversing the selected link are proactively re-routed, to avoid any disruption of the existing lightpaths. The Generalized Multi Protocol Label Switching (GMPLS) control plane is enabled with support of link sleep mode and is evaluated to quantify the achievable power savings and the impact on the network performance, namely the blocking probability, as a function of the newtork load.

Sleep Mode Support in GMPLS Control Plane

Consider a transparent WDM network consisting of N nodes with optical cross-connecting functionalities, and L bi-directional links. Each link supports W wavelengths in each direction and consists of a sequence of link spans, each one terminated by an optical amplifier. Links (i.e., all the devices on the links) are assumed to operate in three different operational modes: *off*, *sleep*, and *active*. Full (working) functionalities are available in active mode, while null, or close to null, functionalities are available in sleep and off modes. Sleep mode permits an automatic transition to active mode (e.g., self-enabled or enabled by the control plane). When *off*, the link is disconnected from both the WDM network and the electrical network. When *sleeping*, the link consumes a negligible amount of power. Events triggering mode transitions are shown in Fig. 1(a).

A distributed GMPLS control plane manages the network using the OSPF-TE routing protocol and the RSVP-TE signaling protocol. OSPF-TE disseminates link state advertisements (LSA) to enable path computation at each node. LSA information are stored in the *traffic engineering database (TED)*. RSVP-TE signaling protocol reserves (releases) the resources for setting up (tearing down) lightpaths.

A distributed strategy for selecting the links to be set in sleep mode is considered and displayed in Fig. 1(b)-(d). The terminating node of a link locally decides the operating mode of the link. If the link supports at most r lightpaths, the link is candidate for being switched to sleep mode. In the figure, node n decides to put to sleep link l (Fig. 1(b)). The node informs the source nodes of the r lightpaths, with RSVP-TE `Notify` message. Upon receiving the `Notify`, the source nodes trigger the lightpath re-routing.

Lightpath re-routing is performed using a make-before-break approach[5]: the old lightpath to be re-routed is still used while a new one is being established (Fig. 1(c)). Once established, the traffic is smoothly re-routed and the source node can tear down the old lightpath. When all the wavelength resources of the selected link are released, the link can be finally put to sleep (Fig. 1(d)). To avoid disruption of the already existing lightpaths, if re-routing is not successful, the old lightpath is preserved and thus the link is forced to operate in active mode.
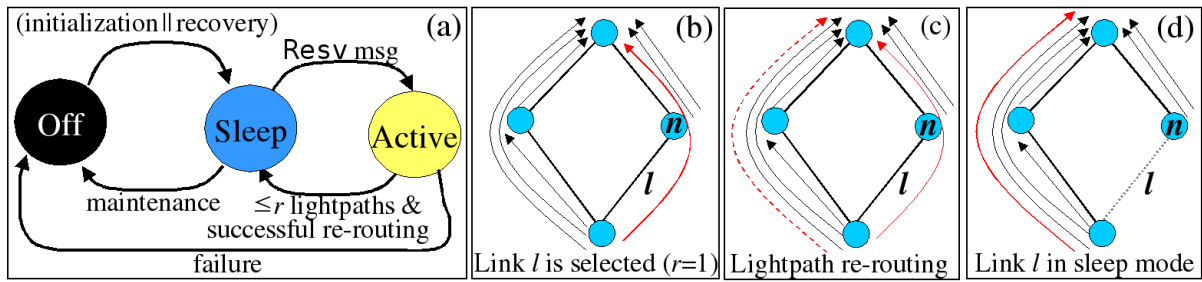
Fig. 1: Mode transition diagram (a) and strategy for lightpath re-routing (b)-(d).

A sleeping link can be reverted to active mode when a newly requested lightpath is routed on such link (e.g., when the node receives a RSVP-TE `Resv` message specifying that the lightpath is reserving resources along the incoming link).

Link sleep mode can be supported by either:

• *sleep mode unaware (SMU)* GMPLS, i.e., the current standard. The LSA messages of OSPF-TE advertise whether a link is active or off. Only nodes receiving a `Notify` message are aware of the link(s) in sleep mode. The other nodes consider the links in sleep mode as active.

• *sleep mode aware (SMA)* GMPLS, i.e., extended for the support of link sleep mode. The LSA messages of OSPF-TE are extended to specify whether a link is in sleep mode. TED is updated accordingly and used for the path computation of incoming lightpath requests. Thanks to this information, all network nodes are informed about the operational mode of the links and can refrain from routing newly requested lightpaths on the sleeping links.

Simulation Results

The performance evaluation is carried out by means of a custom C++ event-driven simulator on the Pan-European network topology[7], with N=27, L=55, W=40. An optical amplifier is placed every 80 km. Link length is equal to the geographical distance. Power drained by each optical device[6] is: transceiver (7 W), optical amplifier (10 W), and OXC (107 mW per internal path). Lightpath requests are uniformly distributed among the source-destination pairs. Both inter-arrival and holding times of the lightpath requests are exponentially distributed with an average of $1/\lambda$ and $1/\mu$, respectively. Network load is, thus, λ/μ.

To balance the traffic, each node randomly selects the routes for the requested lightpaths in the set of shortest routes in terms of hops. In the extended GMPLS, the shortest routes passing on the links in sleep mode are removed from the set. If no route is found, the route at minimum number of links in sleep mode is selected. In both the current and extended GMPLS, re-routing of lightpaths is performed by randomly selecting one of the shortest routes, excluding those passing on the link to be put to sleep. Wavelength selection is first-fit for lightpaths to be set up and is according to a label-preference strategy[7], for the lightpaths to be re-routed. Label-preference strategy maps the wavelengths used by the old lightpath and the new lightpath, in order to minimize the probability of blocking due to contentions during the re-routing.

The performance of SMU-GMPLS and SMA-GMPLS is evaluated in terms of power consumption and lightpath blocking probability and is compared against that of a reference GMPLS network without support of link sleep mode nor re-routing (i.e., links are always active).

Fig. 2 shows the power consumption of SMU-GMPLS (a) and SMA-GMPLS (b) normalized to the case of the reference GMPLS network (i.e., active mode only). Results indicate that the link sleep mode permits to save a considerable amount of power at low loads (e.g., in the range [50,250] Erlang). Indeed, saving of more than 30% can be achieved at 50 Erlang. Moreover, re-routing the lightpaths on scarcely loaded links (i.e., r >0) permits to further reduce the power with respect to the no re-routing case (i.e., r=0), as a larger number of links can be put to sleep. For instance, in a SMU-GMPLS at 150 Erlang, power saving is 5% in absence of re-routing (i.e., r=0) and 12% when up to r=3 lightpaths are re-routed. At higher loads, power savings of SMU-GMPLS and SMA-GMPLS are reduced since links are most likely used by more than three lightpaths and the probability of a succesfull re-routing of all lightpaths is lower.

In the SMA-GMPLS (Fig. 2(b)), power savings are higher than in the SMU-GMPLS (Fig. 2(a)). Indeed, in the SMA-GMPLS, the proposed OSPF-TE extension to advertise the link sleep mode permits to better route the newly requested lightpaths and, at the end, to keep more links in sleep mode. The additional saving is about 7% for r=0 and 5% for r=1, leading to an overall 40% power saving. However, at higher loads, such additional savings are reduced since links are more utilized.
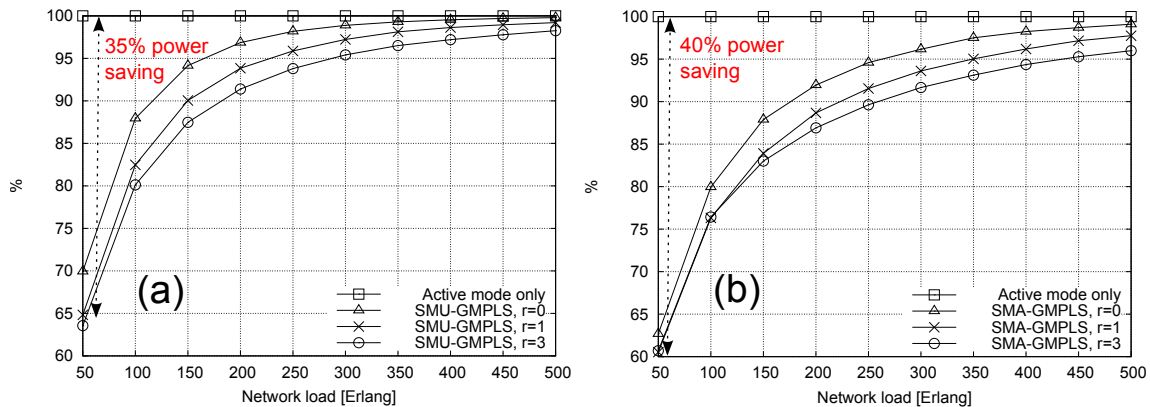
Fig. 3 shows the blocking probability versus

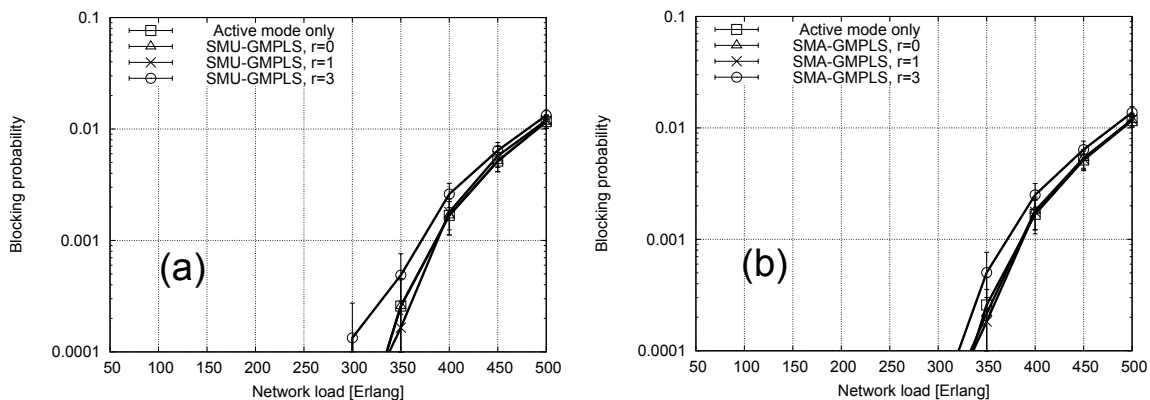**Fig. 2**: Normalized power consumption vs. load for SMU-GMPLS (a) and SMA-GMPLS (b).

Fig. 3: Blocking probability vs. load for SMU-GMPLS (a) and SMA-GMPLS (b).

load, experienced by SMU-GMPLS (a) and SMA-GMPLS (b). Blocking probability is defined as the ratio between the number of blocked requested lightpaths and the number of requested lightpaths. Blocking is experienced if wavelengths are not available or wavelength continuity constraint cannot be satisfied, along the selected route. Interestingly, the blocking probability is not affected by the proposed strategy, especially in the range of [50;300] Erlang, for neither SMU-GMPLS nor SMA-GMPLS. At higher loads, re-routing a single lightpath does not affect blocking. Re-routing up to 3 lightpaths slightly increases the blocking probability, especially in SMU-GMPLS.

Conclusions

A strategy for selecting in a dynamic and distributed way the links to be put to sleep was proposed for a WDM network based on GMPLS control plane. A proactive re-routing of the lightpaths passing on the selected link avoids the disruption of the existing lightpaths and does not impose any tight requirement on the transition time between link operational modes. Considerable power savings are achieved at low loads, leading to an overall power consumption more proportional to the network utilization. More importantly, the network performance (i.e., the blocking of the requested lightpaths) is not influenced by the proposed strategy. This make the sleep mode effective, even with the currently standardized GMPLS suite. By extending GMPLS for avertising link sleep mode, power savings can reach 40%.

References

1 Y. Wu et al., "Power-aware routing and wavelength assignment in optical networks," *Proc. ECOC*, 2009.

2 B. Puype et al., "Energy Efficient Multilayer Traffic Engineering," *Proc. ECOC*, 2009.

3 Energy Star website, www.energystar.gov.

4 Code of Conduct on Energy Consumption of Broadband Equipment, EU Commission, 2008.

5 D. Awduche et al., "RSVP-TE: Extensions to RSVP for LSP Tunnels," RFC 3209, 2001.

6 S. Aleksic, "Analysis of power consumption in future high-capacity network nodes," *IEEE/OSA JOCN*, 1(3), 2009.

7 A. Giorgetti et al., "Label preference schemes for lightpath provisioning and restoration in distributed GMPLS networks," *JLT*, 27(6), 2009.