Experimental Evaluation of Cyclic Sleep with Adaptable Sleep Period Length for PON

Martina Fiammengo[1], Alexander Lindström[1], Paolo Monti[2], Lena Wosinska[2], Björn Skubic[1]
(1) Ericsson Research, Färögatan 6, 164 80 Stockholm, Sweden
(2) Royal Institute of Technology KTH/ICT, Isafjordsgatan 22, Electrum 229, 164 40 Kista, Sweden
bjorn.skubic@ericsson.com

Abstract: An experimental evaluation of cyclic sleep for PON is presented where different control algorithms for adaptable sleep length period are compared.
OCIS codes: (060.2330) Fiber optics communication; (060.4250) Networks

1. Introduction

Reduction of energy consumption in the ICT sector has become an increasingly important issue, due to a growing concern for both global warming and network costs related to energy consumption. In comparison to the metro and core segments, access constitutes a substantial part of the per-subscriber network energy consumption. One of the main research trends today aiming at reducing network energy consumption is introduction of sleep mode techniques where network elements/subsystems are powered off during periods of reduced network load. These techniques are particularly appealing for the access segment, which is characterized by low resource utilization. The combination of large contribution to the overall network power consumption and low resource utilization implies large potential for exploiting sleep mode techniques in access networks. Passive optical networks (PONs) provide a powerful solution for fixed access from the provider central office to the customer. A PON consists of optical network units (ONUs) located at the customer premises, which are connected via a fully passive optical distribution network to an optical line terminal (OLT) at the provider site. The passive network infrastructure makes PON to an attractive candidate for the energy efficient access network solution. Energy efficiency of PON can be further improved by exploiting sleep mode techniques. In this paper we investigate and experimentally evaluate cyclic sleep with adaptable sleep period length for PON.

2. Low power modes in PON

Different low power modes within the context of PON have been described in [1]. Of particular research interest are techniques that involve powering off the ONU transceiver. One major challenge concerns how to arrange timely ONU wake-up with respect to QoS requirements. In order to accommodate OLT-initiated wake-up of the ONU, the sleeping ONU transceiver must be periodically powered on in order to intercept potential wake-up messages from the OLT. This cyclic state of periodic transitions between *sleep* and *awake* is often referred to as *cyclic sleep*. An alternative to cyclic sleep is to power off only the transmitter part and to leave the receiver operational at all times (*doze* mode). Although the residual power dissipation in the doze state is larger than in the sleep state, a complete comparison of the two modes may eventually favor doze if the ONU can spend more time dozing than sleeping [2]. As mentioned previously, the sleep state requires frequent periodic wake-ups, which limits the time that can be spent sleeping. Furthermore, the wake-up time from the doze state is significantly shorter than the wake-up time from the sleep state, implying that doze states can be scheduled at finer granularity. This calls for a better understanding of the power saving potential of the different implementation variants of cyclic sleep.

3. Control intelligence

An important part of cyclic sleep operations is the control intelligence which is where large performance differences can be expected from different implementations. The control intelligence consists of the criteria at the ONU and OLT for initiating transitions or accepting/rejecting transition requests to different sleep mode states. The control intelligence may also consist of control for the sleep period length. The criteria for the state transitions could consist in thresholds for traffic load, buffer occupancy or packet inter-arrival time (IAT). The length of the sleep periods could be adjusted based on several different approaches. One approach is to monitor active services and to set the cyclic sleep parameters according to the minimum requirements for active services. As a minimum requirement, the cyclic sleep mode must always support initialization requirements for any of the services over the PON. A second approach which is evaluated in this paper, consists in adapting the sleep length according to the traffic patterns without any information about active services. By means of this approach the sleep mode could be optimized with respect to power savings or latency degradation, compared to the case with static or semi-static sleep periods.

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

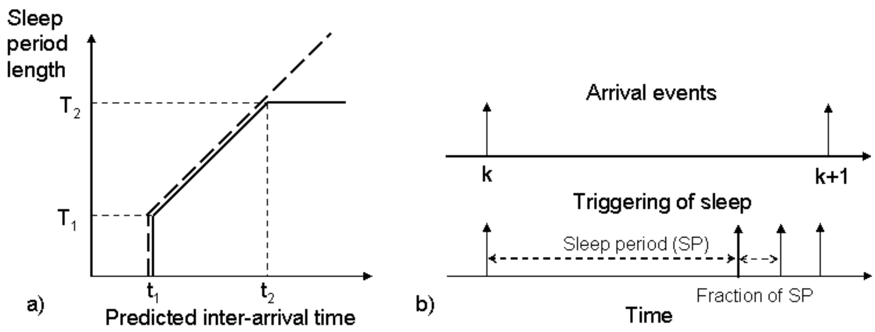

Fig 1. Two versions of adaptive sleep period control where (a) shows the sleep period as a piece-wise linear function of the measured inter-arrival time (e.g. ESA or X) as used in schemes A, C (solid line) and B, D (dashed line) and where (b) shows the incremental approach used in schemes B, D.

Fundamental for the effectiveness of this approach is the scheme used to predict the arrival patterns of new traffic based on the information from the past periods. In this work we collect past packet inter-arrival time information to predict future arrival patterns. Hence, the scheme is based on monitoring packet inter-arrival time and adjusting the sleep length accordingly. Two different ways of collecting inter-arrival time history were considered, i.e. an arithmetic average (AA) here consisting of eight samples and the exponentially smoothed average (ESA) with adaptable smoothing parameter (see [3] for details) .

For each of these two approaches for predicting the inter-arrival time, two ways of choosing the sleep period length have been considered (Fig. 1). The first approach (here referred to as *conventional*) consists of a simple function with two thresholds for a minimum (6 ms) and a maximum (10 ms) sleep length and a linear region in between where the sleep period is set to the predicted inter-arrival time. The second approach (here referred to as *incremental*) [3] consists of a similar function but without the upper threshold to the sleep period and where an incremental scheme for prolonging the sleep period with fractions (here 1/16) of the predicted inter-arrival time has been introduced. Combining all these approaches gives four different adaptable sleep mode schemes that have been evaluated: Scheme A (ESA + conventional approach for sleep period length), Scheme B (ESA + incremental approach for sleep period length), Scheme C (AA + conventional approach for sleep period length), and Scheme D (AA + incremental approach for sleep period length). Details of the signaling scheme used in the experimental setup for supporting the control intelligence is shown in Fig. 2(b).

4. Experimental Setup

The experimental setup for the evaluation of cyclic sleep is shown in Fig. 2(a). Traffic is generated by the Ixia XM12 High Performance Chassis, using a port operating at 1Gbps and transported to an SFP board which converts the received traffic into a differential signal compliant with the 1.25 Gbps serial interface on the Altera board (Transceiver Signal Integrity Development Kit, Stratix IV GT Edition). The Altera FPGA emulates sleep mode operation at both the ONU and OLT. Traffic is returned to the Ixia Chassis for latency measurements and traffic statistics. Several approximations have been made during the experiments: (1) the upstream channel is not included

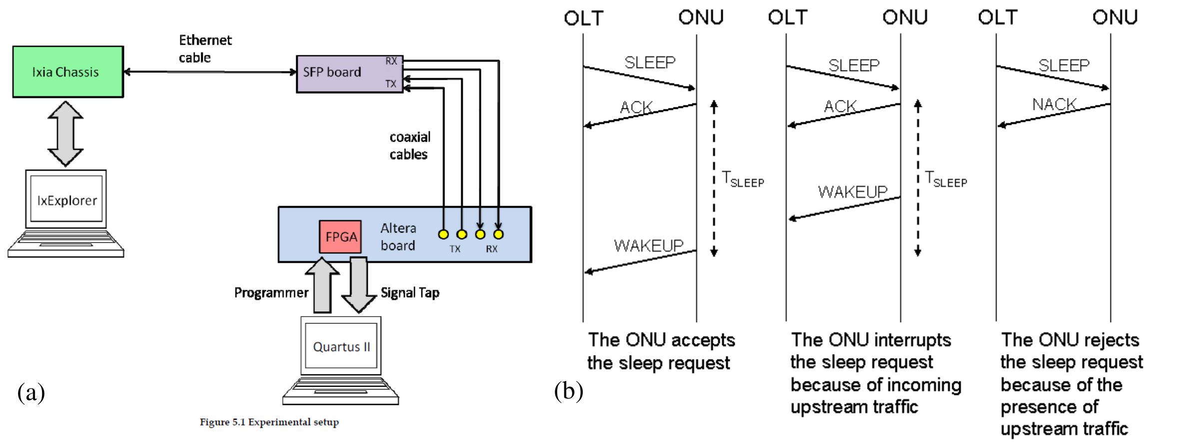

Fig. 2 (a) Schematic figure of the experimental setup and (b) an illustration of the emulated sleep control signaling scheme.

in the emulation and thus the interruption mechanism of the sleep period performed by the ONUs is not included in the evaluation, (2) sleep requests are triggered and the sleep periods are computed only on the basis of downstream traffic, (3) propagation delay in the communication between the ONU and OLT is neglected and (4) the wake-up time is neglected in the implementation although the effect of a 1ms wake-up has been accounted for while evaluating the power savings. Power is estimated by assuming 10W for the entire ONU in full operation and 2W in sleep mode.

5. Results

Several different traffic patterns were considered in the experimental evaluation. Fig. 3 presents results on estimated ONU power consumption, average additional downstream (DS) latency and max additional DS latency. Results refer to the cyclic sleep control schemes A-D, presented in Section 3. Note that traffic load in all cases is very low. Hence, the limits that are being investigated are due to the actual traffic patterns rather than the load. In general, we note that results are quite sensitive to the patterns and that they at first sight can even be quite counter intuitive, which poses certain problems for the implementation.

Two examples of constant rate traffic are presented, i.e. 20 ms and 300 ms IAT. For constant rate traffic there should be no difference between the two IAT prediction schemes. For IAT larger than the upper threshold of schemes A and C, we see an advantage with the incremental approach (although this also depends on the size of the increments). Two examples of bursty traffic are presented, i.e. bursts consisting of 25 and 100 frames respectively, with 1ms inter-frame intervals and 1s inter-burst intervals. These traffic patterns result in overall poor sleep mode performance (either in terms of power savings or latency), since they expose the difficulties of traffic arrival prediction. A number of more complex traffic sequences were also considered such as patterns alternating between two rates (10x1ms intervals + 20x500ms intervals and 20x10ms intervals + 20x30ms) and nearly random patterns consisting of a superposition of multiple carefully chosen constant rate traffic streams. For these complex traffic patterns we see an average power saving of around 60-70%. We further note that schemes B and D exhibit some latency issues due to the absence of an upper bound to the sleep period length.

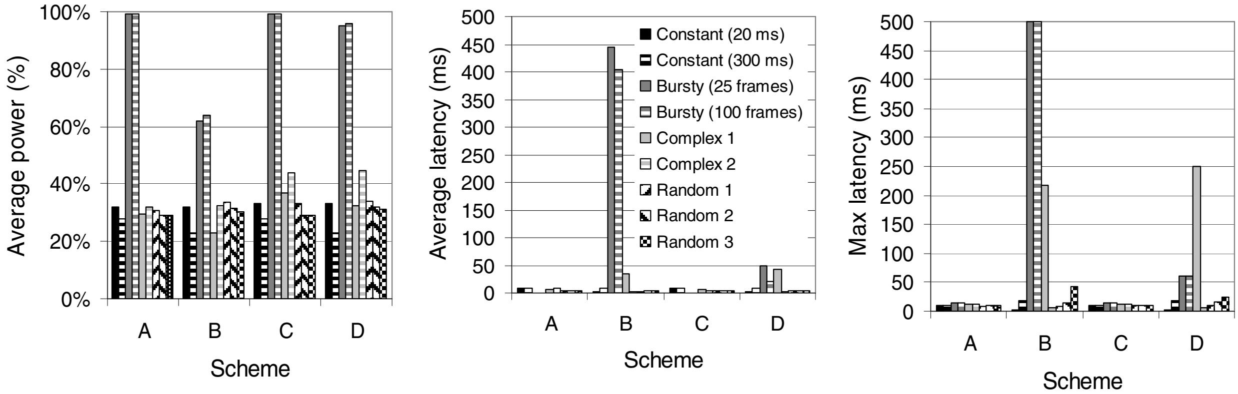

Fig. 3. Evaluation of schemes A-D where (a) shows estimated average power consumption (percentage of full ONU power consumption) derived from measuring time in the sleep state, (b) average additional latency incurred by the sleep mode, and (c) max additional latency incurred by the sleep mode.

6. Conclusions

Results show that performance of the evaluated schemes is sensitive to traffic properties. Bursty traffic exposes the difficulties of cyclic sleep mode schemes with adaptive sleep periods based on traffic prediction. However, for most of the considered traffic patterns, schemes can be constructed that result in an average energy saving of 60-70% with minor latency degradation (using the assumed power model).

References

[1] ITU-T Recommendation G.Sup45, GPON power conservation, 2009.
[2] B. Skubic, D. Hood, “Evaluation of ONU Power Saving Modes for Gigabit-Capable Passive Optical Networks,” IEEE Commun. Mag., vol. 25 no. 2, March 2011.
[3] M. Fiammengo, “Sleep mode scheduling technique for energy saving in TDM-PONs,” MSc thesis, TRITA-ICT-EX-2011:17, KTH, 2011.
[4] R. Kubo, J. Kani, H. Ujikawa, T. Sakamoto, N. Yoshimoto, H. Hadama, "Hybrid Power-saving Mechanism with Cyclic Sleep and Adaptive Link Rate Functions for 10G-EPON Systems," Electronics Letters, Vol.46, No.13, June 2010.
[5] R. Kubo, J. Kani, Y. Fujimoto, N. Yoshimoto, K. Kumozaki, "Proposal and Performance Analysis of a Power-Saving Mechanism for 10 Gigabit Class Passive Optical Network Systems," NOC/OC&I 2009 Proceedings, Valladolid, June 2009.