

Analyzing Local Strategies for Energy-Efficient Networking

Sergio Ricciardi¹, Davide Careglio¹, Ugo Fiore², Francesco Palmieri³,
Germán Santos-Boada¹, and Josep Solé-Pareta¹

¹ CCABA, Universitat Politècnica de Catalunya, Barcelona, Spain

² CSI, Università degli Studi di Napoli Federico II, Napoli, Italy

³ DII, Seconda Università degli Studi di Napoli, Aversa, Italy

{sergiior, careglio, german, pareta}@ac.upc.edu,
{fpalmier, ufiore}@unina.it

Abstract. Power management strategies that allow network infrastructures to achieve advanced functionalities with limited energy budget are expected to induce significant cost savings and positive effects on the environment, reducing Green House Gases (GHG) emissions. Power consumption can be drastically reduced on individual network elements by temporarily switching off or downclocking unloaded interfaces and line cards. At the state-of-the-art, Adaptive Link Rate (ALR) and Low Power Idle (LPI) are the most effective local-level techniques for lowering power demands during low utilization periods. In this paper, by modeling and analyzing in detail the aforementioned local strategies, we point out that the energy consumption does not depend on the data being transmitted but only depends on the interface link rate, and hence is throughput-independent. In particular, faster interfaces require lower energy per bit than slower interfaces, although, with ALR, slower interfaces require less energy per throughput than faster interfaces. We also note that for current technologies the energy/bit is the same both at 1 Gbps and 10 Gbps, meaning that the increase in the link rate has not been compensated at the same pace by a decrease in the energy consumption.

Keywords: sleep mode, energy-efficiency, power consumption, low power idle, adaptive link rate.

1 Introduction

Most of the currently known non-renewable primary energy sources are becoming scarcer and will get exhausted in only some decades. On the other hand, they are highly polluting as their burning process emits large quantity of GHGs causing climate changes and global warming phenomena. The current growth scenario is not sustainable, and international initiatives are trying to decrease the energy consumption and the GHG emissions by 20% for 2020 [1]. In order to achieve such drastic reductions, it is necessary to adopt a radical change in the current lifestyle and business as usual model. For such a transformation, the key factor is the use of energy-efficient processes together with energy-aware solutions and policies that

increasingly exploit renewable energy sources in providing all the public utility services. In the networking scenario, miniaturization and ICT growing dynamics, effectively described by the Moore's and Gilder's laws [2][3], have not had the expected counterpart in power consumption reduction. Miniaturization has reduced unit-power consumption but has allowed more logic ports to be put into the same space, thus increasing performances and, concomitantly, power utilization (a phenomenon called rebound effect already known as Jevons paradox [4]). As a consequence, the total power required per node is growing faster and faster. Nevertheless, traffic dynamics often result in a significantly different network usage which presents peaks alternated by low load periods, making room for power management techniques that, while satisfying the users' demand, exploit the low load periods for saving as much energy (and, thus, money) as possible. Accordingly, adaptive power management strategies that can be implemented independently and at different levels of granularity on each network device can be introduced at the local equipment-level to decrease power consumption in the operational phase and bring positive effects for the environment and significant cost savings. Power consumption can be drastically reduced by temporarily switching off or downclocking unloaded interfaces. In this work, such local energy containment strategies have been properly modeled and their behavior analyzed through simulations with the goal of better understanding their operating dynamics, strengths and weaknesses.

2 Active Local Strategies for Network Energy Efficiency

Experimental measurements collected from several network devices [5] show that in current architectures half of the energy consumption is associated to the base system and the other half to the number of installed line interface cards (even if *idle*). Furthermore, the power consumption of the actual electronic routing/switching matrix and line cards is, quite surprisingly, almost independent from the network load, so that the energy demand of heavily loaded devices is only about 3% greater than that of idle ones. These results suggest that it is necessary to develop energy-efficient architectures exploiting the ability of temporarily switching off or putting into energy saving mode devices or subsystems (e.g. switching fabrics, line cards, I/O ports, etc.) in order to minimize energy consumption whenever possible. Putting entire nodes into sleep mode (*per node sleep mode*) may be unpractical, especially for large and highly connected ones, since many very expensive transmission links become unused, hence negating significant capital investments (CAPEX) for the entire duration of the sleep interval. Furthermore, per node sleep mode drastically reduces the overall meshing degree, by limiting the network reliability and partially negates the possibility of balancing the load on multiple available links/paths. On the other hand, putting into sleep mode only single interfaces (*per interface sleep mode*) may introduce considerable energy savings in particular when operating at high speeds, since, for example, in a commercial off-the-shelf (COTS) Ethernet switch (Catalyst 2970 24-port LAN switch) a 1000baseT interface adds about 1.8 W to the overall consumption [6] (Table 1). Per-interface sleeping mechanisms (ALR and LPI) have been identified [6][7] as viable and effective solutions. In ALR, the ability to dynamically modify the link rate according to the real traffic needs is used as a technique to reduce the power

consumption. Operating a device at a lower frequency can enable reductions in the energy consumption and also allows the use of dynamic voltage scaling (DVS) for reducing the operating voltage. This allows power to scale cubically and hence energy consumption quadratically with operating frequency [8].

Table 1. COTS switch power consumption with varying number of interfaces

# Active Interfaces	10BaseT	100baseTX	1000baseT
0	69.1 W	69.1 W	69.1 W
2	70.2 W	70.1 W	72.9 W
4	71.1 W	70.0 W	76.7 W
6	71.6 W	71.1 W	80.2 W
8	71.9 W	71.9W	83.7 W

For example, the Intel 82541PI Gigabit Ethernet Controller draining about 1 W at 1 Gbps full operation is able to support a smart power down feature by turning off PHY if no signal is present on link and drops the link rate to 10 Mbps when a reduction of energy consumption is required [9]. Also in the last mile, the ADSL2 standard (ITU G.992.3, G.922.4, G.992.5) is able to support multiple data rates corresponding to different link states (L0: full rate, L2: reduced rate, L3: link off) for power management sake [10]. In LPI, transmission on single interface is stopped when there is no data to send and quickly resumed when new packets arrive, in contrast with the continuous IDLE signal used in legacy systems. LPI defines large periods over which no signal is transmitted and small periods during which a signal is transmitted to synchronize the receiver. When operating in low-power mode, the elements in the receiver can be frozen, and then awakened within a few microseconds, as reported in Table 2 [11].

Table 2. Common wake-on-arrival strategy parameters for different interfaces technologies

Technology	Wakeup Time	Sleep Time	Average Power savings
100baseTX	30 μ s	100 μ s	90%
1000baseT	16 μ s	182 μ s	90%
10GbaseT	4.16 μ s	2.88 μ s	90%

Significant energy savings can be obtained when the involved devices spend a considerable fraction of their time in the low power mode. Although the savings vary from device to device, the energy consumption, when the device is in low power mode, can be as low as 10% that the one in active mode. During the transitions back and forth from low power mode there is a considerable increase in energy consumption as many elements in the transceiver have to be active. The actual value will depend on the implementation and possibly ranges from 50% to 100% of the active mode energy consumption. In network environments where packet arrival rates can be highly non-uniform, allowing interface transitions between different operating rates or sleep/active modes can introduce additional packet delay, or even loss, due to the associated transition times. The main issues to be addressed are the coordination

among nodes during the transitions from and to a low power consumption state or from a transmission rate to another one. In line of principle, these transitions should be kept as transparent as possible to upper layer protocols and applications. Several solutions can usefully exploit the tradeoff between potential energy savings, performance and transparency. For example, buffering, packet coalescing and coordinated Ethernet strategies may be introduced, to collect packets into small bursts and thereby creating gaps long enough to profitably sleep [6][12][13]. Potential concerns are that buffering will add too much delay across the network and that traffic burstiness will exacerbate the loss.

3 Modeling and Analyzing Local Energy Containment Methods

In this section, we present a model of the aforementioned local techniques built by interpolating realistic data obtained from the available literature and experimental measurement on available state-of-the-art hardware. We exploited and analyzed through simulation some of the most interesting properties and operational features of these techniques when applied to individual non-cooperating network devices. Let $G(V,E)$ be a directed graph representing the physical network topology; V the set of vertices that represent the network nodes and E the set of edges that represent the network links. Note that, as a (unidirectional) link is attached to each interface, the set of links E actually coincides with the set of interfaces. Each interface has its own native speed: $\forall i \in E, v_i \in R = \{10 \text{ Mbps}, 100 \text{ Mbps}, 1000 \text{ Mbps}, 10000 \text{ Mbps}\}$ represents the *native link rate* of interface i . The energy/power consumption of interfaces working at their native link rates [6][14] are illustrated in Table 3.

Table 3. Energy and power consumption of interfaces working at native speeds

Native link rate v_i	Power per interface	Energy Scaling Index (ESI) - Energy per bit	Energy Consumption Rate (ECR) - Power per Gbps
10 Mbps	0.1 W	10 nJ/bit	10 W/Gbps
100 Mbps	0.2 W	2 nJ/bit	2 W/Gbps
1,000 Mbps	0.5 W	0.5 nJ/bit	0.5 W/Gbps
10,000 Mbps	5.0 W	0.5 nJ/bit	0.5 W/Gbps

ESI and ECR are different energy/power consumption metrics that may be reduced to equivalent values, in fact it holds that: $W/Gbps = (J/s) / (Gbit/s) = J/Gbit = nJ/bit$.

Scaling the energy consumption per bit (ESI metric) reveals that the energy consumption for forwarding one bit is not the same for every interface but depends on its native link rate. In particular, the *energy per bit* is lower for faster interfaces, meaning that forwarding one bit on a slower interface requires more energy than on a faster one (besides occupying the link resource for a longer time). We also note how the energy/bit ratio is the same both at 1 Gbps and 10 Gbps, that is, there is no gain in the energy/bit at 10 Gbps (as instead occurs when switching between 10/100 Mbps and between 100/1000 Mbps). This behavior is due to the current 10 Gbps technology, whose increase in the link rate (achieved through advanced modulation techniques [15]) has not been compensated at the same pace by a decrease in the

energy consumption. As a result, 10 Gbps interfaces consume 10 times more energy than 1 Gbps ones, i.e. the power consumption scales linearly from 1 to 10 Gbps. Consequently, the best balance between power consumption and bit rate is reached at 1 Gbps (see Fig. 1). This situation is further stressed when the throughput is not equal to the link rate, which corresponds to an underutilized channel. Our observations confirm that an interface consumes the same power whatever its *current* throughput is: power consumption is *throughput-independent*. For this reason, the link rate can be adapted to the current throughput by using ALR with consequent energy savings. However, the IEEE Energy Efficient Ethernet working group, when analyzing the opportunity to adopt ALR or LPI in the 802.3az standard, decided in favor of LPI [16] since the two strategies have been considered as alternative to be included in the standard. Instead, we evaluate the advantages offered by a combination of them and advocate the complementary use of the two strategies.

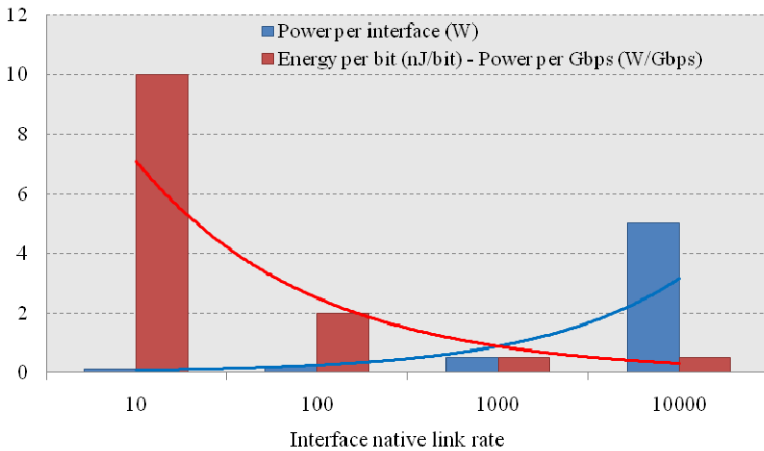


Fig. 1. Energy and power consumption of different interfaces working at their native speeds

To model the ALR, let us consider the set $R=\{r_1, \dots, r_m\}$ set of available link rates; $\forall i \in E, r_i \in R$ represents the *working link rate*, that is the link rate at which the interface i is currently operating; obviously, it holds that $\forall i \in E, r_i \leq v_i$. In the actual standard an interface may switch only to the set of existing link rates, thus $R=\{10 \text{ Mbps}, 100 \text{ Mbps}, 1000 \text{ Mbps}, 10000 \text{ Mbps}\}$. According to current technologies, we consider three possible operating modes for interfaces: (1) Off, occurs when the interface is *down*; (2) LPI, interface up in low power mode and there is no data to transmit; (3) ALR, there is data to transmit and the interface is up at working rate $r_i \in R$. Let us consider an interface i with native link rate v_i and a constant data throughput t_d ; then, with ALR, the interface will switch its current link rate to r_i ; $r_{i-1} < t_d \leq r_i$. In our simulations we found that, when using the ALR, the power consumption of an interface i depends not only on the working link rate r_i but also on the native link rate v_i . In other words, transmitting a fixed data throughput t_d has different power consumption depending on the interface native link rate v_i ; in this case, slower

interfaces consume less power than faster ones for the same throughput t_d , even if they work at the same rate r_i . This result, quite surprising if we consider that slower interfaces consume more energy per bit than faster ones, may be explained considering that the different technologies adopted for reaching higher link rates [15] lead to greater fixed power consumption for faster interfaces. In fact, as routers, also the interfaces have fixed and variable power consumption. The fixed part is always present just for the interface to stay up and accounts for the control circuits, while the variable traffic-proportional power consumption is due to the transceivers. In the following we model such energy consumptions and show a breakdown of the different energy components in a 10 Gbps interface.

Table 4. Power consumptions of interfaces working at different rates in $\{10,10^2,10^3,10^4\}$ Mbps

v_i	r_i	Mbps	Power consumption
$\forall v_i \in R$	Off / r_0	0/0	$\Psi(v_i, \text{Off} / r_0) \cong 0^*$
$v_1: 10$	r_1	10	$\Psi(v_1, r_1)$
$v_2: 10^2$	r_1/r_2	10/10 ²	$\Psi(v_2, r_1) / \Psi(v_2, r_2)$
$v_3: 10^3$	$r_1/r_2/r_3$	10/10 ² /10 ³	$\Psi(v_3, r_1) / \Psi(v_3, r_2) / \Psi(v_3, r_3)$
$v_4: 10^4$	$r_1/r_2/r_3/r_4$	10/10 ² /10 ³ /10 ⁴	$\Psi(v_4, r_1) / \Psi(v_4, r_2) / \Psi(v_4, r_3) / \Psi(v_4, r_4)$

* In LPI, the device only sends signals during short refresh intervals and stays quite during large intervals so the power consumption in the LPI mode is almost 0.

In general, to model the fixed and the variable power consumption, we define $\{\Psi(v_i, r_j) \mid j = 1, 2, \dots, m\}$ where $\Psi(v_i, r_j)$ is the power consumption of the interface $i \in E$ with native speed $v_i \in R$ operating at link rate $r_j \in R$ and $\Psi(v_i, r_j) < \Psi(v_i, r_k) \forall j < k$. Also, we define Θ_n as the fixed power consumption of node $n \in V$ accounting for its base system, switching matrix, control circuits, etc. Note that Θ_n does not include the power consumption of the node interfaces, which is given by the Ψ term. Let's see how the term Ψ models the power consumption of the interfaces. Let $i, j \in E$ be two interfaces with respectively native and working link rates (v_i, r_i) and (v_j, r_j) . We can observe that Ψ can be characterized from the following properties:

$$\Psi(v_i, r_i) : R \times R \rightarrow \Re \quad (v_i, r_i) \mapsto \Psi(v_i, r_i) \in \Re \tag{1}$$

$$\Psi(v_i, r_i) \propto v_i, \forall i \in E \tag{2}$$

$$\Psi(v_i, r_i) \propto r_i, \forall i \in E \tag{3}$$

Where (1) is the functional definition of Ψ , (2) and (3) state the proportionality of Ψ to the native and working link rates respectively;

$$\forall i \in E : (r_i = \text{off} \vee r_i = r_0) \Leftrightarrow \Psi(v_i, r_i) \cong 0 \tag{4}$$

the energy consumption of an interface Off or in LPI is nearly 0;

$$\forall i, j \in E : (v_i < v_j \wedge r_i = r_j) \Rightarrow \Psi(v_i, r_i) < \Psi(v_j, r_j) \tag{5}$$

two interfaces with the same working link rate but different native link rates have different power consumptions;

$$\forall i, j \in E : (v_i = v_j \wedge r_i < r_j) \Rightarrow \Psi(v_i, r_i) < \Psi(v_j, r_j) \quad (6)$$

two interfaces with the same native link rate but different working link rates have different power consumptions;

$$\forall i, j \in E : \left(\frac{r_i}{v_i} = \frac{r_j}{v_j} \wedge (r_i \neq r_j \vee v_i \neq v_j) \right) \Rightarrow \Psi(v_i, r_i) \neq \Psi(v_j, r_j) \cdot \quad (7)$$

two interfaces with the same r/v ratio, but with different native/working link rates, have different energy consumptions. In order to model the interfaces power consumption in a realistic case, we first consider a two-level system, that is a system in which interfaces may work only at two link rates: high and low power modes. The total energy consumption is therefore given by the sum of the energy cost spent in low power mode plus the cost in high mode. That is, the sum of the low-power mode instantaneous cost times the total time spent in low-power mode plus the high-power mode instantaneous cost times the total time spent in high-power mode. In the general case, when more than one link rate is available, we divide the time in intervals so that the link rate stays constant during each interval and record the duration of each interval. We indicate as N the number of time intervals with unchanging state so that there are $N-1$ link rate transitions with changing states; if t_i is the duration of the i -th time interval (in seconds) then the total time considered is given by:

$$T = \sum_{i=1}^N t_i. \quad (8)$$

Let r_i be the link rate in the i -th time interval; τ the time needed for the link rate transition (assume that every transition requires the same time); c_j the instantaneous power consumption at link rate j ; ζ a proportionality constant between the instantaneous power demand c_j and the corresponding link rate j so that $c_j = \Theta + \zeta \cdot j$; X_{hk} the power consumption when transitioning from link rate h to link rate k . Then, for a single interface:

$$\begin{aligned} \Psi &= \sum_{i=1}^N c_{r_i} t_i + \sum_{i=2}^N \tau X_{r_{i-1}r_i} = \Theta \sum_{i=1}^N t_i + \zeta \sum_{i=1}^N r_i t_i + \tau \sum_{i=2}^N X_{r_{i-1}r_i} = \\ &\Theta T + \zeta \cdot \bar{r} \cdot T + \tau \sum_{i=2}^N X_{r_{i-1}r_i} = (\Theta + \zeta \cdot \bar{r})T + \tau (N-1)\bar{X}. \end{aligned} \quad (9)$$

For sake of simplicity, we consider all interfaces to behave in the same way. In the realistic hypothesis [6][17] that $X_{hk} \propto c_{\max\{h,k\}} = \zeta \cdot \max\{h,k\}$, $\bar{X} \propto \bar{r}$ and eq. (9) becomes:

$$\Psi \approx (\Theta + \bar{r}\zeta)(T + N\tau). \quad (10)$$

Starting from the above energy model, combined with several real energy consumption observations available in literature [6][9][14][17][18], we simulated

some native speed interfaces working at different link rates. The associated per bit energy consumption values are shown in the chart of Fig. 2. We can see how the energy per bit depends both from the native link rate of the interfaces and on their actual working link rate. Furthermore, we can notice how the energy consumption of native high-speed interfaces does not vary much when switching to lower link rates, whilst the energy consumption of native low-speed interfaces is highly variable, especially when working at low rates.

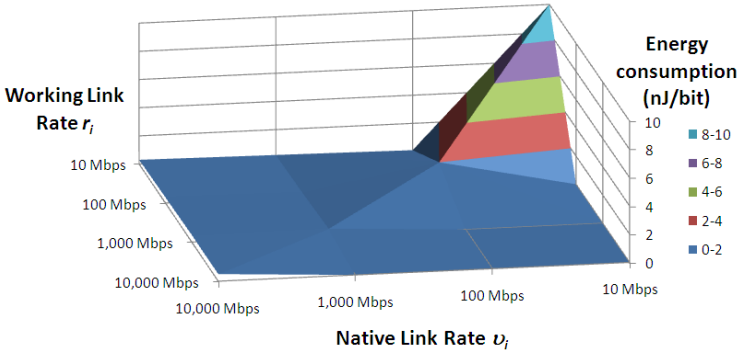


Fig. 2. Energy per bit for native interfaces operating at different link rates (interpolation obtained by putting not defined values to 0)

In Fig. 3 we plotted the power consumption breakdown for a simulated routing device modeled with interfaces at 10 Gbps. The base systems accounts for approximately 50% of the total energy consumption while the interfaces (fixed and variable parts) accounts for the other half.




Fig. 3. Power consumption breakdown for a router with interfaces at native link rate $v = 10$ Gbps and working link rates $r_i = 10^i$ Mbps

As we can see, the larger interfaces energy consumption is due to the fixed part that is independent from the link rate, while only the remaining 15% of energy is due

to the rate-proportional energy consumption. This scenario suggests that between LPI and ALR it is preferable to use LPI: when there is no data to transmit, ALR sends the continuous idle signal (which is everything but *idle*), whilst LPI enters the lower power consumption mode, thus consuming lower energy; furthermore, the auto-negotiation mechanism of ALR may be in the order of *ms*, while the sleep/wake-up transitions of LPI only requires some μs for 10 Gbps interfaces [17][19] and, consequently, lower delay; finally, transmitting the same amount of data in ALR takes longer time than LPI, since in LPI the transmission is realized at the maximum rate, instead in ALR the transmissions is lowered to best fit the throughput, thus occupying the resources for longer time. The Energy-Efficient Ethernet working group has recently adopted the LPI as power management solution for IEEE 802.3az [16]. Furthermore, the difference in energy consumption between interfaces with different native link rates suggests also the possibility for and advanced ALR with circuit over-provisioning, i.e. a network interface may be provisioned with different circuits, say a low and a high speed one, and may switch between one or other according to the required data throughput. This solution, at the expense of increased capital expenditures (CAPEX) for the additional hardware, may lead to decreased operational costs (OPEX) due to the lower fixed power consumption of slow circuits. Finally, also another possibility is given by the heterogeneity of the equipment in a network. In fact, such heterogeneity may be exploited by a global load-balancing schema, implemented as part of a routing and wavelength assignment algorithm, which tries to distribute the connection requests in such a way that the overall network energy consumption is minimized. A best-fit allocation scheme may be implemented in order to close match the bandwidth demands with the interfaces native link rates, so that the fixed power consumption cost is amortized and the maximum efficiency is reached.

4 Conclusions

By modeling and analyzing the available local energy efficiency strategies, we observed that faster interfaces consume less energy per bit than slower ones, but also that, lowering the interfaces working link rates during low utilization periods, does not lead to the same power savings for all interfaces but depends on the interface *native* link rate. Native slower interfaces (e.g., 100 Mbps) consume less power than native faster interfaces (e.g., 1,000 Mbps) for transmitting the same throughput (e.g., 80 Mbps) due to the higher fixed power consumption that comes with faster interfaces. In other words, while the *energy-per-bit* is lower for faster interfaces, with the ALR the *energy-per-throughput* is lower for slower interfaces. Furthermore, we observed that the energy consumption of native high-speed interfaces does not vary much when switching to lower link rates, whilst the energy consumption of native low-speed interfaces is highly variable, especially when working at low rates. Finally, we point out that the different fixed and variable power consumptions of interfaces may be exploited by circuit over-provisioning techniques as well as load balancing schemes for minimizing the overall energy consumption and, thus, network operational costs.

Acknowledgments. This work was supported in part by the COST Action IC0804 on Energy Efficiency in Large Scale Distributed Systems, the Spanish Ministry of

Science and Innovation under the DOMINO project (TEC2010-18522), the Catalan Government under the contract SGR 1140 and the DIUE/ESF under the grant FI-201000740.

References

- [1] EU Spring Summit, Brussels (March 2007)
- [2] Moore, G.E.: Cramming more components onto integrated circuits. *Electronics* 38(8) (April 19, 1965)
- [3] Gilder, G.F.: *Telecosm: How Infinite Bandwidth Will Revolutionize Our World*. The Free Press, NY (2000)
- [4] Jevons, W.S.: *The Coal Question; An Inquiry concerning the Progress of the Nation, and the Probable Exhaustion of our Coalmines*. Macmillan and Co., Basingstoke (1866)
- [5] Chabarek, J., Sommers, J., Barford, P., Estan, C., Tsiang, D., Wright, S.: Power awareness in network design and routing. In: *Proc. IEEE INFOCOM* (2008)
- [6] Christensen, K., Nordman, B.: Reducing the Energy Consumption of Networked Devices. In: *IEEE 802.3 tutorial* (July 19, 2005)
- [7] Hays, R.: Active/Idle Toggling with 0BASE-x for Energy Efficient Ethernet. Presentation to IEEE 802.3az Task Force (November 2007)
- [8] Zhai, B., Blaauw, D., et al.: Theoretical and Practical Limits of Dynamic Voltage Scaling. In: *DAC 2004* (2004)
- [9] Intel® 82541PI Gigabit Ethernet Controller, Intel White Paper, <http://www.intel.com/design/network/products/lan/controllers/82541pi.htm>
- [10] Tzannes, M.: ADSL2 Helps Slash Power in Broadband Designs. *CommDesign.com* (January 30, 2003)
- [11] Reviriego, P., et al.: Performance Evaluation of Energy Efficient Ethernet. *IEEE Commun. Letters* 13(9) (September 9, 2009)
- [12] Kubo, R., Kani, J., Fujimoto, Y., Yoshimoto, N., Kumozaki, K.: Sleep and adaptive link rate control for power saving in 10GEPON systems. In: *Proc. Globecom* (2009)
- [13] Nedeveschi, S., Popa, L., Iannaccone, G., Ratnasamy, S., Wetherall, D.: Reducing network energy consumption via sleeping and rate adaptation. In: *Proc. 5th USENIX Symp. Networked Systems Design and Implementation, NSDI 2008* (April 2008)
- [14] BONE project. WP 21 Topical Project Green Optical Networks: Report on year 1 and updated plan for activities. NoE, FP7-ICT-2007-1 216863 BONE project (December 2009)
- [15] Nortel. A comparison of next-generation 40-Gbps technologies, white paper, <http://www.nortel.com/solutions/collateral/nn122640.pdf>
- [16] IEEE P802.3az. Energy Efficient Ethernet Task Force (2010), <http://grouper.ieee.org/groups/802/3/az>
- [17] Christensen, K., Reviriego, P., Nordman, B., Bennett, M., Mostowfi, M., Maestro, J.A.: IEEE 802.3az: The Road to Energy Efficient Ethernet. *IEEE Communications Magazine* (November 2010)
- [18] Ricciardi, S., Careglio, D., Palmieri, F., Fiore, U., Santos-Boada, G., Solé-Pareta, J.: Energy-aware RWA for WDM networks with dual power sources. In: *Proc. IEEE International Conference on Communications (ICC 2011)*, Kyoto, Japan, June 5-9 (2011)
- [19] Zhang, B., Sabhanatarajan, K., Gordon-Ross, A., George, A.: Real-Time Performance Analysis of Adaptive Link Rate. In: *Proc. Conference on Local Computer Networks* (October 2008)