

Evaluating the Energy-Awareness of Future Internet Devices

Raffaele Bolla

DIST - University of Genoa
Genoa, Italy
raffaele.bolla@unige.it

Chiara Lombardo

DIST - University of Genoa
Genoa, Italy
chiara@reti.dist.unige.it

Roberto Bruschi

CNIT
Genoa, Italy
roberto.bruschi@cnit.it

Diego Suino

TILab - Telecom Italia
Turin, Italy
diego.suino@telecomitalia.it

Abstract—Advanced power management capabilities have been proposed to be included into next-generation green network devices in order to modulate their energy requirements according to the workload. The clear side effect of enabling these new capabilities consists in a performance level reduction of network devices. Starting from some existing benchmarking standards for evaluating energy-efficiency, namely ECR and ATIS-060015, this contribution is devoted to determining a set of parameters, and methodologies that can be applied to correctly and precisely evaluate the tradeoff between energy consumption and network performance. Some experimental results obtained with the proposed indexes and methodologies, as well as a green SW router prototype have been provided.

Keywords—green networking; energy-efficiency evaluation; benchmarking methodologies.

I. INTRODUCTION

The research field of "green" and energy-efficient networking infrastructures has gained a great interest, from both service/network providers and equipment manufacturers, only in the last few years. Despite a more widespread sensitivity to ecological issues, such interest is due to heavy economic reasons, since both energy costs and network electrical requirements show a continuous growing trend over the past years.

Today, fixed and mobile network infrastructures have enormous and heavily increasing requirements in terms of electrical energy. For example, as shown in [1] and in [2], energy consumption of Telecom Italia network in 2006 has reached more than 2TWh (about the 1% of the total Italian energy demand), increasing by a 7.95% with respect to 2005, and by a 12.08% to 2004. Similar trends can be generalized to a large part of the other telecoms and service providers. The European Commission DG INFSO report in [3] estimated European telcos and operators to have an overall network energy requirement equal to 14.2 TWh in 2005, which will rise to 21.4 TWh in 2010, and to 35.8 TWh in 2020 if no green network technologies will be adopted.

Moreover, it is well known that networks, links and devices are provisioned for busy or rush hour load, which typically

exceeds their average utilization by a wide margin. Against such flat energy wastes, the specific challenge for telecoms, network equipment manufacturers and the networking research community regards nowadays mainly the introduction, the exploitation and the control of power management capabilities (i.e., sleeping and rate adaptation) inside architectures and components of network equipment. In this respect, current green networking approaches range from novel traffic engineering and routing criteria [4] [5], to the introduction of energy-aware equipment [6], components [7] and network interfaces [8]. The largest part of approaches undertaken is founded on few basic concepts, which have been generally inspired by energy-saving mechanisms and power management criteria that are already partially adopted in computing systems. These basic concepts can be classified as follows: (i) re-engineering approaches for reducing the maximum energy consumption of a device; (ii) dynamic adaptation approaches for modulating energy consumption according to the actual workload of the device; (iii) sleeping/standby approaches to enter unused parts of a device very low power sleeping states for long time periods.

In this work we will focus on dynamic adaptation schemes for network devices, which are usually achieved by modulating the service capacity in order to meet traffic loads (Adaptive Rate, AR), or forcing links and processing engines to enter low power states when no traffic is received (Low Power Idle, LPI). As better explained in the next section, the adoption of these schemes will make future Internet devices to be able to trade power consumption for network performance (e.g., packet forwarding latency, jitter, etc.).

By explicitly considering such novel energy-aware capabilities, our main objective in this contribution is to review and to extend recent standards, namely ECR 3.0.1 [10] and ATIS 060015.02.2009 [11], which have been proposed for evaluating and rating the energy-efficiency level of network devices. Both these standards simply define energy efficiency in terms of ratio between throughput and average power consumption without taking the effect of power management on packet forwarding performance explicitly into account. Moreover, both these standards use artificial benchmarking

scenarios (e.g., Constant Bit Rate (CBR) traffic flows), which are far from the real features of Internet traffic, and which may lead to misleading energy-efficiency evaluations, given the adaptive nature of AR and LPI schemes.

The first purpose of this paper is the analysis of the existing standards, in order to capture their strongest and weakest points. Then, starting from these considerations, our aim is twofold. On one hand, our goal is to extend current methodologies for evaluating the energy-efficiency of a network device. In more detail, our idea is to consider also testing traffic flows with different burstiness levels, which may better represent real Internet traffic than CBR flows. On the other hand, we propose a set of performance indexes able to represent the trade-off between energy consumption and network performance in an accurate way.

The paper is organized as follows. Section II introduces the main concepts that will be the basis for energy-aware optimizations in future Internet devices. Section III describes the existing standards proposed to evaluate the energy efficiency of a network device. Section IV introduces the proposed extensions to benchmarking methodologies and the novel performance indexes to represent the tradeoff between energy consumption and network performance. Section V shows some experimental results obtained with an energy-aware SW router, which already includes AR and LPI capabilities, performed by means of the proposed methodologies and performance indexes. Finally, the conclusions are drawn in Section VI.

II. GREEN OPTIMIZATIONS FOR FUTURE INTERNET DEVICES

As previously sketched, the largest part of approaches undertaken is founded on few basic concepts, which have been generally inspired by energy-saving mechanisms and power management criteria that are already partially adopted in computing systems. Following the taxonomy proposed in [15], these basic concepts can be classified as follows:

- Re-engineering;
- Dynamic adaptation;
- Sleeping/standby.

Re-engineering approaches aim at introducing and designing more energy-efficient elements for network device architectures, at suitably dimensioning and optimizing the internal organization of devices, as well as at reducing their

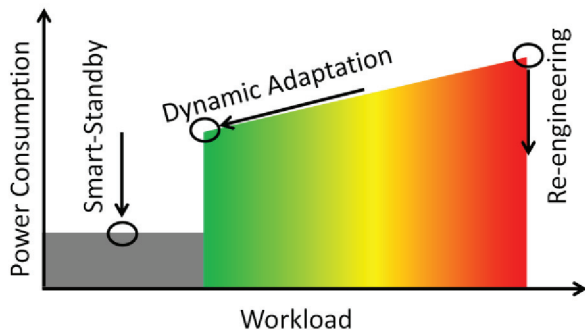


Figure 1. Power profile of future Internet devices and roles of re-engineering, dynamic adaptation and smart standby technologies.

intrinsic complexity levels.

The dynamic adaptation of network/device resources is designed to modulate capacities of packet processing engines and of network interfaces, to meet actual traffic loads and requirements. This can be performed by dynamically trading off device maximum performance for power consumption.

Finally, sleeping/standby approaches are used to smartly and selectively drive unused network/device portions to low standby modes, and to wake them up only if necessary. However, since today's networks and related services and applications are designed to be continuously and always available, standby modes have to be explicitly supported with special proxying techniques able to maintain the "network presence" of sleeping nodes/components.

It is worth noting that all these approaches are not mutually exclusive, and their joint adoption may eventually impact on next-generation devices by providing "energy-aware" profiles [16]. As shown in Fig. 1, re-engineering approaches will result in reducing the maximum power consumption of a device, and dynamic adaptation ones will enable to modulate energy absorption according to the actual workload. Finally, techniques based on smart standby will help to further cut power consumption of unused devices (or parts of them), by allowing HW components entering very low power sleeping states.

The adoption of dynamic power scaling and of smart standby optimizations obviously affect network performance, since these approaches are founded on the ideas of tuning device processing/transmission capacities, and of waking up the hardware upon "work" request. In other words, the use of such green optimizations allows trading energy consumption for network performance, and, in more detail, in terms of packet elaboration/transmission delay. In this respect, let us focus on dynamic adaptation schemes. They are usually founded two basic techniques: Adaptive Rate (AR) and Low Power Idle (LPI). The former allows dynamically modulating the capacity of a link, or of a processing engine, in order to meet traffic loads and service requirements while the latter forces links or processing engines to enter low power states when not sending/processing packets.

These techniques can be jointly adopted in order to adapt system performance to current workload requirements. For instance, The IEEE 802.3az task force [13] considered and evaluated both AR and LPI techniques, and decided to base the new Ethernet standard only on the LPI primitive. This decision stemmed from the need for maintaining the implementation complexity and cost as low as possible. In other network contexts, the evaluations and the resulting decisions may be very different.

In general purpose computing systems, the ACPI standard [14] models AR and LPI by introducing two sets of energy-aware states, namely performance and power states (P-states and C-states), respectively. Regarding the C-states, the C_0 power state is an active power state where the CPU executes instructions, while the C_1 through C_n power states are processor LPI states, where the processor consumes less power and dissipates less heat. On the other hand, as the sleeping




Figure 2. Packet service times and power consumptions in the following cases: (a) no power-aware optimizations, (b) only LPI, (c) only AR, (d) AR + LPI

power state (C_1, \dots, C_n) becomes deeper, the transition between the active and the sleeping state (and vice versa) requires longer time.

As shown in Figure 2, AR (Fig. 2-c) obviously causes a stretching of packet service times (i.e., header processing time in a processing engine, or packet transmission time in a link interface), while the sole adoption of LPI (Fig. 2-b) introduces an additional delay in packet service, due to the wake-up times. Finally, as outlined in Fig. 2-d, the joint adoption of both energy-aware capabilities may not lead to outstanding energy gains, since performance scaling causes larger packet service times, and consequently shorter idle periods.

However, the energy- and network-aware effectiveness of LPI and AR (and their possible joint adoption) must be accurately evaluated by taking HW and traffic features and requirements into account. In this respect, it is worth noting that the overall energy saving and the network performance strictly depend on incoming traffic volumes and statistical features (i.e., interarrival times, burstiness levels, etc.). For instance, idle logic provides top energy- and network-performance when incoming traffic has a high burstiness level. This is because less active-idle transitions (and wake-up times) are needed, and HW can remain longer in low consumption state.

III. EXISTING STANDARDS ON ENERGY EFFICIENCY EVALUATION

This section is meant to discuss and compare the three standard methods for measuring power consumption at varying performance levels of network devices and systems. In the following analysis, each method will be investigated singularly, in order to establish in a complete and, as much as possible, easy way, its specific behavior. Further considerations and comparisons will be provided at the end of the section. Although the last method here described has a different aim, as it does not perform any power consumption measure, it is included because it represents a valid guideline for both performance testing and set up procedures. Moreover, some of the tests it defines are already exploited by the other standards to select the starting parameters for further measurements.

A. ECR 3.0.1

ECR 3.0.1 [10] is a proprietary standard owned by Juniper and IXIA. It allows to measure the ratio between power and

performance as well as the consumption due to single components (in presence of redundancy) and the system energy consumption over a projected lifetime. The main test, defined as mandatory, is performed measuring the maximum offered load the device under test can support without losing any data. A router tester is used for this purpose. Traffic is sent at that speed for 1200 seconds, and a watt-meter measures the average power consumed in this time interval. Traffic is then sent at 50%, 30%, 10% and 0% of the maximum throughput, and average power consumption is measured during each transmission. The chosen packet size is the maximum supported by that kind of SUT. The ECR (Energy Consumption Rating) index can be expressed as follows:

$$ECR = \frac{\Phi_{100}}{T_{100}} \quad (1)$$

where Φ_{100} represents the average power consumption at maximum throughput and T_{100} is the throughput itself.

The ECR index has the clear and simple aim of quantifying how much energy is needed to forward one Gigabit of data at the maximum speed. This performance index is clearly very “unnatural”, since it evaluates the System Under Test (SUT) at an operating condition (maximum speed with Constant Bit Rate (CBR) traffic flows) that is seldom (or never) reached in real-world deployment scenarios.

Notwithstanding these considerations, the ECR index can be obviously useful to evaluate the SUT “green” re-engineering level, with reference to Section II and to Fig. 1. Unfortunately, the ECR does not include any indications regarding power management mechanisms eventually included in the SUT, since their effects become evident only at low levels of traffic offered load.

However, the standard introduces an optional test aimed at estimating a version of the ECR index, called ECR-VL (ECR-Variable Load), weighted for various offered load (0%, 30%, 50%, 100% of the effective throughput).

Unfortunately, all the benchmarking tests use CBR traffic flows with fixed sized packets.

B. ATIS 060015.02.2009

As far as the testing methodologies are concerned, the ATIS 060015.02.2009 standard [11] looks very similar to the ECR 3.0.1. However, it is addressed to transport/core network products, owned by the carrier and explicitly characterized by

TABLE I. COMPARISON OF THE ANALYZED BENCHMARKING STANDARDS

Standards	Benchmarking scenario	Main results obtained
ECR 3.0.1	<ol style="list-style-type: none"> 1. CBR traffic 2. Traffic at different loads 3. Maximum packet size 	<ol style="list-style-type: none"> 1. ECR [W/Gbps] 2. ECR-VL [W/Gbps], ECR-EX [W/Gbps] 3. Energy Bill Estimates [\$]
ATIS 060015.02.2009	<ol style="list-style-type: none"> 1. CBR traffic 2. Traffic at different loads 3. IMIX traffic 	<ol style="list-style-type: none"> 1. DTEER [Mbps/W] 2. CTEER [Mbps/W]
RFC 2544	<ol style="list-style-type: none"> 1. CBR traffic 2. Bursty traffic 3. Traffic at different loads 4. Range of packet sizes 	<ol style="list-style-type: none"> 1. Maximum throughput [Gbps] 2. Max, average, min latency [us] 3. Maximum burst length [bursts/frames]

redundancy.

The test result, in fact, is an index called TEER (Telecommunication Energy Efficiency Ratio) that can be obtained putting together the results of each component, collected in a database. It is also possible to provide a series of certified configurations, each one characterized by its index.

For each component, the database maintains the maximum throughput and the average power consumed at 100%, 50% and 0% of that throughput. TEER is obtained as the ratio between the sum of the maximum rate for each component and the sum of the consumptions, for each component, at the three traffic loads previously mentioned. In order to simulate a real world behavior, Internet MIX (IMIX) traffic profile [17] is used for the transmissions, which means that packets of different lengths are generated by the test equipment according to a statistic distribution and transmitted to the SUT.

The main difference between the two standards can be seen in the amount of produced results: ATIS 060015.02.2009 only provides TEER, while ECR 3.0.1, aside from the index obtained with the mandatory test, gives a wider data volume, including average consumption measured on different traffic rates. The devices taken into consideration by the second standard, however, are more specific, and they focus mainly on modularity.

C. RFC 2544

Notwithstanding the RFC 2544 [12] is not oriented to evaluate energy-efficiency, we decided to include it in this section, since it is the “mother” of all standards for benchmarking network devices. Thus, recalling its base concepts may be useful to extend ECR 3.0.1 and ATIS 060015.02.2009.

This standard provide a very complete set of methodologies and performance indexes for evaluating network performance. The same tests can be applied to all kinds of network interconnection devices, in all their possible configurations. For the sake of this study, we can take into consideration only the following tests:

- *Throughput Test*: the throughput is the fastest rate at which the count of test frames transmitted by the SUT is equal to the number of test frames sent to it by the tester. For each available frame length the maximum throughput is calculated; this provides not only the first test result, but also a binding datum for the next measures.

- *Latency Test*: the procedure that determines latency uses the previous result as a starting point. Latency represents the interval between the time at which a frame is fully transmitted and the time at which the same frame is received.
- *Back-to-back Test*: bursty traffic characterized by the maximum burst length allowed is generated by the tester. The back-to-back value is the number of frames in the longest burst that the SUT can handle without losing any data.

D. Comparisons Between the Standards

The main features and performance indexes obtained with the considered standards are summarized in Table I.

The main characteristic of RFC 2544 is its multi-purpose nature: although it is designed for evaluating energy efficiency, its concepts and methodologies can be deeply exploited within other contexts. The fact that it can be applied to all kinds of network systems, embracing different protocols, topologies and traffic nodes testifies for its versatility.

Considering ECR 3.0.1 and ATIS 060015.02.2009, both of them are focused only on the measurements of energy consumption, and then they lack clear indications on evaluating the trade-off between power absorption and network performance. In addition, since none of their methodologies takes into consideration bursty traffic, the obtained results cannot provide an accurate representation of real traffic.

IV. EXTENDING CURRENT EVALUATION METHODOLOGIES

As already sketched in the previous section, the existing standards present some interesting features, but at the same time they lack of important indications that would be necessary for a thorough evaluation of future network devices. Latency is not taken into consideration neither in ECR 3.0.1 nor in ATIS 060015.02.2009, as they both use only throughput to characterize performance. But in energy-aware systems power reduction tends to cause an increasing of service times, generating delays. For this reason, a correct exploitation of energy efficiency techniques cannot be evaluated regardless of latency. In addition, both standards employ CBR traffic to determine the average power consumed during the transmission. This assumption can be unrealistic if we consider the fact that real traffic generally has a bursty composition. However, CBR traffic could be considered as a limit case and included in a wider set of measurements.

Regarding results' representation, determining a global index would be desirable, since it is more synthetic and gives an immediate idea on a system behavior. On the other hand, some critical aspects could be neglected without any further information.

Starting from these considerations, our method consists in determining two indexes: one represents the amount of energy that can be saved through power management capabilities, while the other one quantifies the performance degradation that these capabilities can provoke. These indexes will be explained in a more detailed way in the next Subsection, followed by a description of the tests needed to obtain the desired parameters.

A. Performance Indexes

Apart from its limitations, that have already been discussed, the index provided by ECR 3.0.1 shows in a simple way the amount of energy consumed to forward one Gigabit of data. For this reason, it can be usefully taken into consideration to represent how a device reacts to a limit behavior of traffic such CBR is.

The information obtained with this result becomes not sufficient especially in presence of power management policies. As previously sketched, both AR and LPI techniques have the increasing of latency as a drawback. Power scaling causes a stretching of packet service times, while idle logic introduces an additional delay in packet service, due to wake-up times.

Moreover, power scaling causes larger packet service times, and consequently shorter idle periods. For all these reasons, it is necessary to collect and put data together to quantify what we gain in terms of energy saving and what we lose in terms of service delays.

Energy gain represents the power saving obtained thanks to power management in comparison to a scenario with no such capabilities:

$$\Phi_{\%} = \frac{\Phi_{max} - \Phi_c}{\Phi_{max}} \quad (2)$$

where Φ_c is the current power consumption, Φ_{max} is the maximum consumption reached by the device. The result gets closer to 1 as consumption decreases.

Performance degradation is expressed as a ratio between the values of packet latency in the case of an ideal network device (i.e., with an infinite processing capacity), and the ones measured with the real SUT:

$$L_{\%} = \frac{\tilde{L}_i + \frac{1}{2}(L_i^{max} - L_i^{min})}{L_r + \frac{1}{2}(L_r^{max} - L_r^{min})} \quad (3)$$

where parameters with the i index represent the latencies for the ideal device, and the ones with the r index the ones measured on the real SUT. L^{min} , \tilde{L} and L^{max} are the minimum, average and maximum values, respectively, of packet latencies.

The values for minimum, average and maximum latency in the ideal case have to be computed by starting from the packet transmission times on input and output links.




Figure 3. Packet burst processing in an "ideal" SUT.

As first, it is worth noting that if the packet transmission time on the input link is longer or equal to the one of output link, then $L_i^{min} = \tilde{L}_i = L_i^{max}$. Otherwise, as shown in Fig. 3, these latency values have to be estimated by considering traffic burstiness in an explicit way.

We can simply compute the L_i^{min} value as the sum of the packet transmission times on the input and output links (τ_{rx}, τ_{tx}):

$$L_i^{min} = \tau_{rx} + \tau_{tx} \quad (4)$$

With reference again to the simple example in Fig. 3, the L_i^{max} parameter can be calculated as follows:

$$L_i^{max} = \tau_{rx} + n\tau_{tx} - (n-1)\tau_{rx} = n\tau_{tx} - (n-2)\tau_{rx} \quad (5)$$

where n is the length of the incoming bursts.

Thus, the \tilde{L}_i value can be simply expressed as:

$$\tilde{L}_i = \frac{n+1}{2}\tau_{tx} - \frac{n-3}{2}\tau_{rx} \quad (6)$$

B. Benchmarking Methodologies

As in the in ECR and ATIS standards, we suggest to start benchmarking tests by exploiting RFC 2544 Throughput Test. But differently from the above methodologies, we propose to perform test sessions with:

- different loads, namely 100%, 50%, 30%, 10% and 0% of output link bandwidth (and not of the effective throughput – this will make more comparable results obtained with different SUTs or with different configuration of the same SUT);
- different packet burst sizes, namely 1 (CBR flows), 10, 50, 100, 200 packets;
- different packet sizes (as suggested by the RFC 2544), and where possible with IMIX profiles.

It is worth noting that, when possible, packets shall have different IP source and destination addresses, since devices often include multiplexer and de-multiplexer elements working on a per-flow basis. So that different IP couples can be classified as different flows, and then trigger different parallel HW pipelines internally to the SUT.

On each session, the average power is measured (we stabilized measured values over 5 minutes' long observations, after 2 minutes of session startup, when the traffic is injected but we do not collect measurements). Regarding network performance we collect different values: the average

throughput, the minimum, average and maximum value of packet latency, and packet loss rates. If the SUT has the possibility of entering various power management configurations, measures are repeated for each of them.

V. NUMERICAL RESULTS

In this Section, a selection of the results obtained with the methodology introduced in Section III is reported. The SUT used to perform the tests is a Linux-based software router [18]. It supports two Intel Xeon X5550 processors, each one with 4 physical cores. The network adapters are two dual Intel Gigabit Ethernet interfaces. The operating system is Linux Debian 5.0.6 and the kernel version is 2.6.30.1. The power management mechanism is performed by the ACPI (Advanced Configuration and Power Interface) technology [14]. This means that, for each CPU, the user can choose the working frequency and idle state, called P- and C-states, according to a desired performance level. The frequencies used to perform these tests are $P_0=2.67$ GHz and $P_8=1.60$ GHz, idle states are C_1 and C_3 (deepest idle state).

For each P- and C-state configuration, tests have been repeated varying the throughput (at traffic loads indicated in Section IV B) for different packet and burst lengths. We used packets of 64 and 1500 B and bursts of 1 (i.e. CBR traffic), 10, 50, 100 and 200 packets.

An Ixia Router Tester [19] is connected to the software router and sends traffic to it; when traffic is sent back the tester can determine traffic statistics. An Agilent watt-meter [20] is also connected to the software router, in order to register the power consumed during the tests. Both the tester and the watt-meter report their results to a controller.

Table II shows energy efficiency results expressed in terms of ECR index for the four test cases. The first thing to be noticed is that, at high offered loads, there is no difference between continuous and bursty traffic: transmitted data are so fast that there is not a relevant interarrival time between the reception of two bursts, turning it into an actual CBR traffic. Another evidence is that packets of 1500 B always give the best results, with no difference among traffic loads or device configurations. In fact, transmitting the same amount of data using longer packet sizes means less headers to be processed by the system, hence a faster service, so consumption is lower for the same offered load and ECR decreases.

Focusing again on higher throughputs, the best results for packet size 64 B are obtained in P_0 , while at the lowest frequency packets of 1500 B are more efficient. A higher working frequency allows us to reach the best performance

even for short packets, at the price of more power consumed. Instead, when the frequency is low, energy reduction is not sufficient to compensate a worse service rate, in accordance to the example shown in Figure 2c. This behavior does not apply to packets of 1500 B, where the lowest frequency still allows both power reduction and a good performance level. If we consider the lower throughputs, difference between results for continuous and bursty traffic is more evident. When the traffic is low, in fact, the interarrival time between two bursts is longer and the system can enter idle states and save energy.

Taking again the whole set of results of Table II into account, it is clear that the working frequency has a stronger impact than the idle state on energy efficiency calculated in terms of ECR. In fact, as we have seen, throughput is only influenced by the P-state, while the contribution of C-states on power consumption is evident only for low traffic loads. But, as already discussed, the main limit of ECR is the absence of latency among its parameters. The results of performance degradation and energy gain in Figures 4-11 show how our methodology allows characterizing and quantifying the impact of both AR and LPI optimizations. The predominant effect of the C-state over latency, with respect to the frequency, is plain considering Figures 4 and 8: we have the same results whether we work at 2.67 GHz or at 1.60 GHz. The same consideration is valid for Figures 6 and 10. In detail, for the cases characterized by C_1 , we can see that the worst results are obtained for packet size 64 B at throughput 100% and 50%: short packets mean more headers, hence the service time turns out to be increased.

Degradations for packet size 1500 B have similar values at varying offered load, since the service capacity is never pushed to the limit and so delays are controlled. At throughput 30% and 10%, the difference between short and long packets becomes more evident for bursts longer than 50 packets, where 1500 B gives visibly best results.

Similar consideration can be applied to Figures 6 and 10: only, in these cases we see very low values for 1500 B packets at full rate. Going back to Table II, for these configurations we had the lowest ECR. Degradation results outline how the deepest idle state allows reaching the maximum throughput with a huge effort from the SUT, which is translated into higher latencies.

As power is influenced more by P-states, we can see that energy gain results in Figures 5, 7, 9 and 11 show similar trends for the same frequency. As an example, let us consider the lines representing packet size 64 B, throughput 10% in Figures 4 and 5. For bursts of 1 and 10 packets we have a good

TABLE II. ECR VARIABLE LOAD VALUES ACCORDING TO DIFFERENT ENERGY-AWARE CONFIGURATIONS, PACKET AND BURST SIZES

Offered Load	$P_0 C_1$				$P_0 C_3$				$P_8 C_1$				$P_8 C_3$			
	64 B		1500 B		64 B		1500 B		64 B		1500 B		64 B		1500 B	
	1 pkt	200 pkt	1 pkt	200 pkt	1 pkt	200 pkt	1 pkt	200 pkt	1 pkt	200 pkt	1 pkt	200 pkt	1 pkt	200 pkt	1 pkt	200 pkt
100%	258	258	199	199	231	231	182	182	490	490	196	196	450	450	182	182
50%	508	505	398	394	463	460	361	339	490	490	392	392	450	450	363	339
30%	769	762	669	662	708	696	614	510	746	746	662	662	692	692	610	510
10%	2463	2413	1940	1910	2275	1850	1770	1360	2400	2400	1920	1920	2238	1850	1760	1360

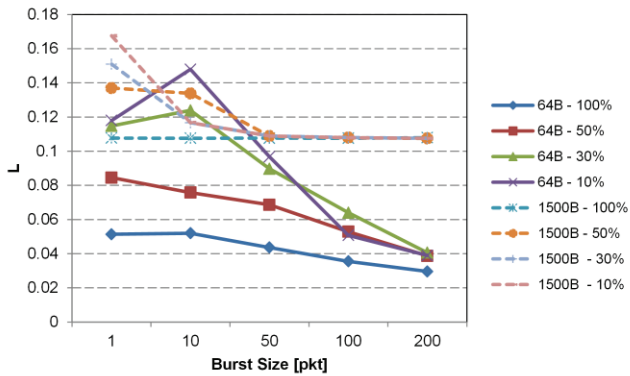


Figure 4. $L\%$ values for the P_0, C_1 configuration according to 64 and 1500 B sized datagrams and different traffic load and burstiness levels.




Figure 5. Φ values for the P_0, C_1 configuration according to 64 and 1500 B sized datagrams and different traffic load and burstiness levels.

level of performance, while energy gain is very low. As the burst length grows, we detect a decay in performance in association with higher power savings.

The same considerations can be extended to all other test cases. In order to give a global idea on how this trade-off appears, besides the decreasing trend of performance degradations and the growing on of energy gains, it is possible to see that results in Figures 4, 6, 8, and 10 cover a wider range of values for short bursts, whereas in Figures 5, 7, 9 and 11 this spreading appears for bursts longer than 50 packets.

VI. CONCLUSIONS

In this work we considered AR and LPI adaptation schemes for network devices, which may lead future Internet devices to trade power consumption for network performance (e.g., packet forwarding latency, jitter, etc.).

By explicitly considering such novel energy-aware capabilities, our main objective in this contribution was to review and to extend recent standards, namely ECR 3.0.1 [10] and ATIS 060015.02.2009 [11], which have been proposed for evaluating and rating the energy-efficiency level of network devices. In more detail, we outlined that both the ECR and ATIS standard simply provide some performance indexes, guidelines and base configurations that could be enough for evaluating the power management features of a device in real-world scenarios. As first we extended their testing

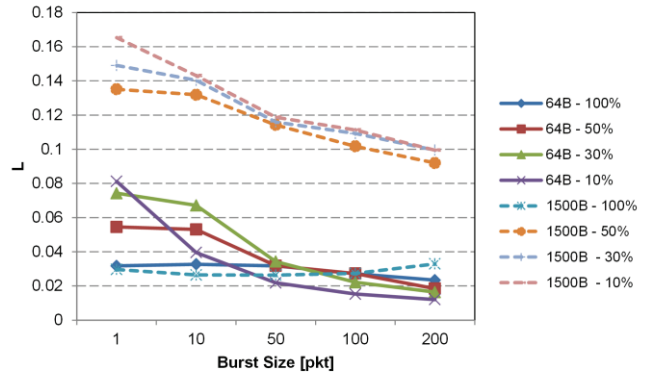


Figure 6. $L\%$ values for the P_0, C_3 configuration according to 64 and 1500 B sized datagrams and different traffic load and burstiness levels.

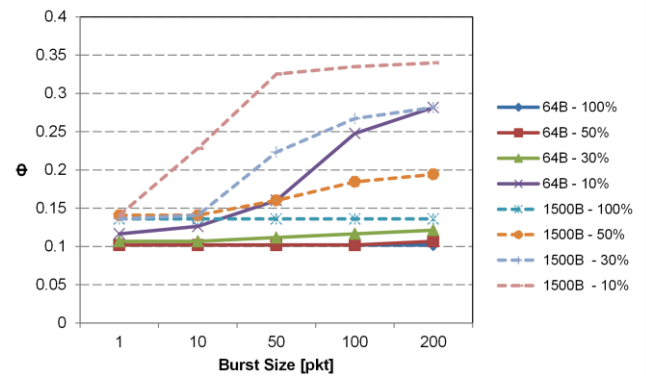


Figure 7. Φ values for the P_0, C_3 configuration according to 64 and 1500 B sized datagrams and different traffic load and burstiness levels.

methodologies for evaluating the performance of the SUT under traffic flows with various burstiness levels. Secondly, we proposed a set of performance indexes for evaluating the energy consumption gains with respect to the decrease of network performance (especially in terms of packet forwarding latency).

Tests performed on an energy-aware SW router (AR and LPI capable) showed how the proposed indexes allow to represent the exchange between power consumption and network performance in presence of power management capabilities in a clear and, at the same time, complete way. Furthermore, results in terms of energy gain and performance degradation also allowed to capture the different effects of each power management strategy on the device global behavior.

REFERENCES

- [1] C. Bianco, F. Cucchietti, G. Griffla, "Energy Consumption Trends in the Next Generation Access Network - A Telco Perspective", Proc of the 29th International Telecommunications Energy Conference, 2007. (INTELEC 2007), Rome, Italy, Sept. 2007, pp. 737-742.
- [2] Telecom Italia Website, "The Environment", URL: <http://www.telecomitalia.it/sostenibilita2006/english/B05.html>.
- [3] European Commission DG INFSO, "Impact Of Information and Communication Technologies on Energy Efficiency", final report, Sept. 2008, http://ec.europa.eu/information_society/newsroom/cf/itemdetail.cfm?item_id=4441.




Figure 8. $L\%$ values for the P_8, C_1 configuration according to 64 and 1500 B sized datagrams and different traffic load and burstiness levels.




Figure 10. $L\%$ values for the P_8, C_3 configuration according to 64 and 1500 B sized datagrams and different traffic load and burstiness levels.




Figure 9. Φ values for the P_8, C_1 configuration according to 64 and 1500 B sized datagrams and different traffic load and burstiness levels.




Figure 11. Φ values for the P_8, C_3 configuration according to 64 and 1500 B sized datagrams and different traffic load and burstiness levels.

- [4] L. Chiaraviglio, M. Mellia, F. Neri, "Reducing Power Consumption in Backbone Networks", Proc. 2009 IEEE Internat. Conf. on Communications (ICC 09), Dresden, Germany, June 2009.
- [5] J. C. Cardona Restrepo, C. G. Gruber, C. Mas Machuca, "Energy Profile Aware Routing", Proc. Green Communications Workshop in conjunction with IEEE ICC 09 (GreenComm09), Dresden, Germany, June 09.
- [6] S. Nedeveschi, L. Popa, G. Iannacore, D. Wetherall, S. Ratnasamy, "Reducing Network Energy Consumption Via Sleeping and Rate-Adaptation", Proc. of the 5th USENIX Symposium on Networked Systems Design and Implementation, San Francisco, CA, 2008, pp. 323-336.
- [7] J. Noguera, I.O. Kennedy, "Power Reduction in Network Equipment Through Adaptive Partial Reconfiguration", Proc. of the 2007 International Conference on Field Programmable Logic and Applications (FPL 2007), Aug. 2007, pp. 240-245.
- [8] M. Gupta, S. Singh, "Using Low-Power Modes for Energy Conservation in Ethernet LANs", Proc. of the 26th IEEE Conf. on Computer Comm. (Infocom 07), Anchorage, Alaska, USA, May 2007.
- [9] The Home Gateway Initiative (HGI), URL: www.homegatewayinitiative.org/
- [10] A. Alimian, B. Nordman, D. Kharitonov, "Network and Telecom Equipment - Energy and Performance Assessment", Draft 3.0.1, October 04, 2009, URL: http://www.ecrinitiative.org/pdfs/ECR_3_0_1.pdf.
- [11] E. Eckert, K. Sievert, B. Hunck, C. Underkoffler, "Energy Efficiency for Telecommunication Equipment: Methodology for Measurement and Reporting Transport Requirements", February 10, 2009.

- [12] S. Bradner, J. McQuaid, "Benchmarking Methodology for Network Interconnect Devices", IETF Request for Comments no. 2544 (RFC 2544), <http://www.ietf.org/rfc/rfc2544.txt>.
- [13] IEEE P802.3az Energy Efficient Ethernet Task Force. URL: <http://www.ieee802.org/3/az/>
- [14] The Advanced Configuration & Power Interface (ACPI) Specification - Revision 4.0, <http://www.acpi.info/>.
- [15] R. Bolla, R. Bruschi, F. Davoli, F. Cucchietti, "Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures," IEEE Communication Surveys and Tutorials (to appear).
- [16] R. Bolla, R. Bruschi, K. Christensen, F. Cucchietti, F. Davoli, S. Singh, "The Potential Impact of Green Technologies in Next-Generation Wireline Networks—Is There Room for Energy Saving Optimization?," IEEE Communication Magazine (to appear).
- [17] "IMIX (Internet MIX)," Test Methodology Journal, Spirent Communications, March 2006, URL: <http://spcprev.spirentcom.com/documents/4079.pdf>
- [18] R. Bolla, R. Bruschi, A. Ranieri, "Energy Saving Support for PC-based Software Router: Performance Evaluation and Modeling," Proc. of the 2009 IEEE International Conference on Communications (IEEE ICC 2009), Dresden, Germany, June 2009.
- [19] The Ixia XM2 router tester, URL: http://www.ixiacom.com/products/chassis/display?key=ch_optixia_xm2.
- [20] Agilent U2700 modular system, URL: http://www.home.agilent.com/agilent/redirector.jsp?key=1537556&action=ref&lc=ita&cc=IT&cna=AGILENT_EDITORIAL