The Hidden Cost of Network Low Power Idle

Raffaele Bolla[1,2], Roberto Bruschi[2], Paolo Lago[1,2]
1. DITEN – University of Genoa
Genoa, Italy
{name.surname}@unige.it
2. CNIT – Research Unit of the University of Genoa
Genoa, Italy
{name.surname}@cnit.it

***Abstract*—This paper deeply and experimentally analyzes the efficiency of low power idle techniques when applied to packet processing engines of network devices. To this purpose, we set up a complex testbed that allowed us to perform several measurements on energy- and network-performance indexes. The reference device platforms that we selected for this evaluation are new generation software routers based on component-off-the-shelf hardware, since they already include advanced power management capabilities, and can be considered as a significant example for next-generation green network devices. The results collected in the measurement campaign allowed us not only (*i*) to provide an in-depth energy consumption profiling of SR data-plane, but also (*ii*) to clearly outline energy costs due to the use of low power idle techniques. Among other interesting aspects, we completely characterized the energy consumption due to wakeup transitions, which may cause instantaneous consumption spikes higher than 4 times the power energy requirement when active.**

Keywords- green networking; low power idle; experimental testbed.

I. INTRODUCTION

The energy efficiency issue in network devices has recently gained great interest on the part of network providers and equipment manufactures. Despite obvious environmental reasons, this interest springs from heavy and critical economical needs, since both energy cost and network electrical requirements show a continuous growing with an alarming trend over the past years [1,2]. In this scenario, several works proposed in literature addressing the possibility of introducing power saving strategies in network devices [3-5]. Such strategies are generally founded on two main kinds of power management capabilities provided by the HW level, namely *Adaptive Rate* (AR) and *Low Power Idle* (LPI). The former allows dynamically modulating the capacity of network device resource according to current traffic loads by reducing the service capacity (e.g., link bandwidths, computational capacities of packet processing engines, etc.). On the other hand LPI reduces power consumption by powering off sub-components when no activities are performed, and by re-waking them up when the system receives new activities.

At the state-of-the-art of green technologies, LPI technique is the most wide-spread strategy implemented in processors and network interfaces. For example, the novel IEEE 802.3az standard (Energy Efficiency Ethernet) supports LPI, which allows to partially turn off the physical layer during link inactivity periods. The initial approach for Energy Efficient Ethernet was based on AR, but it was abandoned in favor of LPI primitives. In fact, applying AR strategies in network devices may raise different technological/applicability issues. The first one is related to the time spent to switch components to different speed, which can introduce additional delay in packet service time. Moreover, if switching is exceedingly greater than the packet time scale, an optimization procedure may not quickly react to sudden increases in traffic volumes. Notwithstanding adopting only LPI techniques in network devices is the simplest solution, it may result in not so efficient operational behavior. For instance, several studies focused on evaluating IEEE 802.3az [6-9] pointed out a not negligible waste of energy due to excessive overhead time required to enter and exit from LPI modes. In fact, most of the energy is spent for triggering transitions between link active and sleeping states, rather than on actual data transmission. In this respect and for this reason, some studies [11,12] outlined that the relationship between traffic load and power consumption is very distant from a linear proportionality. In order to mitigate LPI transition overhead, several packet coalescing policies have been proposed [3,6,8,9] to improve energy efficiency of LPI-capable network devices through traffic burstification. In fact, LPI provides top energy and network performance when incoming traffic has high burstiness levels, since, fixed the incoming load, it has to perform less active-idle transitions.

Starting from this scenario, this paper aims at providing an in-depth experimental analysis of LPI techniques when applied to network forwarding hardware, trying to outline their main benefits and drawbacks under heterogeneous traffic patterns. To this purpose and without making this analysis losing of generality, we focused on packet processing engines, which represent one of the most energy-harvesting elements in many network platforms [10]. In more detail, we chose to use new generation Software Routers (SR) [13], based on open source SW and on multi-core Commercial Off-The-Shelves (COTS) hardware, for a number of different reasons. As first COTS SR already provides mature, advanced, and fully configurable power management capabilities by means of the Advanced Configuration and Power Interface (ACPI) [14]. Moreover, multiple LPI configurations are usually available in COTS processors. The last, but not the least important, aspect consists of the "open" access to every SR hardware internal interface and software package details. This open access has allowed us to develop ad-hoc testbed instruments and probes able to capture internal energy parameters in an effective and fine-grained way.

The collected results demonstrate that LPI performance are affected not only by the overhead time to move between active

and idle modes, but also by energy absorption peaks needed to wake-up and re-power on sleeping HW components. These peaks can reach also four times the nominal power consumption in active mode, and their size may depend on the operating clock frequency of the processor, as well. We also estimated that these wake-up energy spikes might contribute up to more than 15% of the average energy consumption under realistic scenarios.

The remainder of this paper is organized as follows. The experimental testbed and methodology is described in Section II, and the measurements results are in Section III. Finally, the conclusions are drawn in Section IV.

II. EXPERIMENTAL TEST-BED AND METHODOLOGY

In order to characterize LPI capabilities on SR, we decided to setup a complex experimental testbed, which involves three main elements (Fig. 1): the System Under Test (SUT – a Linux SR), DC and AC power measurement instruments, and the IP network with Ixia N2X router tester, useful to generate and to measure the performance of traffic flows stressing the SUT. The rest of this section is devoted to introduce the above-mentioned testbed elements into detail.

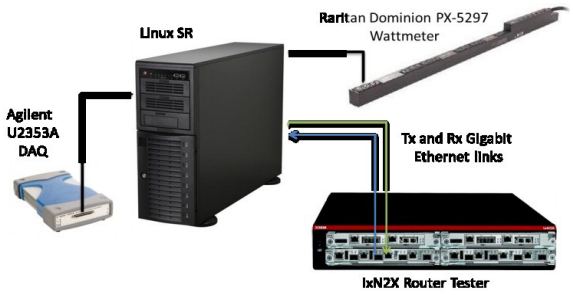

Figure 1. Experimental test-bed.

A. The System Under Test

The measurements were performed on Linux-based Software Router (SR) [13] based on commodity hardware and the Linux operating system. In more detail, the SR is equipped with an Intel Core i5 processor, with 4 physical cores and a maximum power consumption of 95W [17]. The network interface card (NIC) is an Intel PRO 1Gigabit Ethernet adapter. The operating system is the Linux Debian 5.0.6, and the kernel version is the "vanilla" 2.6.30.1, which supports symmetric multi-processor (SMP).

The power management of the Intel i5 processor is configurable through the ACPI [14], and consequently the processor's working frequency and idle state (P- and C-states, respectively) can be set according to desired performance profiles. The idle modes we applied in the tests correspond to some ACPI idle states, which can be considered as lightest sleeping ones, and are:

- C1 state: in this state, no instructions are being executed. When a core is in C1, the Intel Core i5 controller turns off the clocks of all clock domains pertaining to the Core pipeline [17].
- C3 state: the Core Phase-Locked-Loops (PLLs) are turned off, and the Cores flush the contents of their Level 1 (L1) instruction cache, L1 data cache, and Level 2 (L2) cache to the shared Level 3 (L3) cache. The wakeup time for idle state C3 is significantly longer than in state C1 since the L1 and L2 caches must be restored [17].

Regarding the P-states, the Intel Core i5 supports Dynamic Voltage and Frequency Scaling (DVFS) by means of Enhancement Intel speed-step technology, which provides multiple frequency and voltage points for optimal performance and power efficiency. For sake of simplicity, but without losing generality, we decided to analyze the power management of just a single Core.

B. Power measurments

In order to provide a complete characterization of the energy absorption dynamics, we decided to collect two main kinds of data, the energy absorption of the entire SUT on the AC plug and the energy absorption of the most significant HW subcomponents of the SUT (e.g., CPU, RAM, PCI network cards, etc.). On one hand, the measures on the AC plug were simply collected by means of a Raritan Dominion PX-5297 wattmeter [18]; on the other hand, the measures on the internal subcomponents required a certain effort. In fact, PC HW does not usually include power probes, and the power supply to subcomponents is generally carried within multiple standard physical bus interfaces (e.g., different sockets for CPUs, DIMM and SO-DIMM slots for memories, PCI-e slots for network interface cards, etc.).

In order to measure the absorption of certain sub-components, where possible, we used some special bus risers [19] equipped with current and voltage measurement probes. Unfortunately, this approach cannot be applied to many physical interfaces, due to issues related to the much miniaturized physical connector dimensions, to the stability of electrical signals, etc. This last case applies for CPU sockets. However, to measure the energy consumption of CPU, we exploited the ATX standard [20] that defines the internal power supply for general purpose computing systems. ATX motherboards are supplied by means of a 24 pin connector, which includes among the others three main voltage rails at 3.3, 5 and 12 V, and a ground rail. The 12 V rail supplies many components of the motherboard and part of the CPU. An additional 6/8 pin ATX connector, carrying a second 12 V rail, is generally available on the motherboard to provide additional power to the CPU. In more detail, the 12 V rail on the 6/8 pin connector is commonly devoted to provide isolated power supply to the CPU Cores only. Other subcomponents of the CPU (e.g., cache, bus and DDR controllers and other internal control units) are supplied by means of the 12 V rail on the 24 pin connector.

Owing the above mentioned standard characteristics of PC supplies, we decided to monitor the CPU energy consumption by sensing both the 12 V ATX rails. To this purpose, we designed and developed a riser board for ATX power connectors, which allows putting some current and voltage probes on the available supply rails (mainly the two 12 V rails, and the 5 and 3.3 V ones). Figure 2 shows the high-level architecture implemented by the ATX riser board. The current

probes are realized by means of very small (10 mΩ) precise Current Sensing Resistors (CSR). The voltage drop V_{sense}^{n} across CSRs (where n denote the n^{th} ATX voltage rail) is proportional to the incoming current. To reduce the noise and eventual cross-talk level on the measured V_{sense}^{n} values, we decided to closely connect the CSRs to some current-sense amplifiers. The amplified CSR voltage drops and the ATX rail voltages are finally measured by an external Data AcQuisition (DAQ) device. In our tests, we used An Agilent U2353A multifunction Data Acquisition (DAQ), which can simultaneously sample 4 channels at 250 KHz with a 16 bit resolution [22].

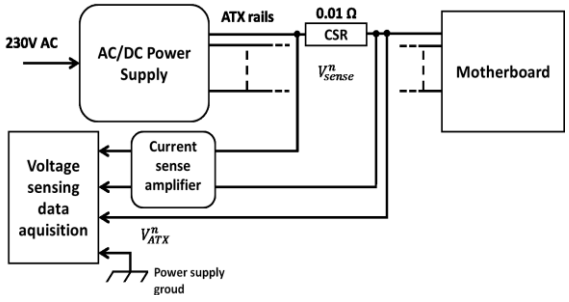

Figure 2. High-level scheme of the riser board for ATX voltage rails.

III. EVALUATION RESULTS.

By using the experimental test-bed described in Section II, we performed several power consumption and performance measurements. Subsection A reports a general characterization of the SR data-plane under an energy perspective. Subsection B deeply analyses the energy consumption overhead of LPI transitions.

A. Energy profile of SR Data plane

The first measurements we collected were aimed at evaluating the power consumptions of different SUT HW components, and their proportionality to the actual workload. To this purpose, the power consumption of these internal elements were measured both when the system is idle at the C3 power state (turned on, but performing no operations), and when one Core is active and performing operations at the maximum speed. The obtained results are in Fig. 3, and, as expected, outline how the CPU is the most energy-hungry component, exhibiting also a clear dependency on the workload since it passes to weight from 28% of the overall consumption in idle mode to 40% when active. Except for the AC/DC power supply (whose efficiency depends on the electric load of internal SUT components), all the other components exhibit negligible energy absorption variations.

Figure 4 shows the global power consumption (performed by the RDPX) and breakdown power consumption for the different ATX power rails at the C1 and C3 power states. As we can see, changing the idle level of only one Core from C1 to C3 provides a significant power save of about 10W. Moreover, we can notice that only the 12V 24-pin rail changes the power consumption with the C-state level. This is because, as previously sketched in section II.B, 12V 24-pin rail supply among the others, the CPU cache, which is turned off when entered in C3. To estimate the active power dissipation of the CPU, we stress the SUT by generating Constant Bit Rate (CBR) traffic at the maximum bandwidth of a 1 Gigabit Ethernet composed by 64B IP packets. In this way the Core 0 work for the 100% of the time, so power consumption only depends on working frequency. Figure 5 shows the active power consumption for a sub-set of available frequencies and the related maximum throughput.

Obviously, scaling frequency allows the reduction of the power consumption at the price of slower maximum forwarding rates. In detail, differently from the idle case, the power supply from 12V 8-pin rail change with respect of the working frequency, while the 12V 24-pin power remains almost constant. While frequency scaling allow to modulate the active power consumption with respect the device workload, LPI reduces the idle power consumption at price of introducing an additional delay in the packet service, due to the wake-up times. In order to evaluate the joint adoption of these power saving strategies in a realistic network traffic scenario, we perform several power consumption measurements in presence of bursty traffic by changing the AR and LPI configuration.

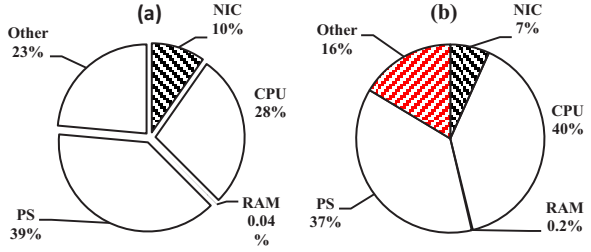

Figure 3. Energy consumption of the different internal HW components of the SUT during idle at C3 power state (a) and active (b) periods (RAM: memory, NIC: network interface card; PS: AC/DC power supply; Other: all the other components in the SUT, from motherboard chips to disks).

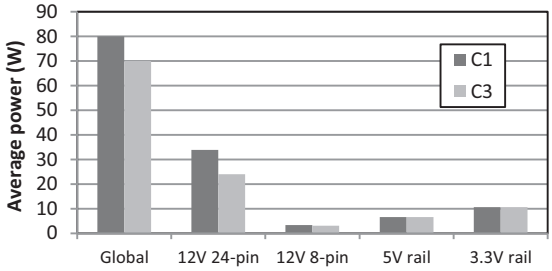

Figure 4. Idle power consumption for C1 and C3 LPI states.

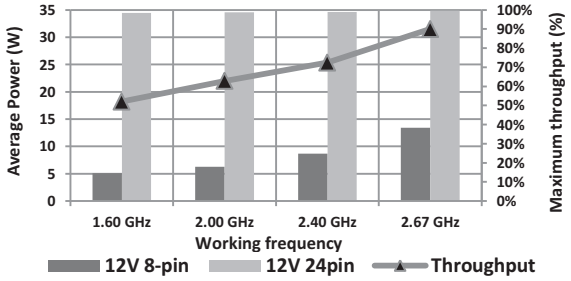

Figure 5. Active power consumption of the 12V ATX rails for a subset of available frequencies and related maximum throughput.

Figs. 6-9 show the instantaneous power of 12V 8-pin and of 12V 24-pin during the reception and forwarding process of two bursts. A burst of 500 IP packets with a size of 64B arrived every 2ms at a NIC, which wake-up the CPU by sending an interrupt signal. Then, when the CPU is active, it starts to

process the incoming packets [13]. In detail, Figs. 6 and 7 show the instantaneous power for 2.67GHz and 1.60GHz working frequencies at C1 state level respectively, while Figs. 8 and 9 show the instantaneous power in the C3 state case for the same working frequencies. After a transition interval in all cases, the 12V 8-pin rail's power consumption during the packet processing is constant at the active power already showed in Figure 5 (where no active-idle transition happens).

Regarding the transition period, it is worth noting that the power-on cost is proportional to the working frequency. In detail, decreasing frequency decreases the power spike during wake-up; as previously sketched, this effect is due to the fact that the Core i5 utilizes DVFS mechanisms to implement the various working P-States, which limit the maximum power adsorption by the CPU. Moreover, in the C3 cases, we can observe another power-on spike during the idle-active interval, which corresponds to the additional processes needed to restore the Core's cache. We estimated a gap of about 100 µs between the idle-active transition times of C1 and C3, which correspond to the same gap achieved by measuring the packet latency with the N2X router tester.

Considering the 12V 24-pin rail, the power supplied in the C1 case remains constant except during the transitions, where we can see two opposite spikes, which do not influence the average power consumption. Otherwise, in the C3 cases, after the a burst service time, the 12V 24-pin rail power progressively decrees to the idle power consumption, since the L1 instruction cache, L1 data cache, and L2 cache are turned off progressively [21]. From the obtained results, it is obvious that grouping packets in burst obviously lead fewer transitions and consequently less energy waste due to the power spikes, at price of adding a buffering delay. However, power spikes can be reduced by slowing down the working frequency. In this respect, the joint adoption of AR and coalescing policies can improve the energy efficiency of devices.

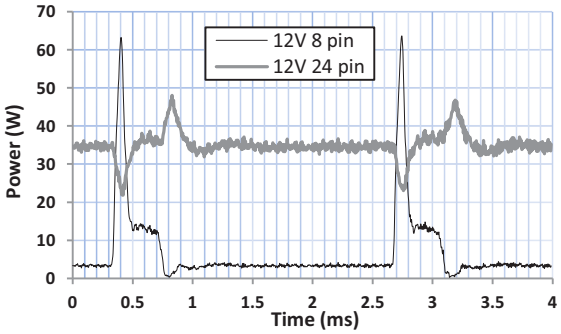

Figure 6. Instantaneous power during the reception and forwarding process of two IP bursts at 2.67GHz and C1 idle state

B. *Weight of power-on transitions*

Starting from the measurement samples collected according to burst sizes and C-state and P-state configurations, from probe on the 12V 8-pin rail, we obtained an estimate of the energy transition cost of the CPU. In this respect, Figure 10 shows the wake-up energy cost estimated by the power consumption samples shown in the previous section. The transition energy due to the C3-C0 transition is higher than the C1-C0 one, since energy transition cost increases as deeper is the C-state entered [15]. Observing the energy transition costs here reported (in the order of magnitude of mJ), one could wrongly conclude that power spikes can be neglected when developing a power consumption model. However, as we will see in the following of this section, under realistic traffic patterns power spikes can have an important weight on the overall power consumption.

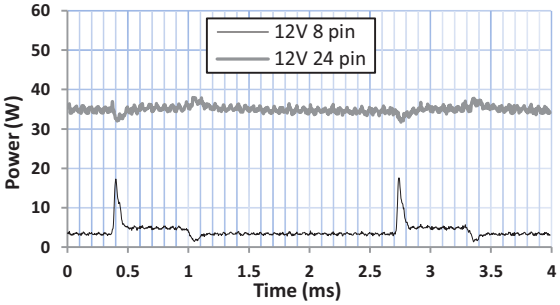

Figure 7. Instantaneous power during the reception and forwarding process of two IP bursts at 1.60GHz and C1 idle state.

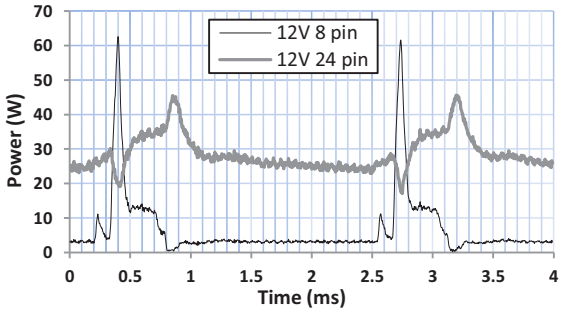

Figure 8. Instantaneous power during the reception and forwarding process of two IP bursts at 2.67GHz and C3 idle state.

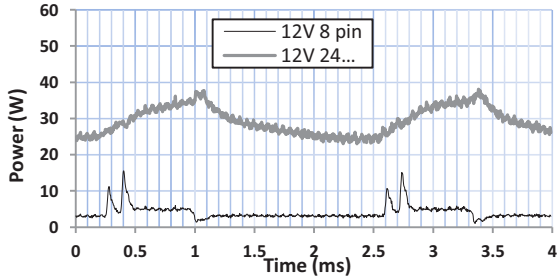

Figure 9. Instantaneous power during the reception and forwarding process of two IP bursts at 1.60GHz and C3 idle state.

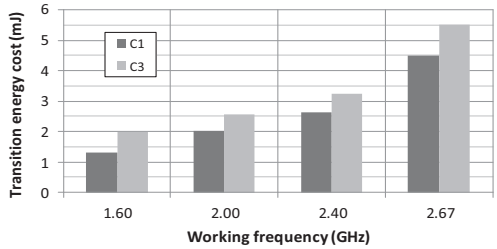

Figure 10. Energy transition costs for a subset of working frequency.

In order to characterize the impact of the wake-up transitions, we perform further measurements on the average power consumption for different burst size and average packet rate of incoming traffic according different ACPI

configurations. Since the 12V 24-pin rail power is also affected by idle-active transitions when the C3 state is configured, we decided to take the aggregated power supplied by both the 12V rails in to account. In more detail, such aggregated power can be considered as the total (dynamic) power consumption of the Core i5 processor plus a constant power offset, due to other SR internal components. Moreover, having the performance results according to different level of traffic burstiness gives us the possibility of estimating how much energy might be potentially saved, if a packet coalescer would be applied to optimize LPI performance.

Figs. 11-14 show the average power consumption results obtained with an incoming packet rate of 1, 10, 100, and 500 kpps, respectively. All these packet rates can be handled by all the used frequency without losing any packets. We chose these values of incoming traffic rates, since for higher volumes transitions are much rarer, and the power consumption consequently depends only on the working frequency.

Fixed the C-state, as the number of packets in a burst increases, the energy consumption provided by different working frequencies tends to become very similar, especially in the case of lower packet rates. In fact, as depicted in Figure 11, a power consumption level near to the minimum possible value can be reached with bursts of only 10 packets or more.

At the same time, AR (i.e., P-states in the ACPI terminology) can be jointly exploited to save energy, especially for higher data rates. Observing Figure 13, it is evident that a joint use of burst aggregation and working frequency decrease allow cutting the energy absorption in a very effective way.

For packet rates equal or higher than 500 kpps (Figure 13), power consumption mainly depends on the processor speed, and the eventual usage of a packet coalescer seems to provide no energy consumption reductions. Notwithstanding deeper sleep states (i.e., C3) are usually thought as more energy-efficient, when incoming traffic rate increases they do not impact on energy consumption in an evident manner. This effect obviously depends on the burst interarrival times, which become too short for allowing the Core to enter in the LPI mode.

Finally, we evaluated the weight of power spikes for waking up from LPI modes. To this purpose, we compare the collected experimental results with the ones obtained by several GreenSim simulations [23]. The simulations were performed by using the same traffic rates, burst sizes, and energy consumption levels of experimental measurements, but the GreenSim simulator was configured in order to neglect power consumption spikes during idle-active transitions[1].

From the difference between experimental results and simulations reported in Figure 15, it is possible to evaluate the contribution of the wake-up process on the overall Core energy consumption, when working under different incoming traffic patterns. As one can expect, the simulations exhibited, in all the analyzed cases, lower energy consumption levels than the ones measured in the experimental testbed. It is a clear indication of the no-negligible overhead introduced by transitions consumption spikes. Observing the obtained results that exhibit differences up to 15%, we can conclude that wake-up spikes have a significant role in the operating energy consumption of networking hardware, and cannot be neglected in performance evaluation and design activities.

In more detail, Figure 15 outlines how the average error is

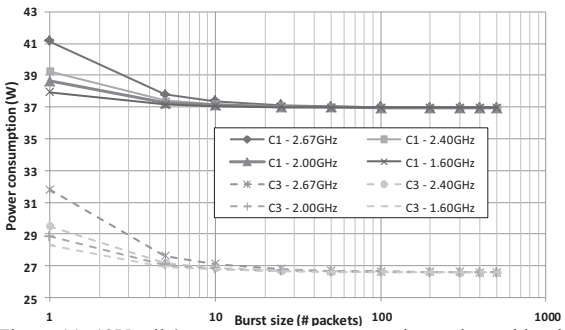

Figure 11. 12V rails' average power consumption estimated by the model at a fixed rate of 1kpps for various burst size according for different working frequency and idle level.

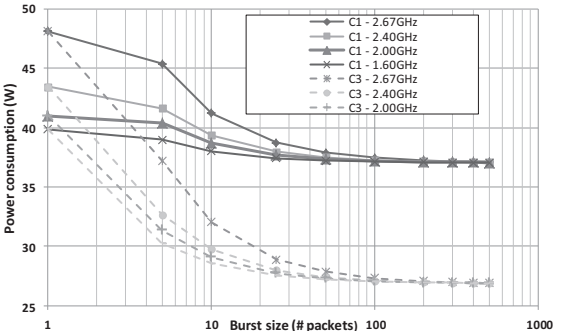

Figure 12. 12V rails' average power consumption estimated by the model at a fixed rate of 10kpps for various burst size according for different working frequency and idle level.

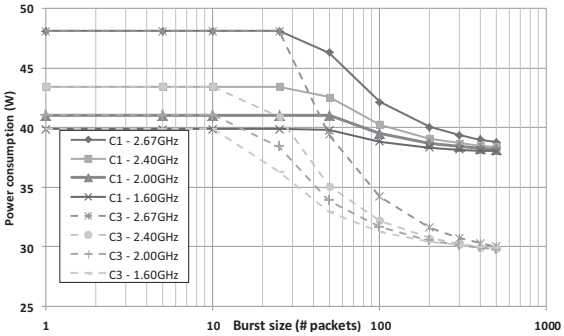

Figure 13. 12V rails' average power consumption estimated by the model at a fixed rate of 100kpps for various burst size according for different working frequency and idle level.

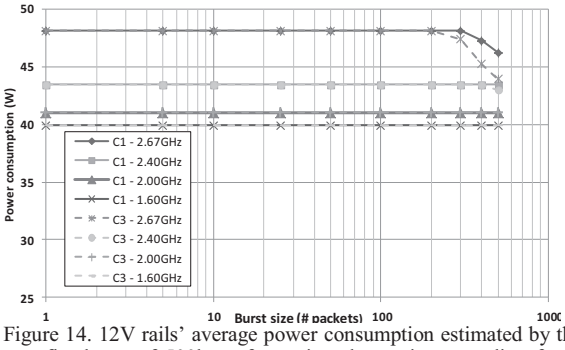

Figure 14. 12V rails' average power consumption estimated by the model at a fixed rate of 500kpps for various burst size according for different working frequency and idle level.

[1] When GreenSim is configured to consider also transition consumption spikes, it provides a maximum accuracy error lower than 0.5%.

larger in C3 since, as previously sketched, the transition cost increases as deeper is the LPI mode. For lower traffic rates (e.g., 1 and 10 kpps), the error tends to decrease according burst sizes: this behavior is clearly due to the lower number of transitions (and then of consumption spikes) occurring when traffic burstiness increases. On the contrary, with higher traffic rates, the error remains almost negligible up to a certain burst size value, which depends on the incoming rate (e.g., 50 pkts for 100kpps and 200 for 500kpps). Here, the difference between experimental and simulative results is caused by the fact that, in experimental tests, the CPU cannot enter into LPI modes, given the shorter burst interarrival times.

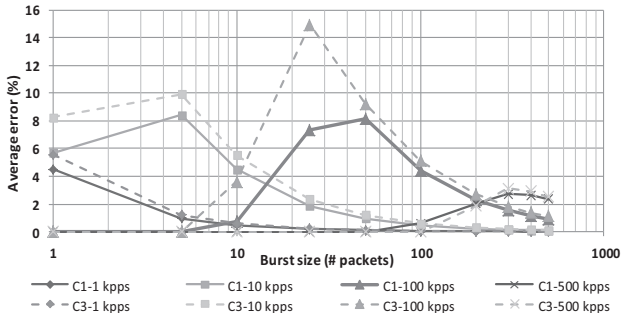

Figure 15. Average error between experimental results and GreenSim ones.

IV. CONCLUSIONS

In this paper, we focused on a deep and experimental analysis of the efficiency of low power idle techniques when applied to packet processing engines of network devices. To this purpose, we set up a complex testbed that allowed us to perform several measurements on energy- and network-performance indexes. The results that we collected in this measurement campaign allowed us not only (*i*) to provide an in-depth energy consumption profiling of SR data-plane, but also (*ii*) to clearly outline energy costs due to the use of low power idle techniques. Among other interesting aspects, we completely characterized the energy consumption due to wakeup transitions, which may cause instantaneous consumption spikes higher than 4 times the power energy requirement when active, and may impact up to 15% of the overall energy consumption of the Core when forwarding traffic.

ACKNOWLEDGMENTS

This work has been supported by the ECONET (low Energy Consumption NETworks) project, co-funded by the European Commission under the 7th Framework Programme (FP7), Grant Agreement no. 258454, and by the GreenNet project funded by the Italian Ministry of Research and Education under the FIRB "Fututo in Ricerca" program.

REFERENCES

[1] R. Bolla, R. Bruschi, A. Carrega, F. Davoli, D. Suino, C. Vassilakis, A. Zafeiropoulos, "Cutting the Energy Bills of Internet Service Providers and Telecoms through Power Management: an Impact Analysis," Computer Networks, Elsevier, vol. 56, no. 10, pp. 2320–2342, Jul. 2012.

[2] Global e-Sustainibility Initiative (GeSI), "SMART 2020: Enabling the Low Carbon Economy in the Information Age", http://www.theclimategroup.org/assets/resources/publications/Smart2020Report.pdf

[3] S. Nedevschi, L. Popa, G. Iannaccone, D. Wetherall S. Ratnasamy, "Reducing Network Energy Consumption via Sleeping and Rate-Adaptation", *Proc. of the 5th USENIX Symp. on Net. Sys. Des. and Impl. (USENIX)*, San Francisco, CA, USA, Jun. 2008, pp. 323-336.

[4] 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 Comm. Surveys and Tutorials (COMST), vol. 13, no. 2, pp. 223–244, Second Quarter 2011.

[5] R. Bolla, R. Bruschi, A. Carrega, F. Davoli, "Green Net. Technologies and the Art of Trading-off," *Proc. of the 2011 IEEE Infocom Work. on Green Comm. And Net. (GCN)*, Shangai, China, Apr. 2011.

[6] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, J. A. Maestro, "IEEE 802.3az: the road to energy efficient ethernet," *IEEE Comm. Magazine,* vol. 48, no. 11, pp. 50-56, Nov. 2010.

[7] P. Reviriego, J. A. Hern, D. Larrabeiti, and J. A. Maestro "Performance evaluation of Energy efficient Ethernet," *IEEE Comm. Letters,* vol. 13, no. 9, pp. 697-699, Sept. 2009.

[8] M. Mostowfi, K. Christensen, "Saving energy in LAN switches: New methods of packet coalescing for Energy Efficient Ethernet," *Proc. of the Internat. Green Comp. Conf. and Work. (IGCC),* pp.1-8, Orlando, FL, USA, Jul. 2011.

[9] S. Herreria-Alonso, M. Rodriguez-Perez, M. Fernandez-Veiga, C. Lopez-Garcia, "A Power Saving Model for Burst Transmission in Energy-Efficient Ethernet," *IEEE Comm. Letters,* vol. 15, no. 5, pp. 584-586, May 2011.

[10] R.S. Tucker, R. Parthiban, J. Baliga, K. Hinton, R.W.A. Ayre, W.V. Sorin, "Evolution of WDM Optical IP Networks: A Cost and Energy Perspective," *J. of Lightw. Tech.*, vol. 27, no. 3, pp. 243-252, Feb 2009.

[11] C.-H. Lien, Y.-W. Bai, M.-B. Lin, "Estimation by Software for the Power Consumption of Streaming-Media Servers," *IEEE Trans. on Instr. and Measur.,* vol. 56, no. 5, pp. 1859-1870, Oct. 2007.

[12] C.-J. Tang, M.-R. Dai, "Dynamic computing resource adjustment for enhancing energy efficiency of cloud service data centers," *Proc. of IEEE/SICE Internat. Symp. on Sys. Integr. (SII)*, Kyoto, Japan, pp. 1159-1164, Dec. 2011

[13] Bolla, R.; Bruschi, R., "PC-based Software Routers: High Performance and Application Service Support," *Proc. ofthe ACM Sigcomm W. on Programm. Routers for Extens. Serv. of Tomorrow (PRESTO'08)*, Seattle, WA, USA, Aug. 2008

[14] The Advanced Configuration & Power Interface (ACPI) standard, version 5.0, http://www.acpi.info/.

[15] "Energy-Efficient Platforms – Considerations for Application Software and Services," Intel Whitepaper, March 2011, http://download.intel.com/technology/pdf/Green_Hill_Software.pdf .

[16] T.-C. Huang, J.-C. Tzeng, Y.-W. Chao, J.-J. Chen, W.-T. Liu, K.-J. Lee, "A Supply-Gating Scheme for Both Data-Retention and Spike-Reduction in Power Management and Test Scheduling," *Proc. of the Internat. Symp. on VLSI Design, Aut. and Test (VLSI-DAT),* Hsinchu, Tawain, pp. 1-4, Apr. 2006.

[17] Intel® Core™ i5 Processor Datasheet, http://download.intel.com/design/processor/datashts/322164.pdf.

[18] Dominion PX-5297 Technical Specifications, http://www.raritan.com/px-5000/px-5297/tech-specs/

[19] Adex Electronics, PEX16LX PCI-e x16 tall extender, http://www.adexelec.com/pciexp.htm#PEX16LX

[20] ATX standard specification version 2.2, http://www.formfactors.org/developer/specs/atx2_2.pdf

[21] H. Amur, R. Nathuji, M. Ghosh, K. Schwan, H.S. Lee, "Idlepower: Application-Aware Management of Processor Idle States," *Proc. of ACM/IEEE Internat. Symp. on High Perf. Distr. Comp.*, Boston, MA, USA, Jun. 2008.

[22] Agilent U2356A 64-Channel 500kSa/s USB Modular Multifunction Data Acquisition, http://www.home.agilent.com/agilent/product.jspx?nid=-33627.384462.0.

[23] R. Bolla, R. Bruschi, A. Carrega, "GreenSim: An open source tool for evaluating the energy savings through resource dynamic adaptation," *Proc. of Internat Symp. on Perf. Eval. of Computer and Telecom. Sys. (SPECTS),* pp. 89-95, Ottawa, Canada, July 2010.