Energy-Aware Equipment for Next-Generation Networks

Raffaele Bolla and Roberto Bruschi
DIST, University of Genoa
Via all'Opera Pia 13,
16139, Genoa, Italy
{raffaele.bolla, roberto.bruschi}@unige.it

ABSTRACT
Besides a more widespread sensitivity to ecological issues, the interest in energy-efficient network technologies springs from heavy and critical economical needs, since both energy cost and network electrical requirements show a continuous growth, with an alarming trend over the past years. In this contribution, we explore and try to evaluate the feasibility and the impact of power management mechanisms on network equipment performance. We discuss and characterize how these energy-aware enhancements can be applied inside the architectures of devices working at access and core networks. Finally, we show how a state-of-the-art multi-Core SW router based on COTS HW can effectively modulate the trade-off between its energy requirements and network performance.

Categories and Subject Descriptors
C.2.6 [**Routers**]: Energy-aware Router Architectures

General Terms
Design, Measurement, Performance, Experimentation.

Keywords
Green networks, Power Management, Energy-aware Network Devices.

1. INTRODUCTION
As the Future Internet is taking shape, it appears that some basic concepts and key aspects should pervade the network infrastructure as a whole, to such extent as to become part of the network design criteria, and to carry across multiple networking domains for the achievement of a general target. One such aspect is that of energy efficiency. It would be remarkable that, at the moment when the role of ICT is becoming of paramount importance in addressing better energy efficiency in the energy production and distribution sector, as well as in fostering energy-awareness in all aspects of production and services, ICT would not apply the same concepts to itself.

In this perspective, power consumption should be considered as a metric for network performance evaluation, much in the same respect as more "classical" indexes are adopted, like throughput, delay, packet loss and jitter. So as avoiding network congestion is a universally shared target, avoiding energy waste in operating the network at the same level of forwarding efficiency should become a similarly widespread concept.

Recent and official studies estimated that ICT industry accounts for approximately 2% of global CO_2 emissions, overcoming even the carbon footprint of aviation. In detail, focusing on telecommunication networks, they are estimated to produce about 0.6% of the global CO_2 emissions.

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 the Telecom Italia network in 2006 has reached more than 2TWh (about 1% of the total Italian energy demand), increasing by 7.95% with respect to 2005, and by 12.08% to 2004. Another explanatory example is represented by British Telecom, which absorbed about 0.7% of the total UK's energy consumption in the winter of 2007, making it the biggest single power consumer in the nation [3]. As outlined in [4], similar trends can be generalized to a large part of the other telecoms and service providers, since they essentially depend on data traffic volume increase, which appears to follow the Moore's law, and new services being offered. To support new generation network infrastructures and related services for a rapidly increasing customer population, telecoms and service providers need an ever larger number of devices, with sophisticated architectures able to perform more and more complex operations in a scalable way.

In such scenario, besides a more widespread sensitivity to ecological issues, the interest on energy efficient networking springs from heavy and critical economical needs, since both energy cost and network electrical requirements show a continuous growth, with an alarming trend over the past years. 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. While this margin is generally reached rarely and over short time periods [5], the overall power consumption in today's networks remains more or less constant with respect to different traffic volumes.

Against such flat energy wastes, the specific challenge for telecoms, network equipment manufacturers and the networking research community nowadays mainly regards the introduction of innovative criteria and technologies, able to save energy by dynamically adapting network capacities and resources to current traffic loads and requirements.

In more detail, novel green technologies for next-generation networks are expected to enable devices and their components (e.g., network interfaces, processing units, etc.) to change their energy and performance profiles by effectively exploiting available HW capabilities, such as "fast" standby optimizations as well as frequency and voltage scaling primitives. On the top of such HW technologies, novel policies and criteria will be certainly required in order to dynamically control the trade-off between energy consumption and performance requirements with respect to estimated volumes of incoming traffic. Despite some interesting scientific contributions (e.g., [6], [7], [8] [9] and [10] among others), green networking performance and optimization remains an open and very interesting issue.

In this paper, our main objective is to explore the feasibility and the impact of introducing advanced power management capabilities in network device architectures for next-generation networks.

2. Green Technologies for NGN devices

Our basic idea is to provide support for architectural solutions and suitable technologies, which allow network devices, or most likely parts of them (i.e., processing units, physical interfaces, etc.), lowering their energy consumption by adapting their "network" performance, or by exploiting "idle" periods to enter standby states.

In this respect, our analysis starts by considering the power management capabilities and mechanisms, today available in the largest part of commodity hardware and general-purpose processors, will be taken into account as a reference starting point. These mechanisms, included in Commercial Off-The-Shelf (COTS) hardware and under rapid development in other hardware technologies (e.g., network processors, ASIC and FPGA circuits), are generally configurable through the Advanced Configuration and Power Interface (ACPI) standard. They respectively allow to minimize power consumption when no activities are performed (namely "idle" or "active standby" optimizations), and to modify the trade-off between performance and energy when the hardware is active and performing operations (namely "power state" optimizations).

These kinds of power management support are generally realized at the hardware layer by powering off sub-components, or by changing the silicon operating frequency and voltage. Specific control applications, namely governors, are needed to dynamically configure such power profiles through the ACPI standard interfaces. In more detail, the specific objective of such SW governors is to optimize the configuration of network devices, in terms of operating frequency and voltage, with respect to their expected performance. Obviously, such energy-aware adaptations have to match current network traffic loads and requirements, which both heavily change according to the network aggregation layers (e.g., access and transport/core layers). Moreover, as highlighted in Figure 1, which depicts an explanatory example of how energy requirements arise in today's networks, both the network access and the core/transport layers cause a non-negligible contribution to the overall carbon footprint.

Thus, It is reasonable to suppose that green technologies need to be specifically shaped and modelled with respect to the different network scenarios and aggregation levels. For example, though, on one hand, an access device (e.g., a DSLAM, a base station, etc.) can be easily supposed to save energy by exploiting standby modes during low activity periods of end-users, on the other hand, the same packet- and flow-level standby optimizations cannot be adopted for active nodes working at the transport and the core layers, given the strict performance requirements and the statistical features of highly aggregated traffic. However, in the latter case, specific standby modes can be explored for reducing the carbon footprint of redundant core and transport nodes and components.

In more detail, these network redundant elements cannot simply switch off, since current hardware technologies do not guarantee restarting times so low to satisfy resilience and availability constraints. Moreover, restarting redundant devices also causes storms of signalling messages and network instability. Thus, future research targets will certainly address the study and the development of specific mechanisms for keeping redundant devices and/or interfaces in very low energy states, while, at the same time, maintaining connectivity and the device presence in the network.

As far as transport and core devices in "active state" are concerned, viable solutions shall exploit other traffic and network features with respect to the access scenario in order to modulate power consumption on the traffic volumes and requirements.

In this respect, specific mechanisms that allow to modulate the processing capacity of network device components, or the bandwidth of link interfaces [14] can be take into account. In addition, these energy-aware enhancements can take a remarkable advantage in being applied inside the architectures of core/transport equipment, which are generally highly modular and distributed in order to divide and conquer traffic loads. If internal modules (e.g., processing units, network interfaces) of such distributed equipment would have its own and independent power management mechanisms, the processing capacity could be finely tuned in a more effective way. When no needed, modules could independently enter in standby modes, and the traffic processing load could be divided among a minimum and suitable number of components, each one with an adaptive processing capacity.

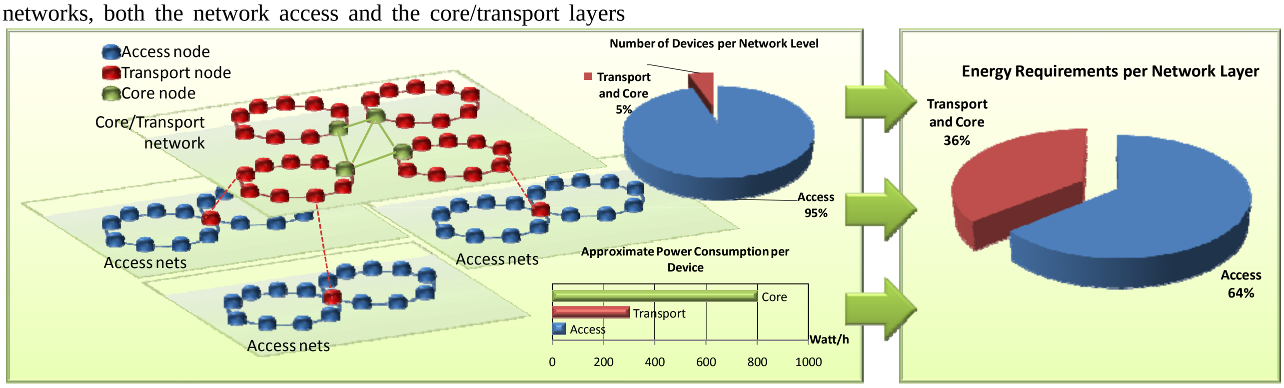

Figure 1. **Core/transport and access layers of a nominal network domain: examples of device density, relative power requirements, and overall power consumption per network layer.**

3. A Proof-of-Concept Case: the SW Router

With the aim of evaluating the impact of the power management features previously introduced, we decided to use a new generation SW Router (SR) based on Linux operating system and COTS HW, able to effectively and scalably exploit multi-CPU/core PC systems. This choice has been guided from the fact that SRs are currently the only network device architectures that include power management capabilities

In detail, the new generation SR architectures include further HW and SW enhancements (as already discussed and evaluated in [11]), which let us go beyond typical performance issues, namely data accessing serialization (i.e., LLTX lock contention) and CPU/core cache coherence, which generally affect SR data-plane operations. The key element of such SR architectural enhancements is the support for "multi-queuing" network adapters, like the Intel PRO 1000 adapters (with 82571 and higher chipsets), which support multiple Tx/Rx Rings and multiple HW interrupts per network interface [13].

Using such innovative network board features, we deployed a new SW architecture/configuration that allows us to optimize SR data-plane performance by reducing memory sharing among CPU/cores and avoiding the LLTX lock contention.

Figure 2, Figure 3 and Figure 4 reports the performance of the SR when a variable number of Cores are used and their working frequency modulated, in terms of maximum throughput, maximum latency and power consumption, respectively.

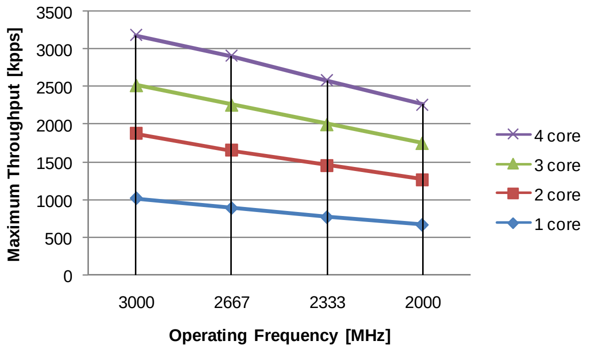

Figure 2. Maximum throughput values with respect to the number of used Cores and their working frequencies.

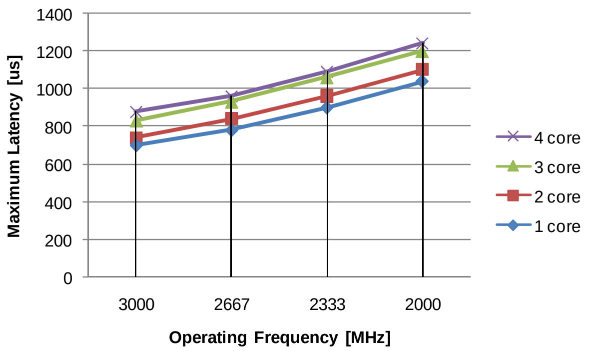

Figure 3. Maximum latency values with respect to the number of used Cores and their working frequencies.

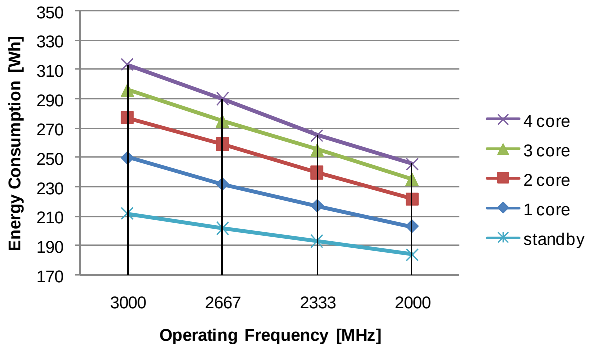

Figure 4. Power consumption values with respect to the number of active Cores and their working frequencies, and power consumption when all the Cores are in standby mode.

The obtained results show that power management mechanisms allow to linearly scale the maximum forwarding rates with respect to power consumption. This linear behavior is maintained not only for different equivalent frequency values, when a single core is used for the SR data plane, but also for a rising number of CPUs/Cores. Accordingly, when the Cores' working frequencies rise, latency values decrease in a quasi-linear way.

4. Conclusions

We focused on the "green" capabilities for next generation network equipment. Our main objective was to analyze, to evaluate the impact of power saving mechanisms, generally included in today's COTS processors, on devices working at the access and the core levels. To this purpose we showed and analyzed the behavior of power management mechanisms of multi-Core SRs on their network performance in terms of throughput and latency. The obtained results show that the adoption of power management mechanisms can effectively be used in order to modulate the trade-off between network performance and energy requirements of next generation equipment.

5. REFERENCES

[1] C. Bianco, F. Cucchietti, G. Griffa, "Energy consumption trends in the next generation access network - a telco perspective," *Proc. 29th Int. Telecomm. Energy Conf (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] BT Press, "BT announces major wind power plans," Oct. 2007, http://www.btplc.com/News/Articles/Showarticle.cfm?ArticleID=dd615e9c-71ad-4daa-951a-55651baae5bb.

[4] ITWales, "Green evangelist to call for big changes in computer use to aid environment at ITWales conference," Nov 2007, http://www.itwales.com/997539.htm.

[5] Sprint IP Monitoring Project. URL: http://ipmon.sprint.com/.

[6] Noguera, J.; Kennedy, I.O., "Power Reduction in Network Equipment through Adaptive Partial Reconfiguration," *Proc. of the 2007 Int. Conf. on Field Programmable Logic and Applications (FPL 2007)*, Aug. 2007, pp. 240-245.

[7] Gupta, M.; Singh, S.; "Using Low-Power Modes for Energy Conservation in Ethernet LANs," *Proc. of the 26th Annual*

IEEE Conf. on Computer Communications (IEEE INFOCOM 2007), Anchorage, Alaska, USA, May 2007.

[8] Nedevschi, S.; Popa, L.; Iannaccone, G.; Wetherall, D.; Ratnasamy, S.; “Reducing network energy consumption via sleeping and rate-adaptation,” *Proc. 5th USENIX Symp. on Networked Systems Design and Implementation*, San Francisco, CA, 2008, pp. 323-336.

[9] R. Bolla, R. Bruschi, A. Ranieri, “Green Support for PC-based Software Router: Performance Evaluation and Modeling”, Proc. *2009 IEEE International Conference on Communications* (ICC 2009), Dresden, Germany, June 2009, to appear.

[10] R. Bolla, R. Bruschi, A. Ranieri, “Performance and Power Consumption Modeling for Green COTS Software Router”, *Proc. 1st International Conference on COMmunication Systems and NETworkS* (COMSNETS 2009), Bangalore, India, Jan. 2009.

[11] Bolla, R.; Bruschi, R., “PC-based Software Routers: High Performance and Application Service Support,” *Proc. of the ACM Sigcomm Workshop on Programmable Routers for Extensible Services of Tomorrow* (PRESTO’08), Seattle, WA, USA, Aug. 2008, pp. 27-32.

[12] K. Argyraki, S. A. Baset, B. G. Chun, K. Fall, G. Iannaccone, A. Knies, E. Kohler, M. Manesh, S. Nedveschi, S. Ratnasamy, “Can Software Router Scale?,” *Proc. of the ACM Sigcomm Workshop on Programmable Routers for Extensible Services of Tomorrow* (PRESTO’08), Seattle, WA, USA, Aug. 2008, pp. 21-26.

[13] Yi, Z.; Waskiewicz, P.J., “Enabling Linux Network Support of Hardware Multiqueue Devices,” *Proc. of the 2007 Linux Symposium*, Ottawa, Canada, June 2007, pp. 305-310.

[14] C. Gunaratne, K. Christensen, S. Suen, B. Nordman, "Reducing the Energy Consumption of Ethernet with an Adaptive Link Rate (ALR)," IEEE Trans. on Computers, vol. 57, no. 4, pp. 448-461, Apr. 2008.