The Potential Impact of Green Technologies in Next-Generation Wireline Networks: Is There Room for Energy Saving Optimization?

Raffaele Bolla and Franco Davoli, DIST-University of Genoa
Roberto Bruschi, National Inter-University Consortium for Telecommunications
Ken Christensen, University of South Florida
Flavio Cucchietti, Telecom Italia
Suresh Singh, Portland State University

ABSTRACT

Recently, network operators around the world reported statistics of network energy requirements and the related carbon footprint, showing an alarming and growing trend. Such high energy consumption can be mainly ascribed to networking equipment designed to work at maximum capacity with high and almost constant dissipation, independent of the traffic load. However, recent developments of green network technologies suggest the chance to build future devices capable of adapting their performance and energy absorption to meet actual workload and operational requirements. In such a scenario, this contribution aims at evaluating the potential impact on next-generation wireline networks of green technologies in economic and environmental terms. We based our impact analysis on the real network energy-efficiency targets of an ongoing European project, and applied them to the expected deployment of Telecom Italia infrastructure by 2015–2020.

INTRODUCTION

Only recently, telecom operators (telcos) and Internet service providers (ISPs) have raised their interest in energy efficiency for wired networks and service infrastructures, making it a high-priority objective. This interest is motivated by the increase in energy prices, the continuing growth of the customer population, the spreading of broadband access, and the expanding services offered. Indeed, the number of new services being offered and the increase in the volume of data traffic follow Moore's law, doubling every 18 months.

To support new-generation network infrastructures and related services for a rapidly growing customer population, telcos and ISPs need an ever larger number of devices, with sophisticated architectures able to perform increasingly complex operations in a scalable way. For instance, high-end routers are increasingly based on complex multirack architectures, which provide more and more network functionalities; historic data from manufacturers' datasheets show continuously raising capacities, by a factor of 2.5 every 18 months [1]. However, silicon technologies improve their energy efficiency at a slower pace following Dennard's law (i.e., by a factor of 1.65 every 18 months) with respect to routers' capacities and traffic volumes.

In the last few years, telcos, ISPs, and public organizations around the world reported statistics of network energy requirements and the related carbon footprint, showing an alarming and growing trend. The Global e-Sustainability Initiative (GeSI) [2] estimated an overall network energy requirement of about 21.4 TWh in 2010 for European telcos, and foresees a figure of 35.8 TWh in 2020 if no green network technologies (GNTs) are adopted. The sole introduction of novel low-consumption silicon technologies clearly cannot cope with such trends and be sufficient for drawing current network equipment toward a greener future Internet.

Moreover, it is well known that network 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 seldom reached, nevertheless the overall power consumption in today's networks is determined by it and remains more or less constant even in the presence of fluctuating traffic loads. This situation suggests the possibility of adapting network energy requirements to actual traffic profiles.

Motivated by these considerations, we recently established a partnership with a group of primary device manufacturers and telcos to launch the ECONET initiative [3], an integrated project

(a)			
	Power consumption (W)	Number of devices	Overall consumption (GWh/year)
Home	10	17,500,000	1533
Access	1280	27,344	307
Metro/transport	6000	1750	92
Core	10,000	175	15
Overall network consumption			1947

(b)		
Home access	Number of customers per DSLAM	640
	Usage of a network access (user up time)	30%
	Link utilization when a user is connected	10%
Core, transport, and metro	Redundancy degree for metro/transport devices (η_t)	13%
	Redundancy degree for core devices (η_c)	100%
	Redundancy degree of metro/transport links (ψ_t)	100%
	Redundancy degree of core device links (ψ_c)	50%
	Link utilization in metro networks	40%
	Link utilization in core networks	40%

(c)			
	Data plane (δ_d)	Control plane (δ_c)	Cooling/power supply (δ_p)
Home	79%	3%	18%
Access	84%	3%	13%
Metro/transport	73%	13%	14%
Core	54%	11%	35%

Table 1. *a) 2015–2020 network forecast: device density and energy requirements in the business as usual (BAU) case, example based on the Italian network; b) traffic and topological data (average figures) for the 2015–2020 prospective network (source: Telecom Italia); c) internal sources of energy consumption (source: the ECONET Consortium).*

The main aim of ECONET is to design and develop innovative solutions and device prototypes for wired network infrastructures (from customer-premises equipment to backbone switches and routers) within 2013.

(IP) funded by the European Commission (EC). The main aim of ECONET is to design and develop innovative solutions and device prototypes for wired network infrastructures (from customer premises equipment to backbone switches and routers) by 2013. The resulting network platforms will adopt GNTs for aggressively modulating power consumption according to actual workloads and service requirements.

In this scenario, the goal of the present article is to evaluate the potential gain to be derived from the application of GNTs in quantitative terms. To this aim, we want to assess how technological solutions able to trade device performance for power consumption can be effectively used in the short term for reducing the carbon footprint and operating expenditures (OPEX) of next-generation wireline networks. In order to make our impact analysis as realistic as possible, we consider the energy efficiency targets of the ECONET project, and apply them to a perspective network of a large-scale telco, which corresponds to an expected deployment of Telecom Italia infrastructure by 2015–2020.

The article is organized as follows. We introduce a reference scenario based on prospective Telecom Italia network development. We briefly describe some of the most promising GNTs we are proposing, along with other contributions by the research community, for disruptively reducing the network carbon footprint. We report an analysis estimating the impact of GNTs in both economical and environmental terms. Conclusions are then drawn.

A Reference Scenario: The Italian Case

The considered reference scenario is shown in Table 1a, which contains the number of devices per network segment and their energy consumption in the business-as-usual case (i.e., the case in which no green enhancements are included in network devices). Both end-user and operator equipment have been taken into account.

We refer to a network with 17.5 million customers, where we assume the presence of broadband-only access technologies, together with suitable overprovisioning in the metro, transport, and core segments.

Devices' energy consumption has been forecast on the basis of present values, high quality specifications (e.g., European Broadband Code of Conduct), and expected "inertial" technological improvements (e.g., Dennard's law [1]). In the devices' energy

The largest part of approaches undertaken is founded on a few basic concepts, which have been generally inspired by energy-saving mechanisms and power management criteria that are already partially adopted in computing systems.

requirements, we included the contribution of site cooling and powering systems, which account for 36 percent of direct device consumption.

Starting from the scenario in Table 1, the per-user average energy requirement consists of about 111 kWh/year, mainly due to home and access networks, for 79 and 16 percent, respectively. Metro/transport and core networks account only for 5 percent, but their joint energy requirement of about 107 GWh/year (about a quarter of the telco's direct energy consumption) can be a convincing driver for reducing the carbon footprint of backbone devices in the near future.

Table 1b defines key parameters that allow synthetically representing the average usage of network devices and links. This has been done by defining the expected (by 2015–2020) average customer up times and loads, the average traffic utilization on metro/transport and core networks, and the number of devices and links that are usually deployed for redundancy purposes. It is worth noting that the traffic load values in Table 1b are significantly larger (and consequently give rise to conservative consumption estimations) than those of the current network and indicated in other studies [1].

Finally, the values in Table 1c, which subdivides the equipment consumption into its main functions/building blocks, outline how the most energy-starving elements in network devices reside in the data plane. In fact, data plane energy shares range between 54 and 84 percent, against 13–35 percent for air cooling and power supply for onboard components, and 3–13 percent for control plane ones. Thus, it is not so surprising that the current GNT proposals in this field are especially devoted to reducing the energy consumption of the network devices' data plane.

GREEN NETWORK TECHNOLOGIES

The largest part of approaches undertaken is founded on a 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 [1], these basic concepts can be classified as *dynamic power scaling* and *smart standby approaches*.

DYNAMIC POWER SCALING

Power scaling capabilities allow dynamically reducing the working rate of processing engines or link interfaces. This is usually accomplished by adopting two basic techniques: *adaptive rate* and Low Power Idle (LPI). The former allows dynamically modulating the capacity of a link or a processing engine in order to meet traffic load and service requirements. The latter forces links or processing engines to enter low-power states when not sending/processing packets and quickly switch to a high-power state when sending one or more packets. These techniques are not exclusive and can be jointly adopted in order to adapt system performance to current workload requirements.

In previous work [4–6], some of the authors of the present article first faced the energy efficiency issue in wireline networks and developed several algorithms that used traffic prediction to put links to low power idle modes. The idea behind these algorithms is simple: the upstream interface on a link maintains a window of inter-packet arrival times. This information is then used to determine the length of time for which an interface can be put to sleep such that, with a high probability, the buffers at that interface will not overflow. The constraint is the non-zero time it takes for the downstream interface to wake up. The results obtained with real traces show that for loads up to 30 percent of link capacity, considerable energy savings can be achieved.

Similarly, the research of Christensen *et al.* has specifically addressed how to reduce direct energy use of Ethernet links, and has contributed to the development of the IEEE 802.3az standard. In [7] they first explored the notion of an adaptive link rate (ALR) for Ethernet, whereby a link would operate at a low data rate during periods of low utilization and at high data rate only for high utilization periods. Given that most links are highly underutilized, with ALR most Ethernet links could operate at a low data rate (and thus reduce energy consumption compared to operation at a high data rate) most of the time [8].

An implementation of ALR would entail an Ethernet interface having two physical layer implementations and switching between them. The time to switch between physical layer implementations was deemed to be a major issue, resulting in an alternative LPI approach [9] proposed by Intel. LPI is the approach specified in the emerging 802.3az standard, and currently allows a 10 Gb/s link to wake up in less than 3 μs.

In [10], Bolla *et al.* analyzed and empirically modeled the energy modulation capabilities of processing engines in Linux-based software routers equipped with general-purpose and multicore processors that already include LPI and adaptive rate primitives. The results achieved were obtained by evaluating several hardware architectures, and they suggest that such technologies permit the trade-off between power consumption and network performance to scale almost linearly.

In [11], the authors extended their approach by introducing a control framework for optimally tuning LPI and adaptive rate mechanisms in order to statistically meet current traffic loads and service requirements. The results obtained on real traffic traces show that energy gains up to 60 percent are feasible.

Working on similar platforms, Intel researchers [12] especially focused on LPI primitives and performed a comprehensive study of the impact of transition times on LPI as a function of load. They showed that as the transition times shrink from the value of 10 ms to 1 ms and then further to 100 μs, the time spent sleeping at 30 percent load goes from 0 at transition time of 10 ms to 40 percent when this time is 1 ms, and to 70 percent when the transition happens in 100 μs.

SMART STANDBY

Sleeping and standby approaches are founded on power management primitives that allow devices or parts of them turning themselves almost completely off and entering very low energy states, while all their functionalities are frozen.

The widespread adoption of this kind of energy-aware capability is generally hindered by the necessity of maintaining the "network presence"

Standby (α)	85%
Performance scaling (β)	50%
Network-wide control (χ)	20%
Air cooling/power supply (γ)	15%

Table 2. *Green efficiency degrees targeted in the ECONET project.*

of the device. Network hosts and devices must maintain network connectivity or they will literally "fall off the network," become unreachable, and network applications and services will fail.

In [13, 14], Christensen *et al.* first explored the idea of using a proxy to "cover" for a network host and thus allow it to go to sleep. Specifically, they addressed requirements for a proxy to be able to respond to ARP packets on behalf of a sleeping host (to maintain reachability from the router), and to respond to other protocol and application messages as needed to maintain full network presence. A specific focus of their research was on peer-to-peer (P2P) protocols and how they might be proxied to allow for a PC sharing files to sleep most of the time. The proxying requirements they developed led to the creation of the European Computer Manufacturers Association (ECMA) TC38-TG4 standard. Apple has recently announced a product with green proxying capabilities.

Nonetheless, the development of such network-specific low-energy modes is a fundamental key factor for reducing the carbon footprint inside networks as well, since it will allow switching some links, entire network devices, or parts thereof to a sleep mode in a smart and effective way. This is the main idea behind emerging approaches to network control, routing, and traffic engineering [15], which aim at dynamically putting network portions to sleep during light utilization periods, in order to minimize the energy requirements of the overall network while meeting the operational constraints and current workloads.

Impact Analysis

This section is organized as follows. We estimate the impact of GNTs on energy consumption of future network devices. We try to assess how much these technologies can really be exploited. Finally, we aggregate all the estimated data for giving a complete overview of the potential impact of GNTs in terms of energy savings, and reduction of both CO_2 emissions and OPEX.

The Device Energy Profiles

Energy-aware devices are meant to dynamically save energy by applying the adaptive capacities of GNTs introduced earlier.

Due to the heterogeneity of functional and performance requirements of the various kinds of networking equipment, different exploitations of green optimizations and technologies can be reasonably expected. Anyway, the joint adoption of these green optimizations will lead future network devices to have a composite energy profile, depending on their current workload and on their actual usage status. From a simplified point of view, we characterized such composite energy profile for a generic network device by considering the following power breakdown:

- P_{full} is the power absorption of an entirely busy device, and corresponds to no green optimization.
- P_{idle} is the power absorption of a device that is active, but not performing any operations.
- $P_{standby}$ is the power absorption of a device in low-energy standby modes; only basic operations to maintain the network presence are performed.

We assume that, as experimentally demonstrated on software router architectures in [10], power consumption values between the "idle" and "full" cases vary linearly with respect to the actual workload.

The specific values of P_{idle} and $P_{standby}$ depend on the efficiency of power scaling and standby primitives that will be adopted in future network devices.

In order to provide a solid basis for our estimation, we decided to use the same efficiency degrees that are targeted by the ECONET project. Table 2 reports such minimum efficiency targets. Taking the full consumption of the device as a term of comparison, the α and β parameters represent how much energy can be saved by the data plane hardware in standby and idle modes, respectively. The γ parameter weighs an indirect gain on the device's power supply and air cooling due to the use of power scaling primitives. Finally, χ is meant to represent the impact of network-wide control strategies (e.g., traffic engineering) for optimizing the energy-aware configuration of the overall network at transport and core levels.

Starting from these considerations and the above definitions, we expressed P_{idle} and $P_{standby}$ as follows:

$$P_{idle} = P_{full}[(1 - \beta)\,\delta_d + \delta_c + (1 - \gamma)\,\delta_p] \quad (1)$$

$$P_{standby} = P_{full}[(1 - \alpha)\,\delta_d + \delta_c + (1 - \alpha)\,\delta_p] \quad (2)$$

where δ_d, δ_c, and δ_p are defined in Table 1c.

Focusing on Eqs. 1 and 2, we can outline how the power requirements of the data plane in idle and standby states are scaled down by α and β, respectively, while those of the control plane elements are maintained constant and equal to the full load condition. This is due to the fact that energy-aware network devices are meant to keep their presence in the network by sending, elaborating, and receiving signaling and control protocol packets.

Figure 1 shows the estimated energy profiles for energy-aware future devices following the 2015–2020 forecast in Table 1a.

Power Scaling and Standby Exploitation

In order to estimate the real energy-savings coming from a massive adoption of GNTs, we need to know the device energy profiles, and we have also to estimate which and how much standby and power scaling primitives would be exploited in real operating scenarios for both access/home networks and transport/core ones.

The widespread adoption of this kind of energy-aware capability is generally hindered by the necessity to maintain the "network presence" of the device. Network hosts and devices must maintain network connectivity or else they will literally "fall off the network."

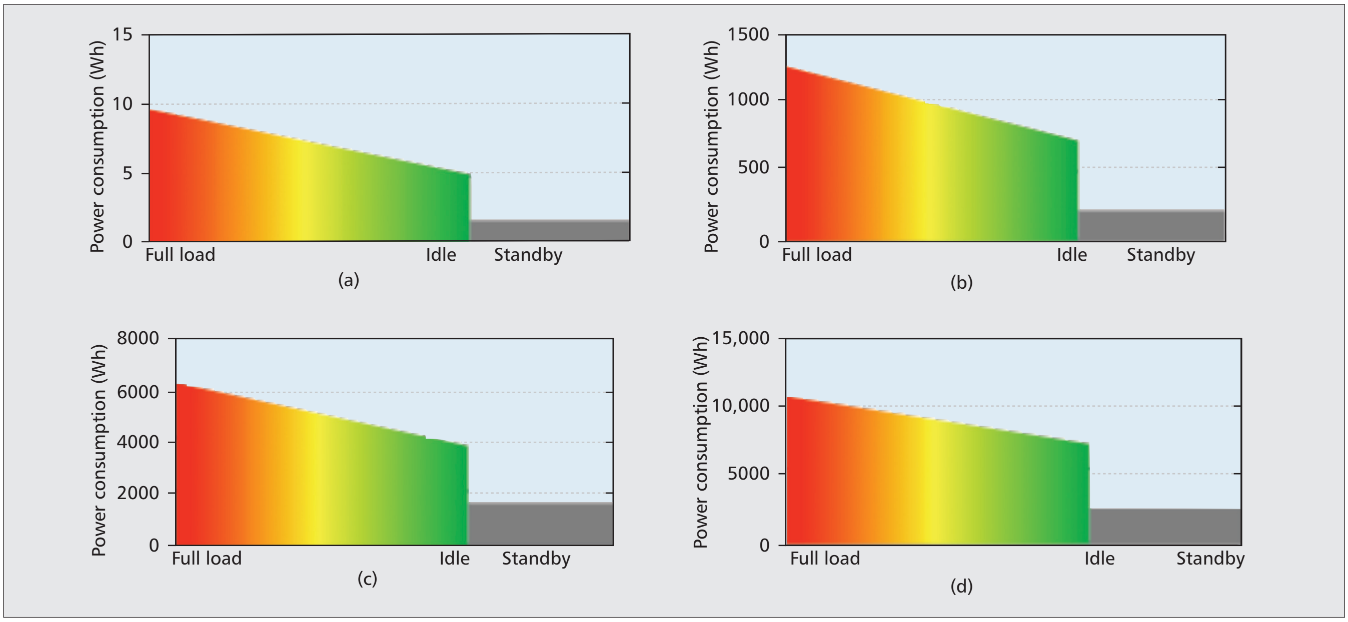

Figure 1. *a) Energy profiles of next generation green home gateways; b), digital subscriber line access multiplexers (DSLAMs); c) metro/transport; d) core devices. The energy profiles have been obtained by using Eqs. 1 and 2 and the input data in Tables 1 and 2.*

Focusing on the standby capabilities, it is reasonable to expect a different exploitation way in the home and access devices with respect to the transport/core ones.

Regarding end-users and network access, standby modes can be thought to suitably follow the end-user behavior — for example, by temporarily putting home gateways or digital subscriber line (DSL) links in standby states when not directly used or not needed. Thus, the average time, in which home devices, or links in an access device can be suitably put in standby modes, mainly depends on end-users' up-times. Recalling the average end-user up-time (i.e., 30 percent, which corresponds to about 7 h/day) from Table 2, we can deduce that home gateways and access links can be put in standby for 70 percent of the time.

Regarding backbone network infrastructures, we have to take tighter performance and operational constraints into account. Core, transport and metro network devices have usually to work on large volumes of traffic, and to guarantee top performance in terms of quality of service (QoS), network reliability, and so on.

Here, while real operating devices and parts of them cannot be put in standby modes, there is a large set of redundant devices and links that is left powered on only to fast recover from faults. Our main proposal is that devices working at these levels should include specific standby support to allow redundant hardware (i.e., entire devices, line cards, etc.) to be woken up upon fault detection with very fast recovery times. For these reasons, such elements cannot be simply shut off, but they need to maintain their network control and signaling activities to always have up-to-date network data and information.

According to Table 1, at the core network level, each active device is coupled with a redundant copy, and furthermore each active device has a quarter of its network links still in redundancy. Thus, the actually operating hardware elements in a core network correspond to about 37.5 percent[1] of the deployed (and powered) ones. Similar remarks can be made also for transport devices, where the redundancy degree is however usually less than the one at the core level. Still according to Table 1, redundant transport/metro nodes are about 13 percent of the deployed ones, while each active device usually has a redundant copy for each of its operating links. In such case, the percentage of operating hardware results being about 43.5 percent of the network deployment. Moreover, green traffic engineering and routing mechanisms can further reduce the number of active links and devices [15], while meeting desired service requirements and traffic loads. Recent studies [15] demonstrate that such policies can lead, on average, to put more than 20 percent of the active devices to sleep. As shown in Fig. 2, by applying this last "energy gain" we obtained the average shares of time with reference to the access and home networks, and the shares of devices and parts of them with reference to the transport and core networks where standby primitives can be exploited fruitfully.

Regarding the exploitation of power scaling mechanisms, they may be applied to active devices, and allow modulating power consumption with respect to the processed traffic loads. Starting from the discussions earlier, we can easily deduce that the average power need of a generic network device using power scaling primitives can be estimated to be equal to $(1 - \rho)P_{idle} + \rho P_{full}$, where ρ is the average link (and device) utilization reported in Table 1b for each network segment.

The Overall Impact

This impact analysis has been performed by considering both the devices working inside the telco network and the customer premises equipment (thus, our estimation covers the whole wireline network, from home gateway to core routers). For all such devices, we used the ener-

[1] *This value has been obtained as $(1 - \psi_c/2)\,\eta_c/2$ with reference to Table 1b.*

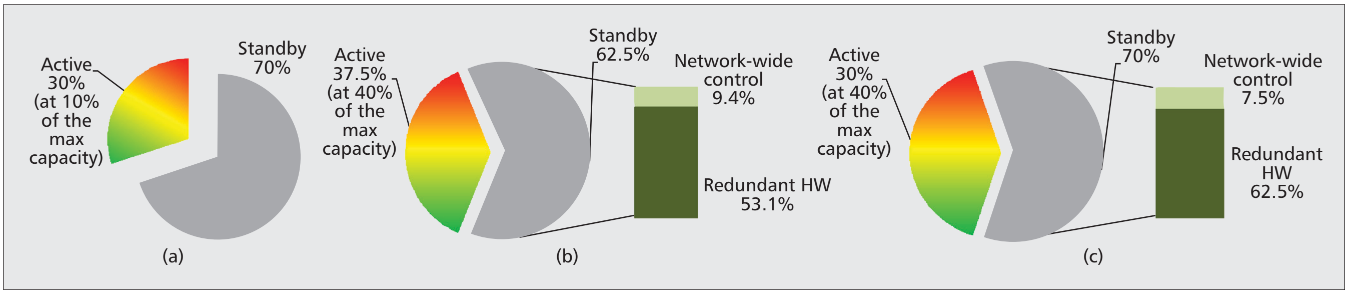

Figure 2. *a) Standby exploitation shares for access and home devices; b) for metro and transport ones; c) for core routers. In case a), such times correspond to the user activity profiles (Table 1b). In the b) and c) cases, standby times arise from both the share of redundant hardware and the devices (and/or parts of them) put in standby modes by green traffic engineering and routing mechanisms. The figures also report the average utilization of devices when active.*

	Full load power consumption (Wh)	Number of devices	Overall full consumption (GWh/year)	Percentage gains	Energy gains (GWh/year)
Home	10	17,500,000	1533	70%	1060
Access	1,280	27,344	307	70%	213
Metro/transport	6,000	1,750	92	54%	49
Core	10,000	175	15	58%	9
			Overall gain	**68%**	
Total BAU (GWh/year)		**1947**	**Total gains with green technologies (GWh/year)**	**1331**	

Table 3. *Impact of green technologies on the 2015–2020 perspective network in terms of energy savings.*

gy profiles in Fig. 1 by weighing them with the exploitation shares in Fig. 2.

Table 3 reports the overall impact of GNTs in terms of energy consumption reduction with respect to the BAU scenario, which was introduced in Table 1a. The obtained values show that the energy requirements of the reference network can be sensibly scaled down by about 1331 GWh/year (which roughly equals 956 ktons/year of CO_2 emissions), corresponding to 68 percent of BAU energy consumption — a significant figure, comparable (and additional!) to the improvement obtainable by the sole increase in hardware efficiency (as predicted by Dennard's law) with respect to the capacity increase of network devices.

This gain especially arises from the customer side, where, by considering only the savings at the home gateways, we obtained a potential reduction equal to 1060 GWh/year, which corresponds to about 70 percent with respect to the BAU requirements. The energy gain would be much larger if we also considered the potential additional savings of GNTs applied to other customer devices like set-top-boxes, VoIP phones, and PCs.

Also, the telco side energy savings, even though corresponding to about a quarter of those at the customer side , show an almost surprising impact: a total gain equal to 271 GWh/year (a reduction of more than 65 percent with respect to the BAU scenario). Using the EIA[2] forecasting on 2020 energy cost, this energy saving can be directly translated to a total OPEX reduction of nearly $33 million/year for the reference Telco. This considerable OPEX reduction especially arises from devices working at the access level, which, thanks to their numerousness, account for 78.6 percent of the overall Telco's gain. Also GNTs applied to metro/transport and core network devices, which account for 18.1 percent and 3.3 percent, respectively, provide a non negligible impact in terms of both OPEX saving and carbon footprint reduction, especially when compared to their number in the network.

Finally, Fig. 3 shows how the overall energy gain varies according to different values of efficiency degrees for standby and power scaling technologies. These results outline that standby primitives have a sensibly higher impact on the final gain than the power scaling ones.

Conclusion

We propose a forecast of the potential impact of GNTs, if massively adopted in a large-scale operator network by 2015–2020. Our estimates consider both the equipment inside the telco network and the customer premises equipment, and we target the same short-term energy efficiency goals as the ECONET European project. The figures resulting from our calculations on this basis outline that GNTs can allow saving about 68 percent of energy requirements of the overall network and an OPEX reduction for the single reference telco of about $33 million/year, due to a gain of 271 GWh/year in the energy consumption of wireline infrastructures. These figures are the upper bounds of what could be attained, whereas the

[2] *EIA, U.S. Energy Information Administration, http://www.eia.doe.gov/.*

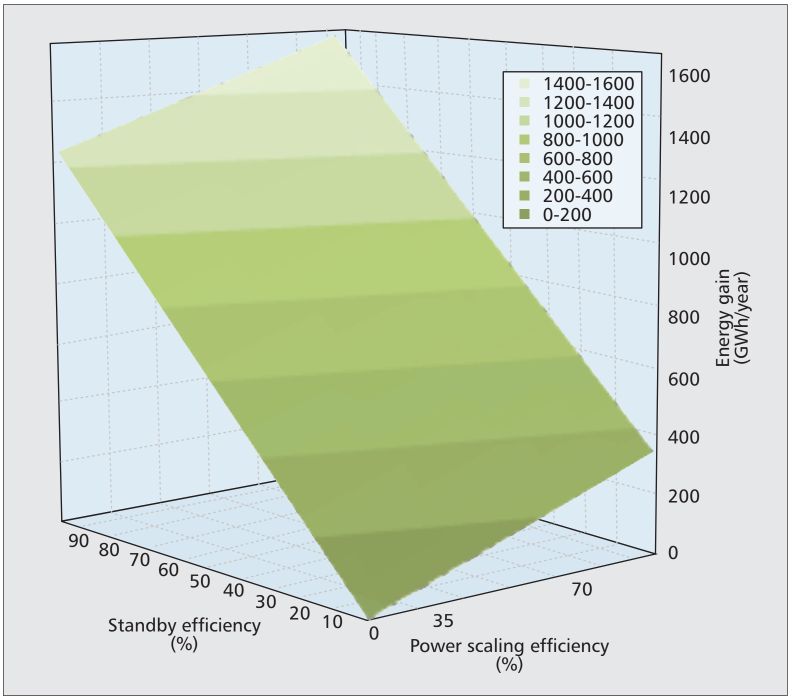

Figure 3. *Overall energy gain for the whole network (including both customer and telco sides) with respect to different efficiency degrees of standby and power scaling mechanisms.*

actual savings will depend on the extent of the adoption of the new methodologies. The technology to achieve these savings is there. At this point, the only question that remains open is whether these numbers will be impressive enough to convince research and industrial communities to set forth the actual development of a greener Internet. Whereas telcos should be encouraged by the potential OPEX reduction, customers and regulators would also support the process once provided with a simple and clear explanation of their own cost savings, especially if guaranteed by certification authorities (e.g., Energy Star[3]).

ACKNOWLEDGMENTS

This work was supported by the European FP7 ECONET project under grant agreement no. 258454.

REFERENCES

[1] R. Bolla *et al.*, "Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures," *IEEE Commun. Surveys and Tutorials*, vol. 13, no. 2, 2nd qtr. 2011, pp. 233–44.

[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] The "low Energy COnsumption NETworks" (ECONET) Project, IP project funded by the EC in ICT-Call 5 of the 7th Framework Programme, grant agreement no. 258454, http://www.econet-project.eu.

[4] M. Gupta and S. Singh, "Greening of the Internet," *Proc. ACM SIGCOMM '03*, Karlsruhe, Germany, Aug. 2003, pp. 19–26.

[5] M. Gupta and S. Singh, "Dynamic Ethernet Link Shutdown for Power Conservation on Ethernet Links," *Proc. IEEE ICC '07*, Glasgow, Scotland, June 2007.

[6] M. Gupta, S. Grover, and S. Singh, "A Feasibility Study for Power Management in LAN Switches," *Proc. IEEE ICNP '04*, Berlin, Germany, Oct. 2004.

[7] C. Gunaratne, K. Christensen, and B. Nordman, "Managing Energy Consumption Costs in Desktop PCs and LAN Switches with Proxying, Split TCP Connections, and Scaling of Link Speed," *Int'l. J. Network Mgmt.*, vol. 15, no. 5, Sept. 2005, pp. 297–310.

[8] C. Gunaratne *et al.*, "Reducing the Energy Consumption of Ethernet with an Adaptive Link Rate (ALR)," *IEEE Trans. Computers*, vol. 57, no. 4, Apr. 2008, pp. 448–61.

[9] R. Hays, "Active/Idle Toggling with 0BASE-x for Energy Efficient Ethernet," presentation to the IEEE 802.3az Task Force, Nov. 2007, http://www.ieee802.org/3/az/public/nov07/hays_1_1107.pdf

[10] R. Bolla, R. Bruschi, and A. Ranieri, "Green Support for PC-based Software Router: Performance Evaluation and Modeling," *Proc. IEEE ICC '09*, Dresden, Germany, June 2009.

[11] R. Bolla, R. Bruschi, and F. Davoli, "Energy-Aware Performance Optimization for Next Generation Green Network Equipment," *Proc. ACM SIGCOMM PRESTO '09*, Barcelona, Spain, Aug. 2009.

[12] S. Nedevschi *et al.*, "Reducing Network Energy Consumption via Sleeping and Rate-Adaptation," *Proc. USENIX NSDI '08*, San Francisco, CA, USA, Apr. 2008.

[13] K. Christensen and F. Gulledge, "Enabling Power Management for Network-Attached Computers," *Int'l. J. Network Mgmt.*, vol. 8, no. 2, Apr. 1998, pp. 120–30.

[14] M. Jimeno, K. Christensen, and B. Nordman, "A Network Connection Proxy to Enable Hosts to Sleep and Save Energy," *Proc. IEEE IPCCC '09*, Dec. 2009, pp. 101–10.

[15] J. Restrepo, C. Gruber, and C. Machoca, "Energy Profile Aware Routing," *Proc. IEEE GreenComm '09*, Dresden, Germany, June 2009.

BIOGRAPHIES

RAFFAELE BOLLA [M] (raffaele.bolla@unige.it) received his Laurea degree in electronic engineering in 1989 and his Ph.D. degree in telecommunications in 1994, both from the University of Genoa, Italy. He is currently an associate professor at the Department of Communications, Computer and Systems Science (DIST) of the University of Genoa. His main current research interests are in the future Internet and green networking. He is the Principal Investigator of the ECONET project.

ROBERTO BRUSCHI [M] (roberto.bruschi@cnit.it) received his M.Sc. degree in telecommunication engineering in 2002 and his Ph.D. degree in electronic engineering in 2006 from the University of Genoa, Italy. He is currently a researcher at the National Inter-University Consortium for Telecommunications (CNIT) at the University of Genoa Research Unit. His main research interests include the future Internet, green networking, and software routers.

KEN CHRISTENSEN [SM] (christen@csee.usf.edu) is a professor in the Department of Computer Science and Engineering at the University of South Florida. He received his Ph.D. in electrical and computer engineering from North Carolina State University in 1991. His primary research interest is in green networks. He is a licensed Professional Engineer in the state of Florida, and a member of ACM and ASEE.

FLAVIO CUCCHIETTI (flavio.cucchietti@telecomitalia.it) is with Telecom Italia R&D Center in Turin, Italy. He is managing innovative projects on infrastructural and powering solutions for the next-generation network. He is currently a member of many standardization groups acting on energy efficiency for telecommunications and IT, and network infrastructure. He co-chairs the Energy Efficiency Inter Operators Collaboration Group (GeSi-EE IOCG). He is a board member of the Italian standardization institute in the electrical, electronic and telecommunication fields (CEI).

FRANCO DAVOLI [SM] (franco.davoli@unige.it) received his Laurea degree in electronic engineering in 1975 from the University of Genoa, Italy. He is currently a full professor of telecommunication networks at DIST, University of Genoa. His current research interests are in multiservice networks, green networking, wireless networks, and multimedia communications and services. He is a member of the CNIT Scientific Board.

SURESH SINGH [M] (singh@cs.pdx.edu) received his Ph.D. in computer science from the University of Massachusetts at Amherst in 1990 and his B.Tech., also in computer science, from the Indian Institute of Technology in 1984. He is currently a professor in the Department of Computer Science at Portland State University. His research interests include green technologies for Internet devices, energy-efficient computation, wireless networking protocols, information theory, and performance analysis.

[3] *http://www.energystar.gov/*