# The GreenNet Project

Roberto Bruschi, Raffaele Bolla, Franco Davoli
CNIT – Research Unit of the University of Genoa
Genoa, Italy
e-mail: {name.surname}@cnit.it

Antonio Cianfrani, Marco Listanti, Marco Polverini
DIET – University of Rome “La Sapienza”
Rome, Italy
e-mail: {surname}@diet.uniroma1.it

Gregorio Procissi, Rosario Garroppo, Stefano Giordano
DIE – University of Pisa
Pisa, Italy
e-mail: {name.surname}@iet.unipi.it

***Abstract*—The GreenNet project aims at re-thinking and re-designing different aspects of the current Internet technologies and network protocols to smartly support hardware power management. The pursued approach exploits two basic features already and largely present in today's networks and devices: the network resource virtualization and the modular architecture of nodes. These features give the opportunity of decoupling physical elements (e.g., a line-card) from their (virtual) functionalities and resources, so that the latter can be migrated towards other active physical elements. In this way, the emptied physical elements may be put in standby mode, while their logical services may continue to work elsewhere. In this respect, GreenNet aims at investigating a coordinated set of architectural solutions, protocol enhancements, control and optimization strategies, and related software developments in order to support such kind of primitives at both core and access networks.**

***Keywords-green networking; virtualization; energy efficiency***

## I. INTRODUCTION

Today, energy efficiency can be considered as one of the biggest technological and financial challenges for developing a better and more sustainable world. This challenge arises from the need of reducing the operating, manufacturing and energy related expenses of enterprises, industries as well as residential buildings, while keeping an eye on targets for the reduction of greenhouse gas emissions. Such "green" trend is speeding up innovation in a wide range of technological fields and services, since it offers the building blocks of innovative efficiency policies for a sustainable but economically viable energy supply, one that would also encourage new industry and employment and generate broad support in society.

In this respect, Information and Communication Technology (ICT) plays a crucial role, since it is widely regarded as an enabler of energy efficiency across the economy. This includes fostering the change in citizens' behavior, as well as improving efficiency in the use of natural resources while reducing pollution and dangerous waste. ICT's nature is so pervasive throughout all kinds of economic and social activities, that a widespread opinion encourages its increasing use for achieving energy savings from other industrial sectors.

Against this background, recent and official studies estimate 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 instance, the overall energy consumption of Telecom Italia in 2006 has reached more than 2 TWh (about 1% of the total Italian energy demand, and the second national consumer after the Italian railways), increasing by 7.95% with respect to 2005, and by 12.08% to 2004 [1]. In this respect, Tucker et al. [2] stated that today's networks rely very strongly on electronics, despite the great progresses of optics in transmission and switching, and outlined how energy consumption of the network equipment is a key factor of growing importance. In this sense, they suggested that the ultimate capacity of the Internet might eventually be constrained by energy density limitations and associated heat dissipation considerations, rather than by the bandwidth of the physical components [3]. In addition, starting from an expected deployment of the Telecom Italia network by 2015-2020, Bolla et al. [4] outlined how the power consumption of end-user network devices (e.g., home-gateways, VoIP phones, etc.) will represent a figure of more than twice the energy requirements of Telco equipment.

The origin of these trends can be certainly found in current Internet infrastructures, technologies and protocols, which are designed to be extremely over-dimensioned and available 24/7. Links and devices at access and core levels 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.

Advanced features for power management are already available in the largest part of hardware technologies, today used for building network devices. Silicon of network interfaces and of device internal chips already has the possibility of entering standby modes, or of scaling its working speed and, consequently, of lowering their energy requirements.

However, the activation of these power management schemes is generally hindered by the network protocols and architectures themselves, since they are specifically designed to be always available at the maximum speeds. Even Personal Computers (PC) are often left powered up 24/7, to maintain their network services and applications correctly working. Indeed, elements in standby (e.g., links, nodes or hosts) do literally "fall off" the network, since they are not able to exchange protocol signaling messages to maintain their "network presence". Moreover, sleeping or wakeup events generally trigger network nodes to exchange signaling traffic, and to re-converge towards new network logical topologies and/or configurations, causing transitory network instabilities and signaling traffic storms.

In access/home networks, if a PC enters standby mode, its applications and network connectivity services (e.g., ARP replies, TCP connections, etc.) literally freeze, and are not able to continue working, or to simply maintain their "network presence". Moreover, when the PC is waked up, it loses the "working states" of its applications and services, as they were before the standby periods. Then, upon wake-up events, the PC resets the states of applications and services, and, for each of them, it triggers new initialization/re-registration procedures. This is the reason why many PCs or networked devices that have advanced power management capabilities are left fully powered on 24/7 by users.

Starting from these considerations, the GreenNet project aims at re-thinking and re-designing different aspects of the current Internet technologies and network protocols to smartly support hardware power management. GreenNet, formally "Greening the Network," is a three-year project funded by the Italian Ministry of Research and Education under the program "FIRB – Futuro in Ricerca." The project started in March 2012, and joins researchers from three research units: CNIT at the University of Genoa (principal investigator: Roberto Bruschi), the University of Rome "La Sapienza" (local coordinator: Antonio Cianfrani) and the University of Pisa (local coordinator: Gregorio Procissi).

The pursued approach exploits two basic features already and largely present in today's networks and devices: the network resource virtualization and the modular architecture of nodes. These features give the opportunity of using the same base concepts already applied in other fields (e.g., data-centers): decoupling physical elements (e.g., a line-card) - which may be put in standby or have their capacities scaled down to perform only base operations - from their (virtual) functionalities and resources, so that the latter can be migrated towards other active physical elements of the same device, or of other neighboring devices. In this way, the emptied physical elements may be put in standby mode, while their logical services may continue to work elsewhere.

This paper is organized as follows. Section II introduces the project approach to green networking, while Section III describes the project's objectives in some detail. The Conclusions are drawn in Section IV.

## II. THE GREENNET APPROACH AND ORGANIZATION

The GreenNet project aims at re-thinking and re-designing different aspects of the current Internet technologies, spanning from the internal organization of network devices to the operating behavior of network protocols. The project is trying to propose a viable approach to effectively activate power management hardware primitives in next-generation devices, and to smartly support them in order to meet network operational and performance constraints.

In order to support hardware power management capabilities, the GreenNet project is mainly founded on the idea of maintaining the presence of network services and/or nodes, against the possibility that some hardware would fall off the network to enter standby mode. The project applies the same simple idea to a heterogeneous set of real-world scenarios, including core and access/home networks. To this purpose, the approach pursued by GreenNet exploits two basic features already and largely present in today's networks and devices: the network resource virtualization and the modular architecture of nodes.

The GreenNet project aims at investigating a coordinated set of architectural solutions, protocol enhancements, control and optimization strategies, and related software developments in order to support such kind of primitives at both core and access networks. From a more general point of view and in extreme synthesis, such approach allows:

- Core network nodes aggregating routing and switching functionalities into a subset of their physical resources (e.g., line-cards), and put their emptied modules in standby mode.
- End-host terminals (PCs, Customer Premises Equipment) migrating the "network presence" of their applications and/or services to network devices, and entering standby while maintaining network services up.

To achieve these objectives, the ideas carried out in the GreenNet project will not be limited to the development of analytical and simulative models, but they will be included in "proof-of-concept" prototypes of green devices and networks (realized on highly programmable platforms, like software (SW) routers, OpenFlow and OpenWRT devices, NetFPGA, etc.), in order to demonstrate their feasibility and main impact.

Regarding the organization of the research activities, they are structured along three research lines, as in the following sub-Sections.

### A. *Enabling Power Management Capabilities on Network Devices*

This research axis is focusing on the activation and on the impact of power management capabilities in hardware (HW) modules (network interfaces, silicon for packet processing, etc.) of the network device data-plane. To this end, this research axis mainly regards the data-plane of network equipment, since this usually includes the most energy starving HW elements.

As first step, GreenNet started by exploiting the main results achieved by other research projects in the same area.

Here, the main idea consists of adapting the advanced power management capabilities studied in the ECONET project [5] (namely power scaling and smart standby), in order to make them suitable for supporting the virtualization mechanisms and paradigms.

Then, the project is trying to exploit the modularity of network device data-plane for supporting advanced strategies for reducing energy requirements. In this respect, the main idea is to investigate the possibility of dynamically activating a set of capabilities/services in a device (or in one of its elements), and to put to sleep and/or to bypass all the elements that are not needed in the current configuration. For example, let us focus on a core router based on line-cards and switching matrix: each line-card usually includes different HW modules for processing the traffic at different layers (from the physical one to the IP, passing through L2/L2.5 network protocols like MPLS, Ethernet, etc.). All these modules are usually over-dimensioned for processing peak traffic volumes. Here the idea simply consists of moving some packet processing to other line-cards, or even to other network nodes. To achieve this goal, the first step is to put internal modules working at higher levels to sleep, and to bypass them by forwarding traffic elsewhere in the device or in the network.

### *B. Network-aware support for hardware Power Management*

Starting from the capabilities enabled in the previous research line, here our main objective is to decouple network services and applications from physical elements, and to study protocol/architectural frameworks for moving them between different hardware elements.

Specific research activities are devoted to the virtualization of applications and protocols at access and home network level. Here the project is studying viable approaches to allow PCs and networked devices to transfer their "network" presence and working status to third-party devices (like OpenWRT home gateways).

These third-party devices maintain the network connectivity through "routing signaling traffic" on behalf of sleeping hosts, and wake them up on the reception of "exception traffic" (i.e., traffic carrying requests that can be processed only by the end-host). The project is exploring (*i*) possible modifications to L4 protocols to transparently support these new features, (*ii*) possible security issues.

Regarding core networks, the project is pursuing a similar aim, but trying to virtualize device functionalities (e.g., the forwarding process at the IP, MPLS or optical levels). Here, the idea is to base such network virtualization on programmability primitives and frameworks coming from recent standard proposals like OpenFlow or IETF ForCES that already offer the possibility of decoupling routing/switching logical instances from the device data-plane hardware.

Then, specific methodologies are going to be proposed in order to migrate these virtual entities from the node/element going to sleep to the one remaining active, and/or bypassing the functionalities disabled (i.e., IP forwarding) in a part of the network/device. Such methodologies have to guarantee zero packet losses, reliability, security and common network performance requirements.

### *C. Controlling and planning Green Networks*

GreenNet is investigating innovative criteria to:

- enable next generation networks and related equipment to locally and jointly control and manage energy consumption according to traffic volumes and performance requirements, and to eventually decide to move virtual service/functionality entities among hardware platforms.
- design next-generation networks with low-energy consumption requirements, by explicitly taking into account the virtualization/migration capabilities introduced in the previous research axes.

These ambitious goals will be achieved by exploiting the novel green capabilities and technologies resulted from the previous research axes. Here, the output will be a new integrated framework for flexible and cognitive network management and operation. Such framework may enable dynamic, ad-hoc and optimized resource allocation, control, deployment, and administration, at both levels of the individual devices and of the network. In the Core-network scenario, example policies can be changing the redundancy schemes, reducing the backbone bandwidth, updating routes, etc. Moreover, specific solutions will be studied in order to statistically guarantee network service performance while reducing the energy consumption, or for a given energy budget assigned to the network or to sub-groups of devices.

## III. THE PROJECT GOALS

As previously introduced, the GreenNet approach tries to focus on the integration of three main aspects that contribute to the whole picture: (*i*) Enabling Green Primitives in Next-Generation Network Devices; (*ii*) Exploiting Virtualization Schemes in Network Protocols and Functionalities; (*iii*) Network-wide "green" optimizations. According to this structure, and in order to achieve the final objective of a significant reduction in energy consumption levels, through the concerted actions of the different parts, the project has three sub-goals, summarized in the following sub-Sections.

### *A. Enabling Green Primitives in Next-Generation Network Devices*

The objective here is to integrate novel highly flexible and modular architectures for next-generation network devices with advanced power management primitives (like standby and power scaling modes) in order to prepare a solid basis for applying advanced virtualization schemes.

In this respect, our aim is to exploit emerging device-internal protocols like the OpenFlow and the IETF ForCES standard proposals, which already give a good chance of abstracting and virtualizing device resources and functional modules. This approach is pursued together with the usage of open-source networking software. By means of all these tools, on which GreenNet partners have long and deep expertize, the idea is to set-up experimental modular platforms with IP routing capabilities composed by off-the-shelves (COTS) and network-specific HW (e.g., Netlogic evaluation boards and OpenFlow switches). In more detail, OpenFlow switches are used for smartly (de)multiplexing and redirecting traffic flows among various packet processing elements, which are realized by means of multiple multi-core SW routers based on both commercial COTS hardware and XLP Netlogic evaluation boards. Such sub-elements are devoted to perform complex and high-speed traffic processing and forwarding operations, and then they certainly represent the most energy starving parts of the entire platform. Here, the main idea is to exploit the modularity at different levels for accurately and separately tuning the tradeoff between power consumption and performance for each device's element and functionality. We exploit the modularity of the overall platform by enabling entire forwarding elements to dynamically enter standby status. Then, we will extend our approach by dynamically managing the internal hardware resources of each active element through the exploitation of the capabilities offered by the Advanced Configuration Power Interface (ACPI). In this respect, our specific objectives include:

- putting CPUs/cores or network adapters that are not needed in deep sleeping mode;
- tuning the operative clock frequencies of active CPUs/cores in order to process incoming traffic with the desired QoS level.

A central supervisor is devoted to perform router control operations, and to dynamically decide the energy-aware configurations of all elements. The exchange of configuration/control data is performed by means of a special version of the IETF ForCES protocol, suitably extended to carry power-management information and commands.

### *B. Exploiting Virtualization Schemes in Network Protocols and Services*

A further step is revisiting parts of the TCP/IP stack in order to re-design some aspects of current network protocols in a more "virtual" way, in the sense of make their operational behaviors more de-coupled from HW hosts and nodes. The impact of a such "virtual" stack is literally a panacea for enabling advanced policies for network power management, since it allows devices (and/or their sub-elements) to migrate and to delegate some of their network tasks to other hardware remaining active. In this respect, specific research efforts will be devoted to both access and core network scenarios.

As far as access networks are concerned, the main idea is to create a network-integrated support for end-host (e.g., PCs) sleeping states. Here, the "network connectivity proxy" (i.e., an always active system that responds to "routine" traffic on behalf of sleeping hosts) proposed by K. Christensen *et al.* [6] represents a solid starting point, and will enable research activities to focus on various aspects and possible solutions. In more detail, we are trying to explore:

- The possible role of OpenFlow home gateways and switches in identifying traffic flows destined to sleeping hosts and redirecting them towards the proxy maintaining their network presence;
- How to make the management of end-host "network presence" as independent as possible from the application layer;
- The impact of asynchronous transport protocols, which may enable network nodes to "cache" packets until the target end-host will re-wake up;
- Possible security issues due to the "delegation" of networking services to external equipment.

Further research activities focus on the network-aware support for virtual machine "live-migration", which is currently supported by most SW virtualization environment (including open source projects such as Xen and Virtualbox).

Regarding the core network scenario, the project is studying efficient mechanisms for handling and migrating virtual functionalities and their working status among different sub-elements of a network node. Here, the main objective is to support standby events or processing re-allocations in a modular router architecture (like the one described in the previous sub-section), without causing network instabilities and performance drawbacks. In detail, the problem is that elements in standby do literally "fall off" the network (together with their links and protocol instances), since they are not able to exchange protocol signaling messages to maintain their "network presence". Moreover, given the features of routing and traffic engineering protocols, the falling off of any elements generally triggers all network nodes to exchange signaling traffic, and to re-converge towards new network logical topologies and/or configurations, causing transitory network instabilities and signaling traffic storms.

The GreenNet objective in this respect is to extend and to manage existing network protocols in order to support standby and wake events in a transparent way with respect to the IP layer. In more detail, our idea is mainly based on the exploitation of today's L2 protocols for backbone networks (e.g., mainly MPLS and Ethernet), since:

- they are specifically used to manage the virtualization of the physical network infrastructure;
- they already include efficient mechanisms for rapidly moving/migrating virtual links across the network (e.g., the fault recovery procedures).

Then, we fully exploit L2 protocols to migrate virtual links from the line-card entering standby to other line-cards. This obviously requires a new L2 links re-mapping on the physical network topology, since each virtual link has to enter the device through other line-cards. At this point, IP interfaces and packet and protocol instances working at higher layers can be "migrated" towards the new line-cards.

The key advantage of this solution with respect to a simple switching-off consists of reduced recovery times, as well as in the possibility of managing device standby and wake-up events in a transparent way with respect to the IP layer, avoiding useless signaling storms and slow network re-convergences. Moreover, a special focus is devoted to network resilience. Novel technologies and mechanisms are going to be studied to make redundant devices and interfaces enter very low power states, while maintaining network presence and connectivity.

## C. Network-wide "green" optimizations

A logical consequence of the first goal is to analyze the network-wide effects of the "green"-optimized devices on the behavior of the network as a whole, and to develop strategies for their coordinated action. To this aim, we need to study new routing and traffic engineering concepts, which can take into account the new possibilities offered by the networking equipment, and even decide whether some pieces of it may be altogether temporarily put in standby mode, in favor of diverting traffic to other parts of the network, to concentrate power effort. Different routing and path computation algorithms for different network layers will be proposed:

- IP: extensions to IP routing protocols (OSPF) that minimize the overall energy consumed by the network for the routing of the IP traffic on the basis of the network topology and of the modular power consumption of network nodes
- RSVP/MPLS: path computation algorithms that minimize the overall power consumption of the network used to route the traffic on the basis of traffic matrix, network topology and energy characterization of the hardware platforms
- Optical/GMPLS: routing and wavelength assignment (RWA) heuristics that minimize the overall energy consumed by the optical devices to route light-paths.

In all these cases, the availability of virtualization techniques and node modularity, introduced in the two previous sections, are key aspects to obtain power savings. These new features provide network device power consumption depending on managed traffic and on the aggregation of network functions in virtualized entities; both attributes will be exploited by the new routing and path computation algorithms in order to minimize the overall power consumed by the network during its operation when traffic conditions change.

The optimization of these strategies requires the definition of new metrics, able to reflect energy-consumption goals with performance-oriented ones, and the adaptation of existing protocols and methodologies.

In addition, GreenNet focuses on the extension of current state-of-the-art methodologies for network design and planning, by explicitly considering modular device architecture and components' energy consumption: in this way it will be possible to gain control of the largest part of telecom operating expenses (OPEX). Thus, these techniques allow to consider "economical" constraints, related to the estimation of energy consumption, besides traditional performance parameters (delay, packet loss, bandwidth) when designing a network. After an analysis of actual methodologies, planning choices causing low energy efficiency in the actual network will be identified. Given these sub-goals, the final aim of the project is to give an initial characterization of a networking infrastructure where all the above-mentioned methodologies and techniques and the corresponding protocol enhancements are active. This environment will be demonstrated through the deployment of an experimental test bed, with the following aims:

- providing an experimental instantiation of the main energy-aware protocol and architectural functionalities;
- analyzing the performance of the system as a whole (and compare the results with those predicted by analytical and simulation tools developed within the course of the research activity);
- producing a first prototype of a platform for the investigation of the problems arising and the possible adaptations required for the transfer of the research result to an industrial deployment.

# IV. CONCLUSIONS

This paper introduced the main motivations, the objectives and the research approaches of the GreenNet project, funded by the Italian Ministry of Research and Education under the program "FIRB – Futuro in Ricerca." The project started in March 2012, and will last for three years. In detail, the project will try to study innovative solutions for enabling advanced power management capabilities in network and networked devices, trying to avoid any possible performance drawbacks in network functionalities. To this end, three research axes will jointly contribute to the whole picture: (*i*) Enabling Green Primitives in Next-Generation Network Devices; (*ii*) Exploiting Virtualization Schemes in Network Protocols and Functionalities; and (*iii*) Network-wide "green" optimizations.

# 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, July 2012.

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

[3] J. Baliga, R. Ayre, K. Hinton, R. S. Tucker, "Photonic Switching and the Energy Bottleneck," Photonics in Switching 2007, San Francisco, CA, USA, 2007.

[4] 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 Commun. Magazine, vol. 49, no. 8, pp. 80-86, Aug. 2011.

[5] The ECONET Integrated Project, funded by the European Comm. under the 7th framework programme; http://www.econet-project.eu.

[6] C Gunaratne, K Christensen, B Nordman, "Managing Energy Consumption Costs in Desktop PCs and LAN Switches with Proxying, Split TCP Connections, and Scaling of Link Speed," Internat. J. of Network Management, Wiley, vol. 15, no. 5, pp. 297-310, Sept./Oct. 2005.