Telecommun Syst
DOI 10.1007/s11235-011-9645-7

Towards an energy-aware Internet: modeling a cross-layer optimization approach

**Sergio Ricciardi · Davide Careglio ·
Germán Santos-Boada · Josep Solé-Pareta · Ugo Fiore ·
Francesco Palmieri**

Abstract The containment of power consumption and the use of alternative *green* sources of energy are the new main goals of telecommunication operators, to cope with the rising energy costs, the increasingly rigid environmental standards, and the growing power requirements of modern high-performance networking devices. To address these challenges, we envision the necessity of introducing *energy-efficiency* and *energy-awareness* in the design, configuration and management of networks, and specifically in the design and implementation of enhanced control-plane protocols to be used in next generation networks. Accordingly, we focus on research and industrial challenges that foster new developments to decrease the carbon footprint while leveraging the capacities of highly dynamic, ultra-high-speed, networking. We critically discuss current approaches, research trends and technological innovations for the coming green era and we outline future perspectives towards new energy-oriented network planning, protocols and algorithms. We also combine all the above elements into a comprehensive energy-oriented network model within the context of a general constrained routing and wavelength assignment problem framework, and analyze and quantify through ILP formulations the savings that can be attained on the next generation networks.

S. Ricciardi · D. Careglio · G. Santos-Boada · J. Solé-Pareta
Dept. d'Arquitectura de Computadors, Universitat Politècnica
de Catalunya, Jordi Girona 1-3, 08034, Barcelona, Spain

S. Ricciardi
e-mail: sergior@ac.upc.edu

D. Careglio
e-mail: careglio@ac.upc.edu

G. Santos-Boada
e-mail: german@ac.upc.edu

J. Solé-Pareta
e-mail: pareta@ac.upc.edu

U. Fiore
Complesso Universitario Monte S.Angelo, Università di Napoli
Federico II, via Cinthia, 80126, Napoli, Italy
e-mail: ufiore@unina.it

F. Palmieri (✉)
Dipartimento di Ingegneria dell'Informazione, Seconda
Università di Napoli, via Roma 29, 81031, Aversa, CE, Italy
e-mail: fpalmier@unina.it

Keywords Energy efficiency · Energy-awareness · Energy-oriented network models · Power consumption minimization · Carbon footprint minimization · Integer Linear Programming

1 Introduction

The growing energy requirements for powering and cooling the various devices enabling the up-to-date ICT infrastructures, together with the rising energy costs consequent to the exhaustion of traditional fossil sources, are drawing an increasing attention to the energetic aspects of ICT in the modern world. In addition to the economic motivation, there is also a strong environmental rationale for energetic concerns. Electricity can be obtained by "dirty" primary energy sources (e.g. burning oil, gas), releasing in the atmosphere large quantities of fine particles (aerosols) and green house gases (GHG) contributing to pollution and climate changes. Alternatively, it can be drawn from "clean" renewable sources (e.g. sun, wind, tide) that do not emit GHG at all during the use phase.[1] Both aerosols and GHG are widely recognized as the major cause for global warming.

---

[1] GHG may be emitted during the construction phase; anyway, renewable energy sources are beneficial over their entire Life-Cycle [34].

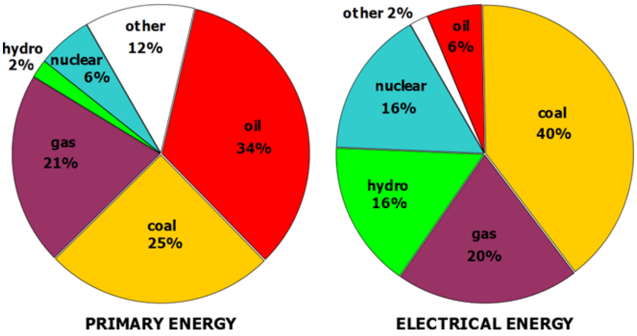

Fig. 1 Primary energy production and electrical energy generation taxonomy [1]

Currently, about 30% of worldwide primary energy is spent for producing electrical energy (with an average yield of 40%), and only a small share comes from renewable sources (Fig. 1) [1].

It has been estimated that ICT worldwide energy consumption amounts to 7% of the global electricity production [1] and the energy requirements of data centers and network equipment are foreseen to grow by 12% per year. Furthermore, with the ever increasing demand for bandwidth, connection quality and end-to-end interactivity, computer networks are requiring more and more sophisticated and power-hungry devices, such as signal regenerators, amplifiers, switches, and routers.

These components tend to increase the energy needs of global communication exponentially. Hence, it can be easily foreseen that in the next years the Internet will be no longer constrained by its transport capacity, but rather by its energy consumption and environmental effects [2]. In this scenario, networking equipment consumes about 1% of the total energy used for ICT, therefore it is important to keep such component into consideration when analyzing sustainable strategies for cutting the energy use. In fact, the amount of power spent worldwide for network infrastructures can be globally quantified in the order of tens of gigawatt [1], and hence limiting power consumption in networks is expected to significantly reduce the overall CO_2 and particle emission, so that the need for a greener, or—better—energy-oriented Internet is rapidly becoming a fundamental political, social and commercial issue.

At the state-of-the-art, miniaturization and ICT growing dynamics (i.e., Moore's and Gilder's laws [3, 4]) have not had the expected counterpart in power consumption reduction in the networking scenario. Miniaturization has reduced unit-power consumption but has allowed more logic ports to be put into the same space, thus increasing performances and, concomitantly, power utilization. Furthermore, the increased energy-efficiency may lead to decreased supply costs which may lead to augmented demand and consequent higher overall consumes that overtake the gained energy savings: a phenomenon called rebound effect (or Jevons paradox/Kazzoom-Brookes postulate [5]). Thus, despite architectural and semiconductor technology improvements, power consumption of network devices is still growing almost linearly with bit-rate and traffic volume [1]. As a consequence, the total power required per node is exploding. It is hence necessary to adopt a systemic approach that comprises state-of-the-art technologies (i.e., *energy-efficiency*) and new operation and management strategies (i.e., *energy-awareness*) exploiting renewable energy sources, acting in a cooperative fashion to achieve *energy-oriented* ICT. Therefore, energy has to be considered as a novel fundamental constraint for design, planning, and operations activities in the networking environment as well as in the whole ICT sector.

The envisioned future technological innovations are presented here together with a holistic view of the research challenges and opportunities that are foreseen to play an essential role in the coming green era towards sustainable (thus scalable) society growth and prosperity. Energy is considered as a novel additional constraint to design, plan and operate in the ICT systems. The semantic network in Fig. 2 illustrates the energy-oriented paradigm where both efficiency and awareness are used to achieve eco-sustainable ICT. It is an effort to visualize a framework in which decrease the energy consumptions and GHG emissions in the green Internet. The main factors that will drive this development are reported, and a visual clue on how these elements are connected with each other is given, helping to identify their relations and co-operations. The paradigm will evolve accordingly as new requirements and technological innovations come into the arena. The energy-oriented paradigm is depicted as an undirected graph, where a number of elements (nodes) and relations (edges) concur to build the complete framework. The connected elements work together and all the elements collectively contribute to achieve a holistic systemic approach. The leftmost part represents the highest-level elements, which *control* the rightmost lower level elements that are part of the global envisioned solution. At higher level, starting from the need to consider the energy consumption as a new constraint, policy should drive the changes both promoting virtuosos practices and discouraging environmental unfriendly approach: cap and trade, carbon offset, carbon taxes and incentives are all viable ways that governments are just starting to explore. In order for any solution to be successful, it is necessary to study its whole life cycle assessment, otherwise it may fall in the rebound effect and get increased energy consumption and concomitant GHG emissions. At lower levels, three main actors are highlighted and discussed: a global distributed energy system (Smart Grids), green data centers and networking.

In our envisioned framework, renewable energy (e.g. solar panels) should be available in every power-consuming site in an effort to provide each Internet service provider (ISP)/data center with its own green energy source. Indeed

Fig. 2 The semantic network for the energy-oriented paradigm in the green ICT era

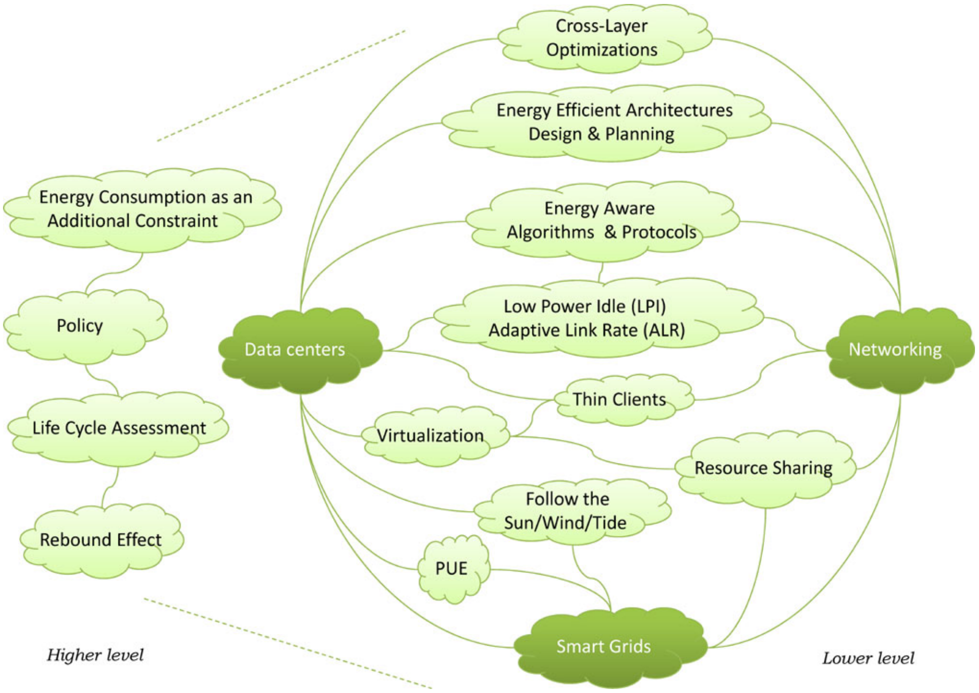

through an initial capital investment (CAPEX) the organizations that will employ this feature, will get reduced operational costs (OPEX), reduced GHG emissions and reduced energy dependency (which in the near future will mean financial survivability); besides, they may take advantage of the establishing green policies. Following these considerations, we envision a change in the way energy is produced and distributed. Energy production/consumption will shift from a current (insufficient and fossil-based) centralized model in which few big plants give energy to a whole geographic region to a distributed model in which a number of small renewable energy plants are spread over the territory (e.g. solar panels on the buildings roofs, mini-eolic wind mills in the streets, etc.): at every site, energy will be produced, exchanged, released when in excess and acquired when needed. This change will be encouraged by a number of factors: the current worldwide energy shortages, the rising costs of energy as fossil sources become scarcer, the need for new alternative and renewable sources of energy and the growing interests of governments and people into eco-concerns. In the same way as the cap and trade system, energy can be bartered and loaned without fees. Energy supply and demand may encounter reciprocally and exchanges may happen between neighborhoods at different hours of the day according to the different demands. The new model will not replace the existing one, but will come alongside, with the central system operating only when the renewable sources are temporarily exhausted or not available at all. Energy consuming facilities in the ICT sector will have access to several sources of energy, and will be able to switch on demand between the different energy sources dynamically acquiring or releasing excessive produced green energy as needed. In this direction, recent initiatives gathering major energy operators and providers started to explore intelligent and automated management of the future electricity distribution networks (see e.g. smart grid initiative in US [6]). To this extent, energy-awareness will become a fundamental operational feature, which can be also applied to other industrial sectors. We argue that this model can be extended to private houses, business premises, university campuses, ISPs, public buildings providing distributed green energy plants that may produce, release and acquire electrical energy (as well as cold and hot flows) from their neighborhood. Today in fact the great need is not for more energy, but for better energy utilization and wastes avoidance should be our primary objective. Data centers need to be cooled while office rooms need to be warmed (at least for several months, depending on the latitude). This supply/demand situation may be exploited by properly exchanging warm and cold flows between data centers and office rooms with both side advantages and without costs, with the enabling technology being an intelligent GMPLS-like control plane. Energy will change from a perceived cost to a revenue (re)source: saving energy reduces OPEX and produced green energy can be traded making revenues. A case in point is that in southern Germany the energy produced by photovoltaic panels exceeds the demand and a way of storing or exchanging energy is strongly advised [7]. Also, Google's 1.6 MW solar installation is the largest in corporate America at the time and recently Google has received the authorization to trade energy [8]. Another promising initiative was initiated by the administration of Iceland [9] that is promoting data centers placement near renewable energy plants. Industries and governments that will move first will obtain the greatest benefits from the new sustainable economy.

In data centers, energy-awareness can be considered from users' equipment to software and middleware level down to hardware resources. In particular, from the user perspective, the ICT trend is moving towards a network-centric paradigm, in which energy-hungry end-user equipment (e.g. PCs) is being substituted by thin clients with low power consumption and high-speed network connectivity (e.g. smart phones and netbooks), notably incrementing the use of network for connecting them to data centers and content delivery networks (CDN). If well managed, this evolution in the form users connect to the Internet can be properly exploited to significantly decrease the power consumption of both data centers and networking. Virtualization may play an important role in moving computations from the client to the server-side where data centers placed near renewable energy plants may execute the computations with lower carbon footprint with respect to traditionally power consuming (almost idle) PCs. Increasing computing density to sites where green energy is available will be the upcoming challenge for data centers. This trend will further increase the use of the network premises and consequently raise the need for energy-aware ICT network paradigms.

As a very complex combination of heterogeneous equipment, a network infrastructure has to be properly designed and managed in order to achieve advanced functionalities with a limited energy budget. Based on the state-of-the-art technological scenario, to decrease energy consumption in such networking infrastructures it is possible to operate on three different dimensions of the problem, according to a cross-layer approach:

- *Including energy-efficiency as the key requirement for the evaluation of technological advancements* is the first fundamental step towards energy-oriented networking. New devices that improve the performance of their predecessors need to be compared with competing technology also on their power consumption levels.
- Second, all the network *designing, building, and operating actions have to consider energy as an additional constraint* for the success of an energy-aware networking paradigm. Energy-awareness in network design is based on the concept of deploying algorithms and protocols that, taking into consideration the energy requirements, minimize the aggregate power demand while satisfying requirements for coverage, robustness and performance. This implies the adoption of an intelligent control plane together with the deploying of physical network topologies that aim at minimizing the number and the consumption of devices that must always be powered on. Also, dynamic power management strategies designed to decrease power consumption in the operational phase may introduce positive effects for the environment and significant cost savings. Accordingly, energy-aware logical network topologies can be dynamically built by making maximum usage of powered-on and highly connected devices exploiting as much as possible resource sharing by cleverly reusing switching nodes and fiber strands.
- The final dimension is the introduction of *energy-oriented control plane protocols* whose goal is to properly accommodate network traffic considering, energy-efficiency, energy-awareness and renewable energy sources. Daily and hourly fluctuations in user demands and green electricity availability may be considered across the involved network infrastructures/areas, in such a way that the overall power consumption and GHG emissions can be optimized. Accordingly, ILP formulations and heuristic methods can be introduced for calculating the routing information subject to power consumption constraints, also by taking into account the specific kind of energy source (dirty or renewable) used for powering the traversed network elements (NE).

Starting from the above operating dimensions, we present a model for the energy consumption aimed at describing and quantifying the savings that can be attained through control plane protocols that affect the route/lightpath choice privileging renewable energy sources and energy efficient links/switching devices. In this endeavor, we tied together into a general constrained Routing and Wavelength Assignment problem framework all the features needed by a comprehensive energy-aware network model. For both contexts where information about the GHG impact of the energy source is available and where it is not, we present Integer Linear Programming formulations to characterize optimization objectives and constraints in a formal and expressive way, and finally discuss some simulation results, analyzing the saving attainable.

2 Current approaches and research trends for energy-oriented networks

In this section we describe the three aforementioned dimensions of the cross-layer approach in network infrastructures, critically discuss the major issues and provide the basic grounds to our approach.

2.1 Evaluating technological advances for energy-efficiency

Technological advances allow having more efficient network devices that consume less and process more bits per second, according to the "do more for less" paradigm. Such solutions are usually referred to as eco-friendly. Currently, the power requirement for electronic network devices is scaling almost linearly with their total aggregate bandwidth.

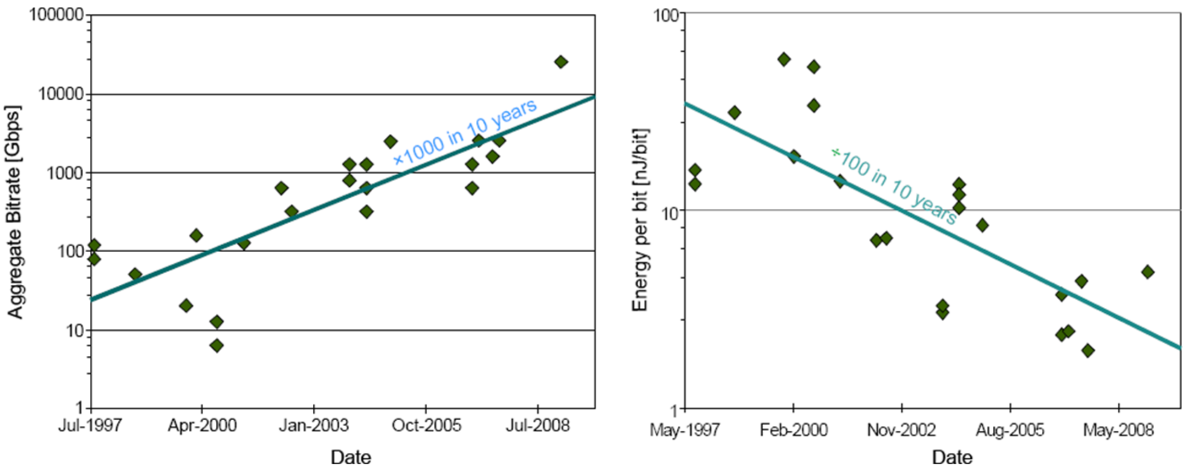

Fig. 3 Evolution of the bandwidth capacity and energy-per-bit consumption [1]

In particular, the energy-per-bit (measured in nJ/bit) decreased approximately by a factor of 100 in the last 10 years while the bandwidth capacity has increased by a factor of 1000 in the same time frame, thus—assuming that the demand almost equals the supply—the energy consumption increment is in the order of 10 times more than 10 years ago (Fig. 3). In other words, technological advancements foreseen by Moore's law have not been fully compensated by the same decrease in the energy-consumption.

The essential reason is that, despite the astonishing performance improvements in terms of transmission and forwarding capabilities observable today in networking devices, the energy dedicated to the primary functions of routing and switching is not exploited in the best possible way. The fundamental cause of energy consumption in networking equipment is the loss effect experienced during the transfer of electric charges, due to imperfect conductors and electrical isolators. Here, the exact consumption rate depends on the transition frequency and the number of gates involved, together with fabrication features (such as architecture, degree of parallelism, operating voltage, etc.). These values can be improved by industrial advancements but only to a certain extent, since there are physical bottlenecks inherent in the electrical switching technology involved. Therefore, traditional electrical networking can be considered as inadequate to support the emerging carbon footprint reduction requirements. Indeed, the power consumption of the actual electronic routing/switching matrix and line interface cards is, quite surprisingly, almost independent from the network load and can reach hundreds of kW for large multi-shelf configurations [10, 11]. Experimental energy consumption measurements [10] on several electronic routers show that in current architectures almost one half of the energy consumption is associated to the base system and up to another half to the active line cards. The traffic load only affects power consumption by a 3%; in other words, the energy demand of heavily loaded devices is only about 3% greater than that of idle ones. These results suggest that it is necessary to develop energy-efficient architectures exploiting the ability of putting into energy saving mode some subsystems (e.g. line cards, input/output ports, switching fabrics, etc.) in order to minimize energy consumption whenever possible. It has been also demonstrated [12] how consumption depends on the packets size and on the bitrates of the links. Traffic flows characterized by bigger packets need less energy than those made of smaller ones, due to the lower number of headers that have to be processed. Analogously, a circuit-based transport layer may reduce energy consumption with respect to a packet-switched one (Fig. 4). Although routers require more power when working at higher throughputs, if the per-bit power consumption is considered, larger routers operating within the core consume less energy per-bit than smaller ones located on the edge [13, 14], thus the power consumption will be greater on the network edge and smaller within the core, due to the higher traffic aggregation in the network core with respect to the edge.

Besides offering huge data rates (theoretically up to 50 terabits per second in a single fiber [15]), limited disturbance, and low cost, optical communication technology requires very low energy for the transmission of signal—light pulses—over the optical fibers. Furthermore, wavelength division multiplexing (WDM) technology (sending several independent optical signals in the same fiber cable using different wavelengths—80 wavelengths devices are commercially available) has dramatically increased the available bandwidth and greatly reduced energy consumption (Fig. 4). For comparison, an Optical Cross-Connect node (OXC) with micro-electro-mechanical system (MEMS) switching logic consumes about 1.2 W per single 10 Gb/s capable interface, whereas a traditional IP router requires about 237 W per port [1].

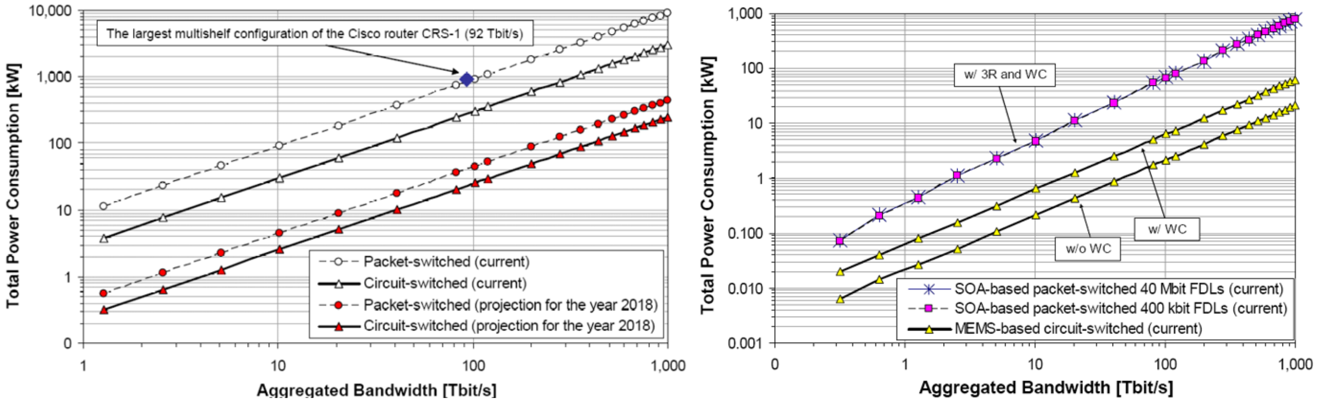

Fig. 4 Comparison of the total power consumption *vs* the aggregated bandwidth for electronic (*left*) and optical (*right*) router/switch technologies [1]

Table 1 Energy/power consumption as function of the aggregated bandwidth

Router technology	Energy Consumption Rate (ECR)	Energy Scaling Index (ESI)	Power consumption (P) as function of the aggregated bandwidth (B)
Electronic DXC	3 W/Gbps	3 nJ/bit	$P = 3 \cdot B$
Optical OXC with conversion	0.062 W/Gbps	0.062 nJ/bit	$P = 0.062 \cdot B$
Optical OXC w/o conversion	0.02 W/Gbps	0.02 nJ/bit	$P = 0.02 \cdot B$

ECR and ESI are different power consumption metrics that may be reduced to equivalent values, in fact it holds that: W/Gbps = (J/s)/(Gbit/s) = J/Gbit = nJ/bit

Many reference metrics can be used when comparing the energy consumption of networking equipment. The Energy Consumption Rate (ECR) [16], targeted towards high-end packet-based network and telecom equipment, defines a testing methodology and expresses the energy consumption per maximum throughput, typically *Watt/Gbps*. ECRW is a similar weighted metric that also takes into account off-peak and idle conditions. The Energy Scaling Index (ESI) is a metric to compare the efficiency of switching devices; the ESI corresponds to the switched aggregate bitrate offered for each *Watt* of energy budget. In Table 1 we report the mean energy/power consumptions for different router technologies under several energy metrics derived from real devices energy consumption values [1]. Electronic routing devices need 150 times more power than optical ones to route the same amount of traffic without wavelength conversion (WC), and nearly 50 times more than optical devices with WC. In electronic routers, an increase of 1 throughput unit corresponds to an increase of 3 units of power consumption; the same increase in an optical node performing WC corresponds only to an increase of 0.062 units of power consumption and to 0.02 units not supporting it.

However, as far as the optical signal still needs to be converted into the electronic domain (such as in current opaque optical network equipment) the power requirement will remain remarkably significant. Therefore, there is much room for improvement towards entirely optical networks, where most of the processing—ideally, all of it—is done at the optical layer, with the associated energy savings. Nevertheless, points of electronic processing are still necessary. At different network levels (access, metro, and core) electronic processing is required in order to aggregate low/medium bandwidth client signals into higher capacity flows—a process known as *traffic grooming*—and achieve high usage percentage of transmission links, reducing the required number of active node ports. At network interface points, the electronic level still seems to be desirable to maintain well distinct and separated Service Level Agreements (SLA) responsibilities between clients and operators. Furthermore, exhaustive network monitoring is currently possible only by converting and analyzing the signal into the electronic domain. Complete optical signal regeneration (“3R” regeneration: re-amplification, re-shaping, re-timing) is commercially available only by means of electronic energy-expensive processing devices (typically 60 W per wavelength/channel). Therefore, 3R regeneration should be avoided as much as possible in planning, designing and managing new paths throughout an optical infrastructure. On the contrary of 3R signal regeneration, optical signal amplification (1R regeneration) may be done entirely in the optical domain and should be used to extend the reach of optical fibers without any electronic conversion. Thus, instead of using energy-hungry 3R regeneration, it is preferable to use optical amplifiers (every 80–100 km) so that the other end of the fiber can be reached without 3R regeneration, with optical add and drop multiplexers (OADM) inserting and extracting client signals when needed. Commercially available optical amplifiers are mainly based on the erbium doped fiber amplifiers (EDFA) technology, while semiconductor optical amplifiers (SOA) are emerging as possible candidates to replace EDFA. Currently, EDFA are more performing (higher gain, lower insertion loss, noise and cross-talk effects) than SOA but have also higher energy consumption (respectively 25 W and 3 W). As SOA technology will evolve and reach EDFA performance, SOA will become the default choice for achieving low-energy optical signal amplification into future long haul optical fibers. In addition to the use of low-power SOA, the use of dispersion compensation fibers (DSF: ITU-T G.653, NZ-DSF ITU-T G.655/656) instead of "simple" single mode fiber (SMF: ITU-T G.652) will reduce the dispersion of the optical signal traversing the fiber and reduce the number of required optical amplifiers.

A research field related to optical networking promising further power savings is optical logic. It consists in incorporating photonic functionality in silicon very-large-scale integrated (VLSI) circuits and it is considered a natural choice for optical networks because it could provide the ability to build optoelectronic systems with integrated control electronics. It is also argued that the energy cost of converting data from the optical to the electronic domain and back is not inherent to the fundamental physics of such conversions, so that a properly designed integrated approach may help reduce this cost. However, optical logic remains challenging and one of the toughest issues is power absorption. Optical logic would indeed lose much of its attractiveness if it would consume more energy than regular silicon. Current silicon CMOS devices operate with energies in the range of femtojoules per operation and future transistors are expected to evolve towards capacitances of tens of attofarads (and therefore energies of tens of attojoules for operation at the expected voltage levels), and matching these values is a demanding target for natively optical devices [17].

New ideas are also emerging in the evolution of core networking and the converged transport and Ethernet for carrier networks. Driven by high-definition video and network computing, the bandwidth requirements are doubling every 18 months and Terabit Ethernet is forecasted to be needed as early as in 2015–2017 [18]. The IEEE 802.3 Working Group and ITU-T Study Group 15 have recently established draft standards for 100 Gigabit Ethernet [19]. These will call for the development of new optical transceivers, whose power requirements have been constrained at an 80 W maximum power consumption and maximum temperature of 70°C. While technology and component reuse is already established as a driver orienting decision about the next leap in speed of Ethernet, power consumption issues does not seem to have yet reached the same importance. The advent of 100 Gb Ethernet will bring many advantages related to the reduction of the required energy. In particular, the price and power consumption of one 100 Gb interface will be significantly lower than those of ten 10 Gb interfaces, and the efficient use of the DWDM links will limit the recourse to parallel links. However, the transition phase must be properly planned and managed. Not every operator will invest into 100 GbE at the same time. Thus, it is important that power awareness will be considered at all levels, according to cross-layer optimization principia, to obtain immediate benefits.

The choice of transport technology in access networks may also be a strong enabler for energy-efficiency. At the state-of-the-art, the vast majority of the energy consumption can be attributed to fixed line access connections. Today, access networks ("the last mile") are mainly implemented with copper based links and technologies such as ADSL and VDSL, whose energy consumption is very sensible to increased bitrates. The trend is to replace such technologies with fiber infrastructure, especially in the emerging countries in which new installations are being deployed from scratch. In access networks, the energy consumption scales basically with the number of subscribers, so that the massive diffusion of fiber to the home (FTTH) in place of old copper xDSL access links would have the dual benefit of dramatically increasing the access bandwidth and decreasing energy consumption. For comparison, a single ADSL link consumes about 2.8 W, while using a gigabit-capable passive optical network (GPON) as the access infrastructure will reduce the consumption to only 0.5 W, an improvement of about 80% for a potentially very high number of users. Such ongoing replacement is moving the problem to the backbone networks, where the energy consumption for IP routers, driven by the ever increasing bandwidth requirement, is becoming a bottleneck [20, 21].

Estimating the global footprint accurately is in many cases highly complex. The specific equipment density and hardware integration, heat dissipation and power supply specifications must also be kept in mind as fundamental parameters for energy efficiency, when considering collateral energy charges such as cooling and power conversion. The Power Usage Effectiveness (PUE), defined by the Green Grid [22], measures the efficiency of an ICT facility as the ratio of total amount of power used by the facility to the power delivered to the equipment, thus assessing the fraction of energy consumption due to, e.g., the HVAC (Heating Ventilation and Air Conditioning), the UPS (Uninterruptible Power Supply) subsystems and the lighting facilities. A PUE value of 2 is the current average, meaning that HVAC and UPS double the energy requirements [13]. In this scenario, overcooling can be considered as the main collateral energy drain and further gains can be obtained by using computational fluid dynamics and introducing contained cooling strategies. The use of cold aisles ducted cooling or in-rack cooling systems help to keep the volume of fluid being cooled at a reasonable minimum. Outlet air can be vented directly to the outside, or preferably reused for space or water heating elsewhere in the building as required with the consequent improvements in energy and carbon footprint.

The above issues become more and more significant when a network is to be built from scratch or network upgrade decisions need to be taken. In these cases, it is necessary to choose equipment and network topology considering not only performance and cost but also the energetic budget: the usual trade-offs in capital (CAPEX) and operational (OPEX) expenditures between design decisions will have to be evaluated under an energy-efficiency perspective. The effort should be twofold: on one side, developing commercially available all-optical devices that perform wavelength conversion and 3R regeneration without the need of energy expensive electronic devices; on the other side, planning the network in such a way that 3R regeneration is not needed at all.

2.2 Designing, building, and operating energy-aware networks

Current network design, configuration and management practices are based on deploying and maintaining infrastructures that are extremely reliable, provide performance that enables competitive SLAs, and offer a set of features and services that are attractive to a broad range of customers. To accomplish these goals, network architects typically conceive network infrastructures that are densely meshed, with many redundant interconnections between nodes, so that many alternative paths can provide multiple reachability options between geographically distant sites. Also, fair load balancing has always been a predilection of network designers and maintainers, because it increases the possibility of putting new traffic into the network. Since the traffic demand may not be known in advance, network operators need to ensure that they have sufficient free capacity for any demand that may reasonably emerge in the operating lifetime of their infrastructure.

When designing the layout of large scale infrastructures, it is desirable to find a good balance between the competing needs to avoid as many electrically-powered hops as possible (to reduce the power consumption at intermediate switching nodes or regenerators) and to not transmit data over excessively long stretches, because it's more energy-expensive to move data farther. Furthermore, traffic dynamics often present notably changes over time, resulting in different network usage between peak hours and the rest of the day. In these cases, the network has to be dimensioned to handle the maximum load, to satisfy the users' demand in peak hours, but the deployed connectivity resources risk to remain under-utilized by a wide margin for most of the time, giving rise to significant energy waste and unnecessary operating costs. It should be considered that it is possible to run only the part of the infrastructure that is really required at any time. This is an opportunity that cannot be missed and hence it is necessary to develop energy-efficient architectures exploiting the ability of selectively shutting down some links or putting into energy saving mode some devices or subsystems (e.g. switching fabrics, line cards, input/output ports, etc.) in order to minimize energy consumption whenever and wherever possible. Accordingly, adaptive power management strategies designed to decrease power consumption in the operational phase may introduce positive effects for the environment and significant cost savings, as a consequence of the reduced energy usage. Maximizing the reuse of energy-conservative transmission links and powered-on, highly connected devices—in contrast to spreading traffic on all the available switching nodes, fibers, and paths—power consumption can be drastically reduced by temporarily switching off unused devices and line cards. Because such "sleep mode" strategy can be implemented at different levels of granularity, the chosen scheme needs to be very flexible and ensure the potential to save energy as soon as few end-customers are disconnected. Nodes may be put into sleep completely (per-node sleep mode) or partially (per-interface sleep mode). However, we deem that a drastic energy containment strategy such as per-node sleep mode is too simplistic and its effectiveness on real world network environments is questionable due to the undeniable impacts on the carrier-level network economy both in terms of capital (making connectivity investments partially useless) and operational expenditures (reducing the meshing degree and hence resiliency and traffic engineering capabilities). Furthermore, state-of-the-art consumer electronics used in broadband infrastructures are typically designed to enable maximum performance in an "always on" mode of operation. By leveraging on hardware equipment similar to those used in laptops, supporting fast "sleep" or "low-power" modes, next generation networking devices will have an outstanding opportunity to efficiently reduce their power consumption when not in use. These may comprise energy proportional computing techniques, meaning slowing down CPU (Central Processing Unit) clock for inactivity periods or "simply" reducing execution speed for energy saving purposes. Fast full-clock return procedures will be needed in order to achieve the desired level of system responsiveness. Introducing these technologies in networking hardware architectures will imply a shift from the "always on" to the "always available" paradigm, where each device can spontaneously enter a sleep or energy saving mode when it is not used for a certain time and quickly wakes up or restores its maximum performance on sensing incoming traffic on its ports. On the other hand, putting into sleep mode at interface level may have some sense, in particular high speed ones, since typical commercial off-the-shelf (COTS) devices drastically improve their consumption when transmitting at their maximum rates. However, the support of sleep mode at the single interface or linecard level also requires modifications to current router architecture and routing protocols. In fact, an interface put into sleep mode may not respond to periodic *hello* messages of the routing protocol and be classified as "down"; consequent link state advertisement messages will spread along the network informing that the interface is down, causing stability problem to the convergence of the routing (or spanning tree) algorithm. For these problems, more than a static sleep mode, a per-interface "wakeup on activity" or "downclocking" mechanism [23] may be more viable and effective solutions. In the first case, the transmission on single interfaces is stopped when there is no data to send and quickly resumed when new packets arrive. To do this, the circuitry that senses packets on a line is left powered on, even in sleep mode. This mechanism can be implemented by introducing the concept of Low Power Idle (LPI) [24], which is used instead of the continuous IDLE signal when there is no data to transmit. LPI defines large periods over which no signal is transmitted and small periods during which a signal is transmitted to refresh the receiver state to align it with current conditions. Alternatively, the ability to dynamically adapt the link rate according to the real traffic needs can be another effective technique to reduce power consumption (Adaptive Link Rate, ALR). Operating a device at a lower frequency can enable reductions in energy consumption for two reasons. First, simply operating more slowly offers some fairly substantial savings. Second, operating at a lower frequency also allows the use of dynamic voltage scaling (DVS) that reduces the operating voltage. This allows power to scale cubically, and hence energy consumption quadratically, with operating frequency [25]. The adaptive link rate speed control mechanisms [26] aims at dynamically adapting the link speeds and interface behavior to the current network load by using some specific thresholds. In [14] it is shown that the energy consumption does not depend on the data being transmitted but only depends on the interface link rate, and hence is *throughput-independent*. In particular, faster interfaces require lower *energy per bit* than slower interfaces, although, with ALR, slower interfaces require less *energy per throughput* than faster interfaces, due to the higher fixed power consumption of faster interfaces circuitry. In such a context, circuit over-provisioning may lead to decreased operational costs (OPEX) at the expense of increased capital expenditures (CAPEX), i.e. a network interface may be provisioned with different circuits, say a low and a high speed one, and may switch between one or the other according to the required data throughput. It is also observed that for current technologies the energy/bit is the same both at 1 Gbps and 10 Gbps, meaning that the increase in the link rate has not been compensated at the same pace by a decrease in the energy consumption. After long discussion about which technology between ALR or LPI should be introduced into the Energy Efficient Ethernet (EEE), the IEEE 802.3 EEE Study Group chose in favor of LPI to reduce the energy consumption of a link when no packets are being sent. In [23] the LPI with packet coaliscing was used to improve the efficiency of EEE while keeping the introduced packet delays under tolerable bounds.

In addition, resiliency to failures can be very energetically inefficient when implemented through the provision of $1+1$ protection, since usually equipment stays always turned on for fast failure recovery, but is used very rarely—only when a failure occurs. Several alternatives to this standard scheme may reduce the energy consumption induced by protection. For example, resiliency can be provided without having redundant equipment stay in the fully operating state all the time, but rather keeping it down-clocked with fast wake-up capability, taking care that it is compatible with industry-standard path/span protection switch times (50 ms).

2.3 Energy-oriented control plane protocols

By considering the problem from the perspective of the topmost layers, it can be envisioned that the future green network will be based on a highly adaptive and reconfigurable transparent optical core. Several optimizations can be performed according to a "cross-layer" approach, whereby issues arising at the physical layer (e.g., energy consumption) can be handled at higher layers (typically within the control plane), through appropriate routing and signaling practices. Introducing energy-awareness in the network control plane is based on the concept of placing network traffic over a specific set of paths (and hence sequences of nodes and communication links) so that the aggregate power demand is minimized while end-to-end connection requirements are still satisfied. Every time a new path is established between any pair of nodes, traffic between these nodes will be routed on it as if in presence of a direct "virtual" connection between them, by creating the abstraction of a logical network topology on top of the physical one. Energy-oriented logical network topologies can be dynamically built by optimizing the choice of energy sources in such a way that renewable sources are preferred when available, possibly with a trade-off between path length and carbon footprint. In fact, network elements may have dual power supply: the always available power coming from dirty energy sources and the not always available power coming from green energy sources. Consider, for instance, the availability of energy produced by solar panels; it is strongly correlated with

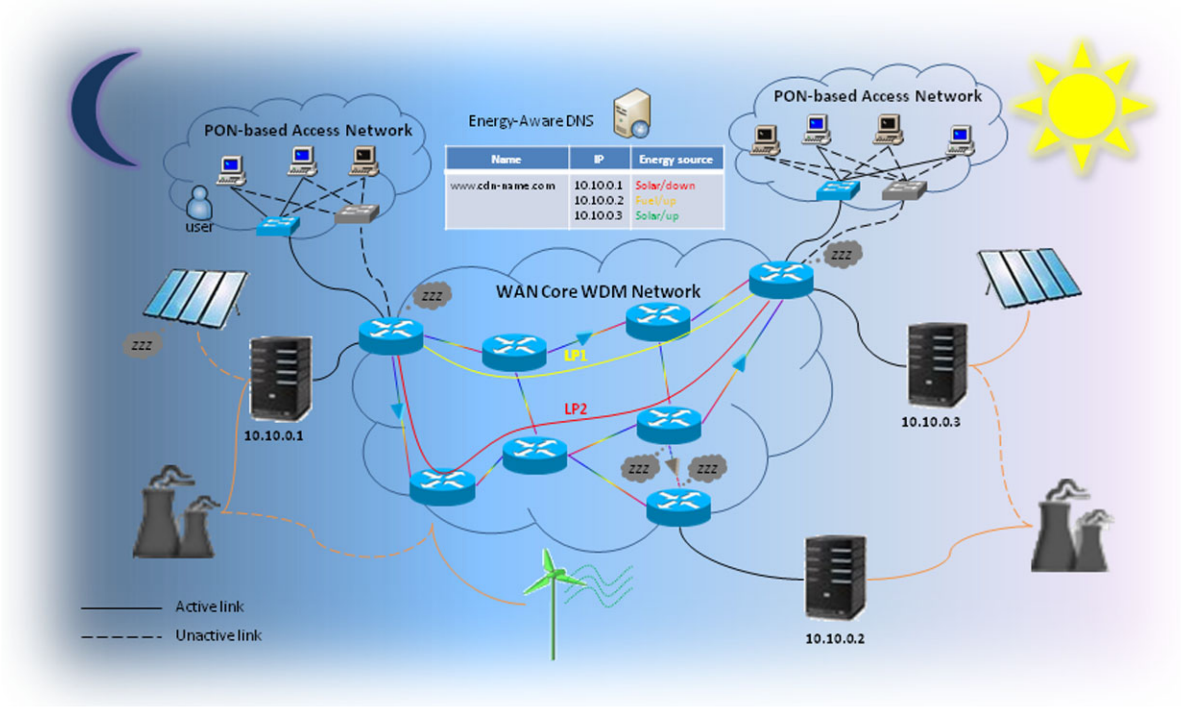

Fig. 5 The energy-oriented network infrastructure

the time of the day, since it is known that no energy will be produced during the night and that some energy is expected to be produced during the day. Such knowledge should be included in the control plane algorithms for energy-aware routing, signaling and resource allocation, implementing an automatic *follow the sun* paradigm. As another example, we can imagine some equipment powered by wind energy where power supply is a pseudo-random process depending on the availability of wind. Due to the inertia of the power generating mechanisms and batteries, a drop in the wind power does not result immediately in a power generation drop. Hence, if wind stops, it is possible to reconfigure the network dynamically, to consider the new distribution of available clean energy and re-optimize its carbon footprint. Differently from the case of the daylight, whose duration is known in advance, a decrease in wind strength is much more unpredictable and the warning time is shorter. This should be only handled with adaptive and efficient rerouting mechanisms implemented within the network core. For this reason, it is necessary to develop novel routing schemes and resource allocation mechanisms that take advantage of the early notification of the forecast power variation of clean sources with time-varying power output [10]. Furthermore, another interesting perspective in energy-aware networking comes from linking traffic routing to the different available electricity prices, dynamically and continuously moving data to areas/devices when and where electricity costs are lower.

Energy-awareness may be implemented at the application layer, e.g. in an energy-aware DNS (Domain Name System), bringing advantages both to data centers and networking. As an example, a content distribution network (CDN) is made up of several data centers located in several distant sites (e.g. Europe and US). Usually, each data center contains replicated data for security and load balancing purposes. So, large bursts of data have to be transferred between the sites. Among the different possible paths, the most energy-efficient ones may be chosen for transferring the data. Similarly, users' requests to access the CDN contents may be redirected to the current lowest carbon footprint data center by the energy-aware DNS constrained on the current energy supplies. These concepts are illustrated in Fig. 5, in which the following scenario is depicted. The WAN core network is a dynamically reconfigurable transparent WDM network and the access network is based on an energy-efficient passive optical network. The data centers sites 10.10.0.1, 10.10.0.2, 10.10.0.3 are part of the same CDN and data is mirrored among them with high-speed data transfers through the optical core using lightpaths (not represented here) chosen by an energy-aware routing protocol. Some data centers and routers are equipped with dual power supply (in sites where renewable energy source are available): the green energy source and the legacy always-available fossil-based energy source. The control plane is aware of the type of energy source that is *currently* powering routers and servers. When the green energy source is temporarily not available or the accumulated energy in the batteries is exhausted (for example because night has fallen or the wind has stopped), the UPS at the site switches to the fossil-based power supply without any energy interruption. Within data centers, a subset of servers is automatically put into sleep mode when the current load allows it. In the core and access networks, the unused network interfaces and the corresponding links and amplifiers/regenerators are dynamically put into sleep mode using an energy-aware control plane. Following a planning stage, end users premises are connected to two access points (or two line cards of a single access point), such that one access point (or one line card) can go to sleep mode as soon as a suitable number of clients are turned off and require no network activity. When a user (top-left corner) needs to download a file from the CDN, a query is made to a green DNS server that knows how the CDN servers that hold the desired file are currently powered up: server 10.10.0.1, although provided with the dual power supply, is currently powered up by the fossil-based energy plant because it is night; server 10.10.0.2 is using exclusively electricity generated by a coal power plant, while server 10.10.0.3 is currently powered up by clean energy and hence its IP address is returned by the DNS server. In this way, a paradigm shift towards energy-oriented networks and data centers is capable of sustaining the growing traffic rates while limiting and even decreasing the power consumption.

In order to support all the above adaptive behaviors, energy-related information associated to devices, interfaces and links need to be introduced as new constraints (in addition, for instance, to delay, bandwidth, physical impairments, etc.) in the formulations of dynamic routing algorithms and heuristics. Down-clocking or enhanced sleep mode features should be handled as new features in the network element status that need to be accounted for at both the routing and traffic engineering layer, and such information must be conveyed to the various network devices within the same energy-management domain. This clearly requires modifications to the current routing protocols and control plane architecture. For example, the existing routing (OSPF-TE, IS-IS) and signaling (RSVP-TE, CR-LDP) protocols within the GMPLS traffic-engineering framework may be extended to include energy-related information such as the power consumption associated to a specific link and the type of energy source currently used by the entire device. This can be easily accomplished by introducing new specific type-length-value (TLV) fields in IS-IS or opaque Link State Advertisements (LSA) in OSPF. Analogously, the same information has to be handled by signaling protocols such as RSVP-TE and CR-LDP to allow the request and the establishment of power-constrained paths across the network (i.e., path traversing only nodes powered by renewable energy sources or crossing only low-power transmission links).

However, in many cases, the carbon footprint improvements may be achieved at the expense of the overall network performance (e.g. survivability, level of service, stability, etc.), which can in turn be compensated through over-designing (increase of CAPEX) or over-provisioning (increase of OPEX). This implies that the new routing algorithm empowering the energy-aware control plane should be driven by smart heuristics that always take into account the trade-off between network performance and energy savings. In fact, by putting equipment or components that consume energy into a low energy consumption mode (e.g., nodes, line cards, links), or creating traffic diversions driven by reasons different from the network load, we implicitly reduce the network available capacity and hence paths tend to be longer and/or more congested, decreasing the overall transmission quality.

3 Modeling a cross-layer energy optimization framework for wavelength routed optical networks

Starting from the above considerations, we propose the introduction of energy-awareness into control plane protocols whose goal is to properly condition all the route/lightpath selection mechanisms on relatively coarse time scales by privileging the use of renewable energy sources and energy efficient links/switching devices, simultaneously taking advantage from the different users demands across modern wavelength routed network infrastructures, in such a way that the overall power consumption can be optimized. In doing this, we tried to combine, on each involved layer, all the notable features that a comprehensive energy-aware network model should have and put them together into a general constrained routing and wavelength assignment problem framework. Such problem has been modeled through integer linear programming to better characterize its formulation in terms of optimization objectives and constraints. Clearly, the ILP formulation, while effective in its formalism and descriptive power, is inherently static and hence implies a-priori knowledge of the entire traffic matrix to be solved at optimum; furthermore, the RWA ILP can be reduced to a single-commodity NP-complete problem. Anyway, for each real implementation perspective we need to assume that each node is capable to interact and cooperate with its neighbors by using a GMPLS or ASON-like control plane intelligence, enabling the exchange of the aforementioned energy-related information.

3.1 Basic modeling choices and assumptions

Defining a sustainable and effective model for energy consumption is the essential prerequisite for introducing power awareness within the wavelength routing context. A broad

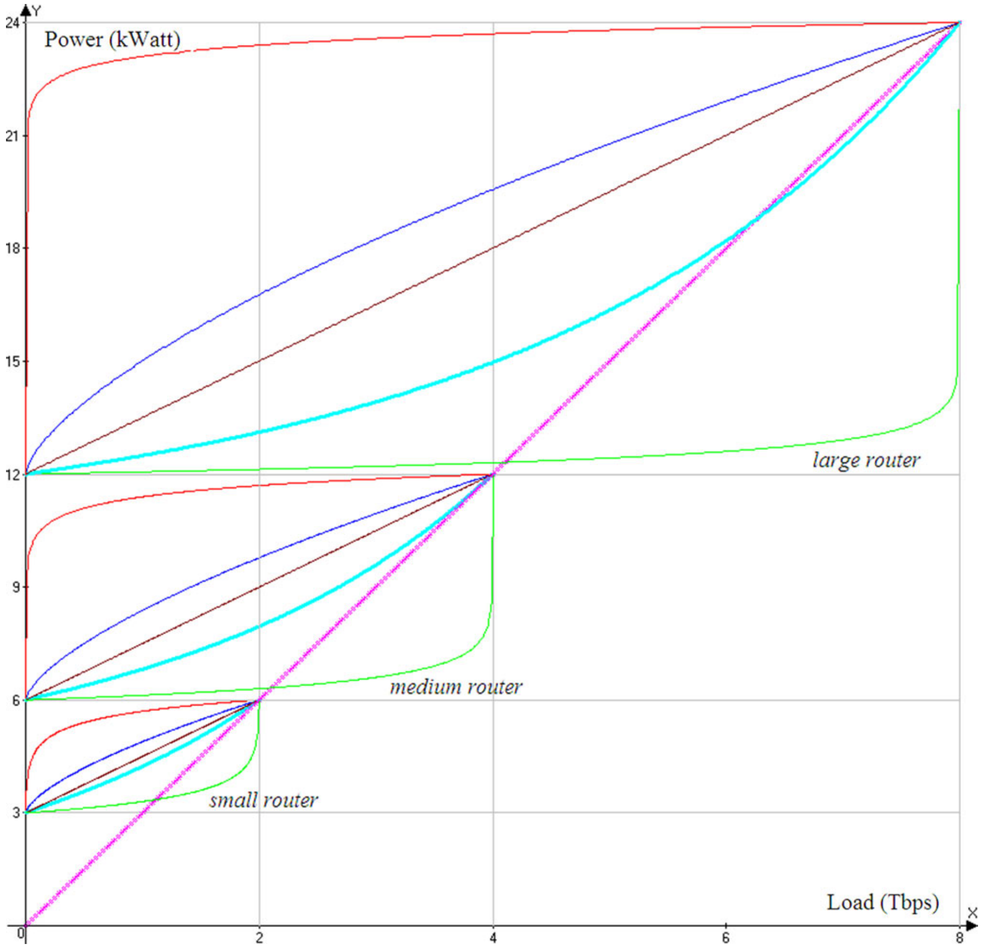

Fig. 6 Power consumption functions for various size electronic routers

variety of devices contribute to power adsorption in a WDM network: OADM, regenerators, amplifiers, "opto-electronic" and totally optical routers and switches. Each of these devices draws power in a specific way, which may also depend on the relationship between different devices or the components of more complex structures such as switches. In addition, NEs may be powered either by green or dirty energy sources statically assigned to each at the network topology definition time, therefore a differentiation between energy sources is required at the control plane level. In order to formally characterize the energy consumption of network elements we propose a comprehensive analytic model based on real energy consumption values and in line with the theoretical grow rate predictions encompassing new energy-aware architectures that adapt their behavior with the traffic load in order to minimize the energy consumption. Such an energy model characterizes the different components and sub-systems of the network elements involved. It provides the energy consumptions of network nodes and links of whatever typology and size and under any traffic load. The efforts in the developing of such an energy model have been focused on realistic energy consumption values. For this scope, the energy model has been fed with real values and the energy consumption behavior of NEs has been crafted in order to match with the state-of-the-art architectures and technologies. At this extent, future energy-efficient architectures with enhanced sleep mode features have been considered and implemented in the energy model. The energy model is based on a linear combinations of energy consumption functions derived from both experimental results [1, 10, 13, 26, 27] and theoretical models [6, 28]. Besides, following the results reported in [13, 29, 30], the power consumption has been divided into a fixed and a variable part; fixed part is always present and is required just for the device to be on; variable part depends on the current traffic load on the device and may vary according to different energy consumption functions. We chose [31] a linear combination of two different functions (logarithmic and line functions) and weighted them depending on both the type of traffic and the size of the NE, in order to obtain a complete gamma of values and thus adapting its behavior to the most different scenarios. In particular, in our energy model we managed to obtain that larger routers consume less energy per bit than the smaller routers (Fig. 6), as reported in [13, 14] and that electronic traffic consumes more energy per bit that optical traffic (Fig. 7), as reported in [1, 29]. Wavelength conversion

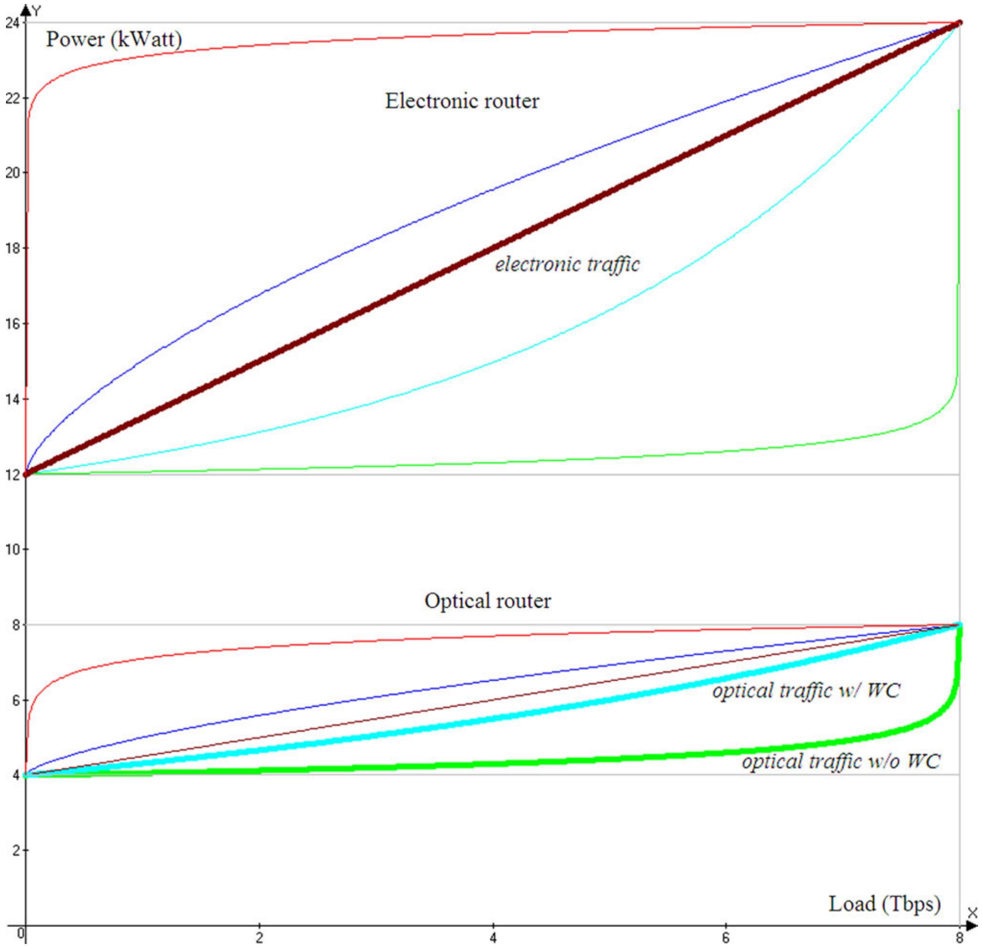

Fig. 7 Power consumption functions for electronic and optical routers

and 3R regenerations have a not negligible power consumption, which is accounted for in the model. Finally, links have an energy consumption that depends on the length of the fiber strands and thus on the number of optical amplification and regeneration needed by the signal to reach the endpoint with an acceptable optical signal-to-noise ratio (OSNR).

The power consumption functions of three routers of different sizes are reported in Fig. 6. Each router may support different types of traffic, each defined by a different curve (Fig. 7). In the example, the thicker lines represent the power required by a given type of traffic (e.g. electronic traffic).

We can observe that, according to our model, the larger the router, the larger the total energy consumption, as the fixed part notably contributes to (half of) the energy consumption. But if we focus only on the variable power consumptions, we observe that, for example, a traffic load of 2 Tbps, requires as much as 3 kW in the smaller router, about 1.5 kW in the medium one and just 1 kW in the larger router. In this way, we managed to obtain that greater routers consume less energy per bit than smaller ones, as reported in [13, 14]. Note also that the overall energy consumption scales linearly with the size of the router and that half of the energy consumption is due to the fixed part and the other half to the variable part, according to literature source [10].

In detail, at the basis of our model we consider wavelength-routed networks and, for the sake of generality, lightpaths that may have different bitrates (i.e., different bandwidth requirements, according to the particular SLA on the QoS of each client). The power required for transporting one lightpath will vary with its bitrate, so we consider as traffic unit the *bps* (bits per second). Network nodes may be electronic routers (digital cross connects, DXC) or optical switches (optical cross connects, OXC) connected by fiber links with up to λ wavelengths on each. The network is represented (Fig. 8) as a multigraph $G = (V, E)$ with $|V| = N$ nodes and $|E| = M$ edges (one for each wavelength λ in the optical layer).

We assume that the traffic is unsplittable, i.e. a traffic demand is routed over a single lightpath. In addition, we explicitly considered the influence of traffic on power consumption by using realistic data for traffic demands, network topologies, link costs, and energy requirements of single network elements. Specifically, the amount of power consumed by the NEs depends on the type of device and on the type and load of traffic that it is currently supporting (e.g. an

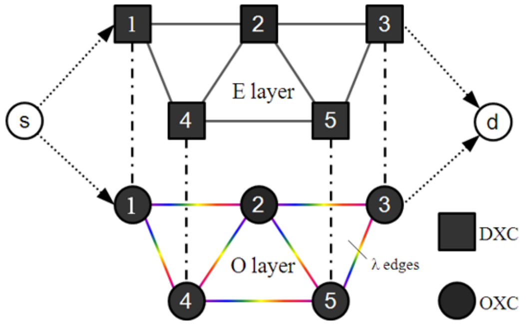

Fig. 8 Schematic representation of the network model with the electronic (E) and optical (O) layers

Table 2 Traffic supported by the devices

Type of device (NE)	Type of traffic
Electronic router[a]	Electronic
Optical switch[b] with WC capability	Optical (with or without WC)
Optical switch[b] without WC capability	Optical (without WC)
Fiber optic	Optical (without WC)
Optical amplifier	Optical (without WC)
3R regenerator	Electronic

[a] DXC

[b] OXC

OXC may support transparent optical traffic with or without WC). Even though the energy consumption of current node architectures does not scale with traffic (the energy demand of heavily loaded devices is only 3% greater than that of idle ones [10]), energy-aware architectures that adapt their performances to the traffic load lowering the power requirement under low traffic loads are strongly advised and are being designed [29]. Consequently, we assume that the power consumption of the NEs, i.e. both network nodes and links, consists of two factors. When turned on, a NE consumes a constant amount of power depending on the router size and technology (measured in $J/s = W$) and independent on the traffic load (fixed power Θ). The second factor (proportional power ε) consists of an amount of power proportional to type and quantity of the traffic load (measured in nJ/bit or, equivalently, in $W/Gbps$). The overall power drained by the WDM network is thus given by the sum of the fixed and proportional powers of the NEs subject to the current traffic load and varies with the routing of the connection requests. This implies that, as the NEs are always turned on, the routing optimization process works "only" on the proportional power. The power consumption functions of an electronic and an optical router are reported in Fig. 7 (optical router values not in scale). Three types of traffic are represented: (1) electronic traffic in the electronic router and (2) optical traffic with and (3) without WC in the optical switches. The Table 2 reports the types of network element and the corresponding supported traffic types. Note that each type of traffic accounts for different power consumption when traversing NEs, as explained in the following. We observed that the electronic traffic grows quickly with respect to the optical traffic and that, among the optical traffic, the WC actually consume a not negligible quantity of energy. As the power consumption functions are obtained by linear combinations of the logarithmic and the line functions, the complete gamma of slopes can be represented by the actual curves.

We also assume that all the nodes have the possibility to convert wavelengths, either in the electronic or in the optical domain, depending on their technology. In the electronic domain, the full range of operations is supported: wavelengths routing/switching, wavelengths add/drop, WC, 3R regeneration; in the optical domain the operations supported are the transparent wavelength switching/pass-through and the WC.

3.2 Energy-aware routing and wavelength assignment

In this section, three ILP formulations of the problem of energy-aware RWA in WDM networks with dual power sources are given, with different objective functions. First, the problem of minimizing the overall GHG emissions (*MinGas-RWA*) is presented. Next, the problem of minimizing the overall network power consumption (*MinPower-RWA*) is discussed. Finally, to obtain a reference for comparison, a minimum cost RWA (*MinCost-RWA*)—i.e., energy-unaware RWA—is derived and used. These formulations extend our previous work [32] to comprise the connection requirements on the guaranteed bandwidth (lightpath bitrate) thus supporting lightpath with different bitrates, and to prove its effectiveness against the well-known NSFNET network topology. In Table 3 we report the problem statements of the three ILP formulations and in Table 4 the associated notation; note that the three ILPs have the same inputs and constraints, but different objective function.

3.2.1 Energy-aware RWA at minimum GHG emissions (MinGas-RWA)

The energy-aware RWA in WDM networks with dual power sources (*MinGas-RWA*) is formalized as an ILP problem. The objective is to route the requested lightpaths so that the overall network GHG emissions are minimized. Only NEs powered by dirty energy sources emit GHGs, whilst NEs powered by green energy sources do not emit any GHG at all. The ILP problem can be mathematically formulated as follows.

The objective function for *MinGas-RWA* is:

$$Minimize\ DC_{G(V,E)} + \log TC_{G(V,E)} \tag{1}$$

Table 3 Problem statements of the three ILP formulations

Schematic view		
Given		
(1) the physical network topology comprising routers and links, in which links have a known capacity and cost,[a]		
(2) the knowledge of the average amount of traffic exchanged by any source/destination node pair,		
(3) the maximum link utilization that can be supported (wavelength link capacity),		
(4) the energy model (power consumption of each link and node),		
(5) a set of k candidate paths (routes) between any source/destination node pair,		
Find		
the routes that must be used		
In order to		
[MinGas-RWA]	[MinPower-RWA]	[MinCost-RWA]
minimize the total GHG emissions, and, as secondary objective, the minimization of the total power consumption,	minimize the total power consumption, and, as secondary objective, the minimization of the total installation cost,	minimize the total installation cost,[b]
Subject to		
traffic demand volume constraint, maximum link utilization constraint.		

[a] The cost, the wavelength conversion and the 3R regeneration are considered in the selection phase of the k candidate paths (routes) between the source/destination node pairs, though they will have different cost impacts on the problem

[b] it is assumed that the installation cost is proportional to the number of wavelengths required and to the length of the chosen lightpaths

Subject to the following constraints:

$$
\begin{aligned}
&DC_{G(V,E)} \\
&= \sum_{n\in V} g_n \cdot \left(\begin{array}{l} \Theta_n + \varepsilon_n^{t_1} \cdot \sum_{sd,k,b:n\in\pi^{sd,k,b},n\neq s,d} w^{sd,k,b} \\ \quad \cdot b \cdot (1 - x_n^{sd,k,b}) \\ \quad + \varepsilon_n^{t_2} \cdot \sum_{sd,k,b:n\in\pi^{sd,k,b},n\neq s,d} w^{sd,k,b} \\ \quad \cdot b \cdot x_n^{sd,k,b} \\ \quad + \varepsilon_n^{t_3} \cdot \sum_{sd,k,b:n=s,d} w^{sd,k,b} \cdot b \end{array} \right) \\
&\quad + \sum_{(i,j)\in E} h_{ij} \cdot \left(\left\lfloor \frac{\ell_{ij}}{\Lambda} \right\rfloor \right. \\
&\quad \left. \times \left(\Psi_{ij} + \delta_{ij} \cdot \sum_{sd,k,b:(i,j)\in\pi^{sd,k,b}} w^{sd,k,b} \cdot b \right) \right) \qquad (2)
\end{aligned}
$$

$$
\begin{aligned}
&TC_{G(V,E)} \\
&= \sum_{n\in V} \left(\begin{array}{l} \Theta_n + \varepsilon_n^{t_1} \cdot \sum_{sd,k,b:n\in\pi^{sd,k,b},n\neq s,d} w^{sd,k,b} \\ \quad \times b \cdot (1 - x_n^{sd,k,b}) \\ \quad + \varepsilon_n^{t_2} \cdot \sum_{sd,k,b:n\in\pi^{sd,k,b},n\neq s,d} w^{sd,k,b} \\ \quad \times b \cdot x_n^{sd,k,b} + \varepsilon_n^{t_3} \cdot \sum_{sd,k,b:n=s,d} w^{sd,k,b} \cdot b \end{array} \right) \\
&\quad + \sum_{(i,j)\in E} \left(\left\lfloor \frac{\ell_{ij}}{\Lambda} \right\rfloor \right. \\
&\quad \left. \times \left(\Psi_{ij} + \delta_{ij} \cdot \sum_{sd,k,b:(i,j)\in\pi^{sd,k,b}} w^{sd,k,b} \cdot b \right) \right) \qquad (3)
\end{aligned}
$$

$$\sum_{k,b} w^{sd,k,b} = t^{sd,b} \quad \forall s,d \in V,\ \forall b \in B \qquad (4)$$

$$\sum_{sd,k,b:(i,j)\in\pi^{sd,k,b}} w^{sd,k,b} \le a_{ij} \quad \forall (i,j) \in E, \forall b \in B \qquad (5)$$

$$w^{sd,k,b} \in N, \quad \forall s,d \in V, \forall k, \forall b \in B \qquad (6)$$

The objective (1) is the minimization of the network power consumption $DC_{G(V,E)}$ due to the network elements powered by dirty energy sources (as we want to minimize GHG emissions) and—among the solutions at minimum power consumption—the minimization of the total power consumption of the network $DC_{G(V,E)}$. Equation (2) sets the power consumption of the network elements in $G(V,E)$ due only to dirty power sources, whilst (3) indicates the total power consumption of the network elements in $G(V,E)$ evaluated in the energy model. Constraint (4) selects the routes for the lightpaths among the k pre-computed ones and assures that the whole traffic demand matrix is satisfied. Constraint (5) ensures that the maximum number of lightpaths passing on a link does not exceed the number of available wavelengths on that link. Constraint (6) imposes the integrality of the ILP problem by forcing integer values for the variables $w^{sd,k,b}$. Note that the fixed power consumptions in (2) and (3) are reported only for completeness sake but they are not involved in the optimization process (as sleep mode is not considered, the optimization is realized only on the variable energy consumptions).

Table 4 Summary of the notation used for the ILP model

Input parameters	Meaning
$G(V, E)$	directed graph representing the physical network topology; V set of vertices that represent the network nodes; E the set of edges that represent the network links; $\lvert V\rvert = N, \lvert E\rvert = M$. Each network link has a different (maximum) bitrate (i.e., bandwidth capacity);
a_{ij}	number of wavelengths available on link (i, j);
b_{ij}	bandwidth capacity of link (i, j);
ℓ_{ij}	length of link (i, j) (in km);
Λ	maximum length of links without need of amplification (in km);
$t^{sd,b}$	number of lightpaths to be established from s to d with required bandwidth $b \in B$; bandwidth ranges from 54 Mbps (1 OC-unit) to 40 Gbps (768 OC-units); $OC\text{-}units = \{1, 3, 12, 24, 48, 192, 768\}$; $\{t^{s,d}\}_{s,d \in V}$ is the traffic matrix;
$\pi^{sd,k,b}$	k-th pre-computed route[a] from s to d satisfying the bandwidth requirement of b bps;
$\rho^{sd,k,b}$	the geographical length of route $\pi^{sd,k,b}$ (in km);
Θ_n	fixed power (W) of node n;
$\varepsilon_n^{t_1}$	proportional energy (nJ/bit) for transporting one bit of *transparent pass-through* traffic at node n;
$\varepsilon_n^{t_2}$	proportional energy (nJ/bit) for transporting one bit of *opaque pass-through* traffic at node n (e.g. 3R regeneration or opaque wavelength conversion);
$\varepsilon_n^{t_3}$	proportional energy (nJ/bit) for *add/drop* one bit at node n;
Ψ_{ij}	fixed power (W) for devices in link (i, j), (e.g. optical amplifiers);
δ_{ij}	proportional energy (nJ/bit) for transporting one bit through link (i, j); it is assumed that each device (e.g. OA) on the same link (i, j) has the same fixed and proportional consumptions;
$x_n^{sd,k,b}$	identify the presence[b] of O/E/O conversion at the node n: $x_n^{sd,k,b} = \begin{cases} 1 & \text{if } n \in \pi^{sd,k,b} \text{ and } \pi^{sd,k,b} \text{ undergoes O/E/O conversion at node } n \\ 0 & \text{if } n \notin \pi^{sd,k,b} \text{ or } \pi^{sd,k,b} \text{ transparently passes through node } n \end{cases}$
g_n	identifies the presence of dirty energy source at node n: $\forall n \in V, g_n = \begin{cases} 0 & \text{if node } n \text{ is powered by a green energy source} \\ 1 & \text{if node } n \text{ is powered by a dirty energy source} \end{cases}$
h_{ij}	identifies the presence of green energy source at link (i, j): $\forall (i, j) \in E, h_{ij} = \begin{cases} 0 & \text{if link } (i, j) \text{ is powered by a green energy source} \\ 1 & \text{if link } (i, j) \text{ is powered by a dirty energy source} \end{cases}$

Variables	Meaning
$w^{sd,k,b}$	integer, indicates the number of lightpaths using route $\pi^{sd,k,b}$ (on the same route there may be several lightpaths using different wavelengths);
$TC_{G(V,E)}$	indicates the total power consumption of the NEs in $G(V, E)$ evaluated in the chosen traffic model;
$DC_{G(V,E)}$	indicates the power consumption of the NEs in $G(V, E)$ due only to dirty power sources.

[a] In this ILP formulation, a set of pre-computed routes is used for routing the demand lightpath in order to reduce the time complexity, leading to a sub-optimal solution of the ILP. The k paths satisfy the requirement on the bandwidth (b bps) since they are found by the *bandwidth constrained* k-shortest paths algorithm.

[b] Note that 3R regeneration and opaque wavelength conversion are implicitly considered in this matrix and this information will be used in the power consumption calculus

3.2.2 *Energy-aware RWA at minimum power consumption (MinPower-RWA)*

The objective of the *MinPower-RWA* problem is to minimize the overall power consumption regardless of the energy sources types and, thus, of the GHG emissions. The set of the input parameters is the same as the *MinGas-RWA* problem except for the g_n and h_{ij} vectors which are no longer necessary; also, an additional constant ξ is considered,

$$\xi : 0 < \xi \cdot \left(\sum_{n \in V} \Theta_n + \sum_{(i,j) \in E} \Psi_{ij} \right) < 1 \tag{7}$$

The mathematical formulation of *MinPower-RWA* is the same as above, with a different objective function:

$$\mathit{Minimize}\ TC_{G(V,E)} + \xi \cdot \sum_{sd,k,b} w^{sd,k,b} \cdot \rho^{sd,k,b} \tag{8}$$

and taking (3) (4) (5) (6) as constraints.

The objective function (8) is the minimization of the total network power consumption due to fixed and proportional power consumed by all the devices installed in the network, and—among the solutions at minimum power consumption—the minimization of the installation cost.

3.2.3 Minimum Cost RWA (MinCost-RWA)

The objective of the *MinCost-RWA* problem is the minimization of the installation cost regardless of the NEs energy consumptions and GHG emissions. It will try to aggregate as much lightpaths as possible while minimizing their physical lengths. The objective function in this case is:

$$\mathit{Minimize} \sum_{sd,k,b} w^{sd,k,b} \cdot \rho^{sd,k,b} \tag{9}$$

and the constraints are those of (4) (5) (6).

4 Model evaluation

In this section, we analyze the model effectiveness through ILP optimizations exploiting minimum power consumption, minimum GHG emissions and mínimum installation cost through simulation on the well-known NSFNET network topology. The obtained results have been briefly discussed to show the potential benefits achievable through the presented cross-layer optimization approach.

4.1 The proof of concept simulation environment

We used the NSFNET core optical network with 14 nodes and 21 bidirectional fiber links each with 16 wavelengths. Simulations were performed under different power distribution systems, with green energy sources powering 25, 50 and 75% of the NEs and randomly generated traffic matrices. Connection requests are fully satisfied, i.e., the blocking probability is kept strictly null. To solve the ILP problems the CPLEX software tool was used on an Intel® Xeon® 2.5 GHz dual processor Linux server. The available memory (physical RAM + swap area) amounted to 16 GBytes. To reduce the notable requirements in terms of computational and memory resources, we first bound the problem dimension by restricting the paths' alternatives to a static set of k pre-computed routes, obtained by using a traditional K-SPF algorithm and hence satisfying the traditional network management objectives without considering any energy-related information. Secondly, we limited the depth of the branch-and-bound/cut algorithms after calculating a pre-definite number of integer solutions. While such simplification techniques are certainly useful to contain the computational burden, the solution they produce is an approximation of the actual optimal (in terms of power consumption) virtual network topology built on the available physical infrastructure. However in these cases the ILP approach maintains its added value, as far as the approximated solutions can be close to the exact one. Some of the selected paths would probably not be the best ones, but the resulting power savings could be substantial without significant losses on the other optimization objectives.

4.2 Results and discussion

The energy consumption (during 1 year time period) resulting from the three ILP RWA strategies with 50% of the NEs powered by green energy sources is reported in Fig. 9. As expected, the *MinCost-RWA* is the most energy consuming strategy, whilst the *MinPower-RWA* is the best strategy as for the energy consumption, but the best one as GHG emissions is the *MinGas-RWA*. Anyway, the difference in energy consumption between the latter two strategies is lower than 14% in the worst case. This result was somehow expected, as the minimum power RWA strategy attempts to save as much energy as possible regardless of the sources of energy, whereas the minimum GHG emissions may route the lightpaths on longer—thus, more energy consuming—paths but preferring those NEs that are powered by green energy sources. At low loads, *MinGas-RWA* attempts to use only green-powered nodes, at the expense of possibly choosing longer paths. The effect of these suboptimal choices is visible at higher loads, when the overall energy consumption rises more steeply that of *MinPower-RWA*. This becomes relevant at network loads as high as 70%, whereas in the 30%–70% operating range the savings achieved by *MinGas-RWA* with respect to *MinCost-RWA* remain consistently substantial. As for the energy consumption, compared with *MinCost-RWA, MinGas-RWA* saved an average of 18% of energy while *MinPower-RWA* reached savings up to 30% of the overall energy consumption.

Besides the saving in energy consumption, *MinGas-RWA* achieves to save also considerable quantity of CO_2. For a medium loaded network (50% of routed lightpaths), where one half of the NEs are powered by green power plants and the other half are powered by fuel-based power plants, *MinGas-RWA* strategy saves an average of 37,500 kg of CO_2 per year (see [1] as a reference value for the emitted CO_2).

In Fig. 10, we compared the estimated CO_2 emissions with the three strategies at different network loads, where one half of the NEs are powered by green energy sources and

Fig. 9 Network energy consumption *vs* traffic load with the three ILP strategies

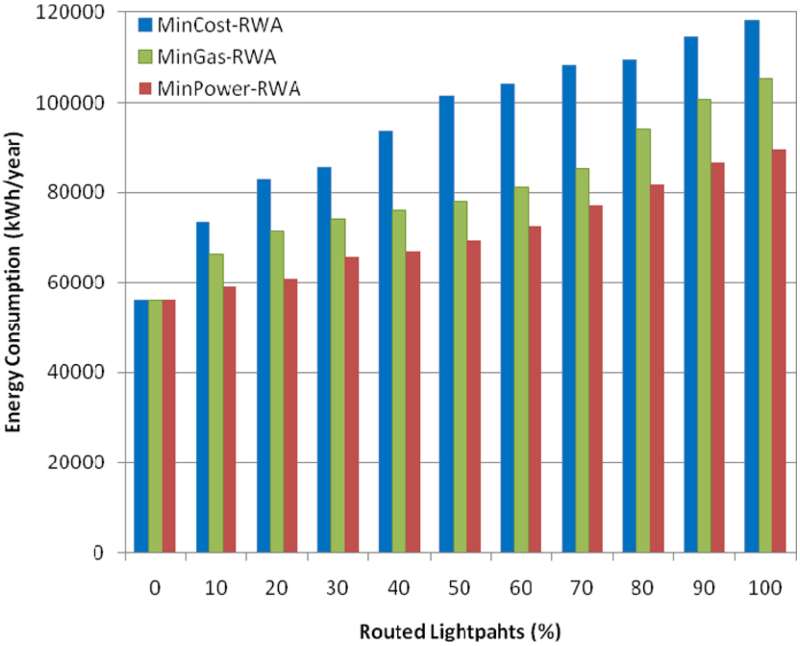

Fig. 10 Network CO_2 emissions *vs* network load with the three ILP strategies

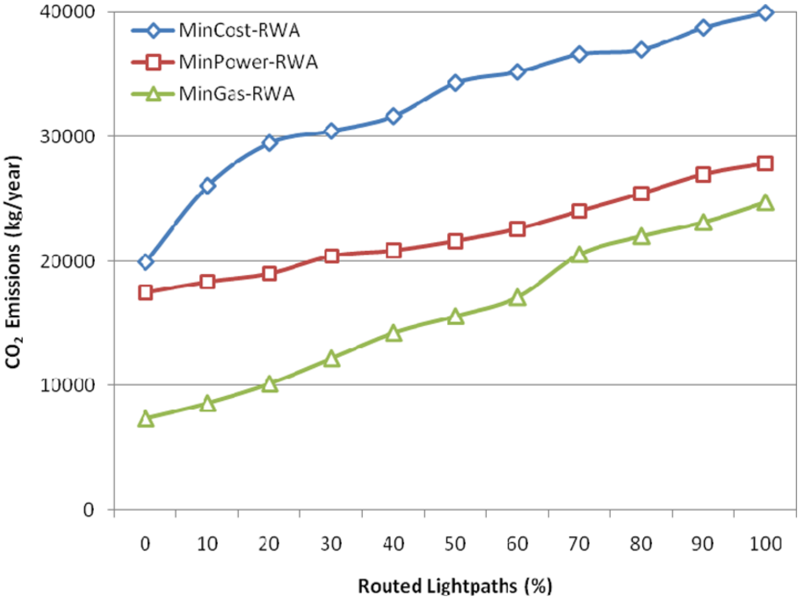

the other half by fuel-based power plants. As can be seen, at low loads the *MinGas-RWA* strategy achieves prominent CO_2 savings (only about 33% of CO_2 were emitted with respect to *MinCost-RWA* and about 50% relative to *MinPower-RWA*), whilst, as the network load increases, the difference between the *MinGas-RWA* and the *MinPower-RWA* strategies decreases, because at higher loads it becomes more and more difficult to satisfy the demand without resorting to dirty-powered nodes. In other words, at high loads, minimizing the overall power consumption implicitly leads to the minimization of the concomitant CO_2 emissions, while at midrange loads the CO_2 savings induced by *MinGas-RWA* are significant.

We also explored the power consumptions when different lightpath bitrates are considered. In the simulation, the network load was kept constant (at 50% of its maximum capacity) while connection requests with different bandwidth requirements (bitrates) are generated, ranging from many low-speed connections to few high-speed ones (Fig. 11). As a general trend, we observe that—though the traffic load is constant—higher bitrates are associated with higher power consumptions. This behavior is due to the fact that smaller connections have more possibilities to be routed over less energy-demanding routes than larger ones which are instead more likely to be routed over high capacity network routes. *MinCost-RWA* power consumption grows quite linearly with

Fig. 11 Network power consumption with different lightpath rates (constant traffic load at 50%)

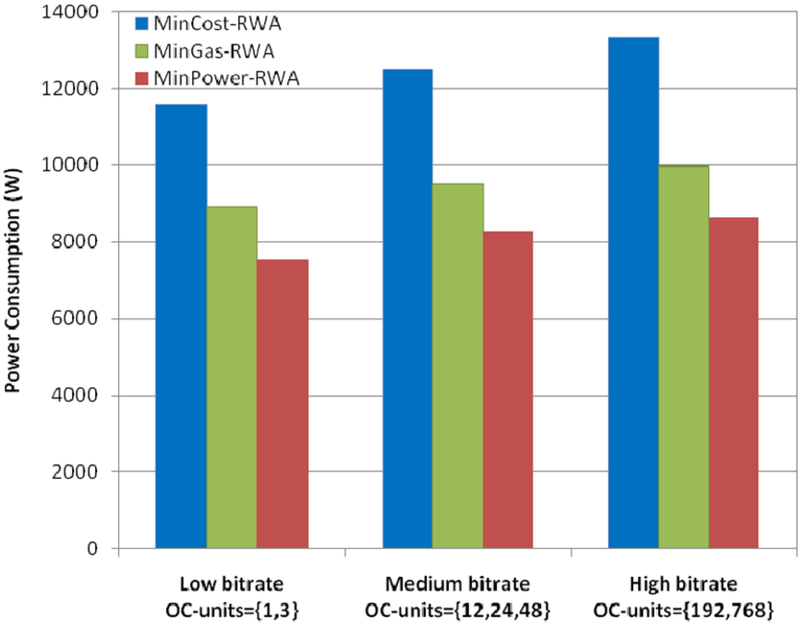

the increasing of the lightpath bitrates, whilst the *MinGas-RWA* and, above all, the *MinPower-RWA* exhibit a more constant behavior: thanks to the energy-awareness of such ILPs, they are able to accommodate more profitably the connection requests, even if also in these strategies an increase in the power consumption is still observed due to the lower eligible routes. We also note that the *MinCost-RWA* power consumption is always higher than with the other two strategies even at higher bitrates, showing that the energy-awareness may help to substantially compensate the higher energy consumption due to higher bitrates.

Finally, we have analyzed the dependency of the power consumption from the actual values of the fixed and variable components of the power draw by an interface, expressed as a function of the link rate at which the interface operates. Note that, since a (unidirectional) link is attached to each interface, the set of links E in the aforementioned network graph representation $G(V,E)$ actually coincides with the set of interfaces. Each interface has its own native speed: $\forall i \in E, v_i \in R = \{10\text{ Mbps}, 100\text{ Mbps}, 1000\text{ Mbps}, 10000\text{ Mbps}\}$ represents the *native link rate* of interface i. If the link rate is fixed, the power draw of an interface will depend mainly on the link rate, with minor variations due to architecture, circuitry, and components. When using an ALR, instead, the power consumption of an interface i depends not only on the *working* link rate r_i but also on the *native* link rate v_i. In other words, a given throughput t_d results in different power consumption depending on the interface native link rate v_i: in this case, slower interfaces consume less power than faster ones for the same throughput t_d, even if they work at the same rate r_i. This result, quite surprising if we consider that slower interfaces consume more energy per bit than faster ones, may be explained by considering the different technologies adopted for reaching higher link rates (mainly based on advanced modulation techniques [33]) that lead to greater fixed power consumption for faster interfaces. In fact, like routers, also the interfaces have fixed and variable power consumption. The fixed part is always present just for the interface to stay up and accounts for the control circuits, while the variable traffic-proportional power consumption is due to the transceivers. In the following we model such energy consumptions and show a breakdown of the different energy components in a 10 Gbps interface.

In general, to model the fixed and the variable energy consumption, we define $\{\Psi(v_i, r_j) | j = 1, 2, \ldots, m\}$ where $\Psi(v_i, r_j)$ (see Table 5) is the power consumption of the interface $i \in E$ with native speed $v_i \in R$ operating at link rate $r_j \in R$ and $\Psi(v_i, r_j) < \Psi(v_i, r_k)\ \forall j < k$.

In Fig. 12 we plotted the CO_2 emissions as a function of the average link rate, assuming a uniform distribution of the working link rates between 0 and the native link rate and a real-world distribution of values for the $\Psi(v_i, r_j)$ as given in [1, 23]. On the horizontal axis there is the percentage of interfaces operating at the native link rate. A notable characteristic is that the savings induced by an extensive use of an adaptive link rate are not as dramatic as one may expect, although sensitivity is slightly higher in the case of *MinPower-RWA* and *MinGas-RWA*.

5 Conclusions

The ICT sector has the fundamental capability of acting as drawing factor to drive the development of energy-efficient technological innovations for both industry and society, although any action in the direction of energy-efficiency may

Table 5 Power consumptions of interfaces working at different rates

v_i	r_i	Mbps	Power consumption
$\forall v_i \in R$	Off $/r_0$	0/0	$\Psi(v_i, \mathit{Off}/r_0) \cong 0$[a]
v_1: 10	r_1	10	$\Psi(v_1, r_1)$
v_2: 10^2	r_1/r_2	$10/10^2$	$\Psi(v_2, r_1)/\Psi(v_2, r_2)$
v_3: 10^3	$r_1/r_2/r_3$	$10/10^2/10^3$	$\Psi(v_3, r_1)/\Psi(v_3, r_2)/\Psi(v_3, r_3)$
v_4: 10^4	$r_1/r_2/r_3/r_4$	$10/10^2/10^3/10^4$	$\Psi(v_4, r_1)/\Psi(v_4, r_2)/\Psi(v_4, r_3)/\Psi(v_4, r_4)$

[a] in LPI, the device only send signals during short refresh intervals and stay quite during large intervals so the power consumption in the LPI mode is almost 0

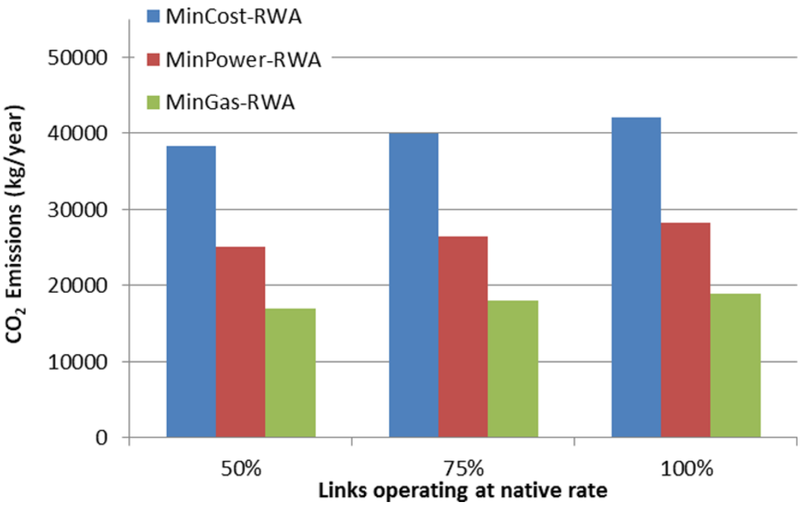

Fig. 12 Average emitted CO_2 at different ratios of links operating at their native rate

result in direct power and cost savings in the short and medium term, while other indirect effects will be only observed on the long term on both the environment and human health. Anyway, the massive introduction of energy efficiency within the network world requires a coordinated effort of equipment vendors, governments, and service providers to identify technological standards, best practices, and solutions to support the necessary changes in the basic construction and functional requirements for network equipment and control plane algorithms. Accordingly, several energy-aware ILP formulations exploiting dual energy sources have been presented along with an energy model in which no sleep mode is available but the optimization relies only on the traffic-variable power consumption of the NEs. Two ILP formulations have been presented: minimum power (*MinPower-RWA*) and minimum GHG emissions (*MinGas-RWA*) strategies with the objectives to minimize respectively the absorbed energy and the emitted GHG. Results show that the *MinPower-RWA* strategy may save a considerable amount of energy by routing the lightpaths on minimum consuming NEs and that the GHG emitted may be notably reduced by the *MinGas-RWA* strategy that prefers NEs powered by green energy sources. As drops are observed in the day/night traffic at core network nodes, there is room for some possible optimizations by putting NEs into sleep mode only partially (per-interface sleep mode). In fact, putting into sleep mode single interfaces or line cards may have some sense, saving up to 50% of the total router power, but modifications to current router architecture and routing protocols need to be investigated. Renewable energy sources may vary their availability with time (e.g. solar panels only generate electricity during the day). While in the current work we handled the availability of green and dirty sources in a static way, in future works statistically variable green energy sources may be considered within a totally dynamic scenario in which the availability of the different types of renewable energy sources can be associated with the variations of the day time and traffic load (e.g. night/day cycle). We are confident that the above efforts, together with incrementing network eco-sustainability, will improve the sustainable growth and—in the long run—the society prosperity.

Acknowledgements This work was supported in part by the COST Action IC0804 on Energy Efficiency in Large Scale Distributed Systems, the Spanish Ministry of Science and Innovation under the DOMINO project (TEC2010-18522), the Catalan Government under the contract SGR 1140 and the DIUE/ESF under the grant FI-201000740.

References

1. BONE project (2009). *WP 21 topical project green optical networks* (Report on year 1 and updated plan for activities). NoE, FP7-ICT-2007-1 216863 BONE project, Dec. 2009.
2. Baliga, J., et al. (2009). Energy consumption in optical IP networks. *Journal of Lightwave Technology*, *27*(13), 2391–2403.
3. Moore, G. E. (1998). Cramming more components onto integrated circuits. *Proceedings of the IEEE*, *86*, 82–85.
4. Gilder, G. F. (2000). *Telecosm: how infinite bandwidth will revolutionize our world*. New York: The Free Press.
5. Jevons, W. S. (1866). *The coal question; an inquiry concerning the progress of the nation, and the probable exhaustion of our coalmines*. London: Macmillan.
6. Smart Grid Task Forces (2011). Office of Electricity Delivery & Energy Reliability, Washington, DC. http://www.smartgrid.gov/.
7. Karg, L. (2010). Consult GmbH, *keynote speech*. In *e-Energy 2010*, Passau, Germany Apr. 2010.
8. Lam, W. (2009). Getting the most out of Google's solar panels. Jul. 2009 [online]. Available: https://docs.google.com/present/view?id=dfhw7d9z_0gtk9bsgc.
9. St Arnaud, B. (2011). ICT and Global Warming: Opportunities for Innovation and Economic Growth. http://docs.google.com/Doc?id=dgbgjrct_2767dxpbdvcf.
10. Chabarek, J., Sommers, J., Barford, P., Estan, C., Tsiang, D., & Wright, S. (2008). Power awareness in network design and routing. In *Proc. IEEE INFOCOM*.
11. Gupta, M., & Singh, S. (2003). Greening of the Internet. In *Proc. of the ACM SIGCOMM*, Karlsruhe, Germany
12. Feng, M. Z., Hilton, K., Ayre, R., & Tucker, R. (2010). Reducing NGN energy consumption with IP/SDH/WDM. In *Proceedings of the 1st international conference on energy-efficient computing and networking*, Passau, Germany (pp. 187–190).
13. Vereecken, W., Van Heddeghem, W., Colle, D., Pickavet, M., & Demeester, P. (2010). Overall ICT footprint and green communication technologies. In *Proc. of ISCCSP 2010*, Limassol, Cyprus Mar. 2010.
14. Ricciardi, S., Careglio, D., Fiore, U., Palmieri, F., Santos-Boada, G., & Solé-Pareta, J. (2011). Analyzing local strategies for energy efficient networking. In *Proceedings of sustainable networking SUNSET 2011, IFIP networking 2011*, Valencia, 9–13 May 2011.
15. Saleh, B. E. A., & Teich, M. C. (1991). *Fundamentals of photonics*. New York: Wiley.
16. The Energy Consumption Rating (ECR) initiative (2011). [online]. Available: http://www.ecrinitiative.org/.
17. Miller, D. A. B. (2010). Are optical transistors the logical next step? *Nature Photonics*, *4*(1), 3–5.
18. D'Ambrosia, J. 100 Gigabit Ethernet and beyond. *IEEE Communications Magazine*, March 2010.
19. Anderson, J., & Traverso, M. (2010). Optical transceivers for 100 gigabit Ethernet and its transport. *IEEE Communications Magazine*, *48*(3), S35–S40.
20. Lange, C. (2009). Energy-related aspects in backbone networks. In *Proc. ECOC 2009*, Vienna, Austria, Sep. 2009.
21. Tucker, R. S., et al. (2009). Evolution of WDM optical IP networks: a cost and energy perspective. *IEEE/OSA Journal of Lightwave Technologies*, *27*(3), 243–252.
22. The Green Grid (2008). *The green grid data center power efficiency metrics: PUE and DCiE*. Technical Committee White Paper.
23. Christensen, K., Reviriego, P., Nordman, B., Bennett, M., Mostowfi, M., & Maestro, J. A. (2010). IEEE 802.3az: the road to energy efficient Ethernet. *IEEE Communications Magazine*, *48*(11), 50–56.
24. Hays, R. (2008). Active/idle toggling with low-power idle. In *IEEE 802.3az task force group meeting*. [online]. Available: http://www.ieee802.org/3/az/public/jan08/hays_01_0108.pdf.
25. Zhai, B., Blaauw, D., et al. (2004). Theoretical and practical limits of dynamic voltage scaling. In *DAC*.
26. Christensen, K., & Nordman, B. (2005). Reducing the energy consumption of networked devices, IEEE 802.3 tutorial, July 19, 2005, San Francisco [online]. Available: http://www.csee.usf.edu/~christen/energy/ieee_tutorial.pdf.
27. Van Heddeghem, W., De Groote, M., Vereecken, W., Colle, D., Pickavet, M., & Demeester, P. (2010). Energy-efficiency in telecommunications networks: link-by-link versus end-to-end grooming. In *Proc. of ONDM 2010*, Kyoto, Japan, Feb. 1–3 2010.
28. Tucker, R. S. (2011). Modelling Energy Consumption in IP Networks, [online]. Available: http://www.cisco.com/web/about/ac50/ac207/crc_new/events/assets/cgrs_energy_consumption_ip.pdf.
29. Aleksic, S. (2009). Analysis of power consumption in future high-capacity network nodes. *Journal of Optical Communications and Networking*, *1*(3), 245–258.
30. Energy Star (2011). Small network equipment [online]. Available: http://www.energystar.gov/index.cfm?c=new_specs.small_network_equip.
31. Ricciardi, S., Careglio, D., Palmieri, F., Fiore, U., Santos-Boada, G., & Solé-Pareta, J. (2010). Energy-oriented models for WDM networks. In *Proceedings of 7th international ICST conference on broadband communications, networks, and systems (Broadnets 2010)*, Athens, Greece, 25–27 Oct. 2010 (pp. 1–4).
32. Ricciardi, S., Careglio, D., Palmieri, F., Fiore, U., Santos-Boada, G., & Solé-Pareta, J. (2011). Energy-aware RWA for WDM networks with dual power sources. In *Proceedings of 2011 IEEE international conference on communications (ICC 2011)*, Kyoto, Japan, June 5–9, 2011.
33. Nortel (2011). *A comparison of next-generation 40-Gbps technologies*. White paper [online]. Available: http://www.nortel.com/solutions/collateral/nn122640.pdf.
34. Koroneos, C. J., & Koroneos, Y. (2007). Renewable energy systems: the environmental impact approach. *International Journal of Global Energy Issues*, *27*(4), 425–441.

Sergio Ricciardi received the degree summa cum laude in Computer Science from the University of Naples Federico II, Italy, in 2006 and the M.Sc. degree with honours in Computer Architecture, Networks and Systems from the Technical University of Catalonia (UPC), Spain, in 2010. He worked with the Federico II University and with the Italian National Institute for Nuclear Physics (INFN) within several national and international projects. From 2008 he is research associate in the Advanced Broadband Communications Center (CCABA) at the Department of Computer Architecture of the UPC. His current activities concern energy-efficient architectures and energy-aware RWA algorithms and protocols for optical networks and grid/cloud computing. His research interests are mainly focused on energy-oriented routing, optimization algorithms and topology management for transparent and opaque optical networks.

Davide Careglio (S'05–M'06) received the M.Sc. and Ph.D. degrees in telecommunications engineering both from Universitat Politècnica de Catalunya (UPC), Barcelona, Spain, in 2000 and 2005, respectively, and the Laurea degree in electrical engineering from Politecnico di Torino, Turin, Italy, in 2001. He is currently an Associate Professor in the Department of Computer Architecture at UPC. Since 2000, he has been a Staff Member of the Advanced Broadband Communication Center. His research interests include networking protocols with emphasis on optical switching technologies, and algorithms and protocols for traffic engineering and QoS provisioning. He is the coauthor of more than 80 publications in international journals and conferences. He has participated in many European and national projects in the field of optical networking and green communication.

Germán Santos-Boada obtained his M.Sc. degree in Telecom Engineering in 1978, and his Ph.D. in 1993, both from the Technical University of Catalonia (UPC). He worked for Telefónica as manager of engineering from 1984 up to 2007 and simultaneously he joined the Computer Architecture Department of UPC as a partial time Assistant Professor. Currently he is full time Assistant Professor with this department. Dr. Santos current research interests are Quality of Service provisioning in next generation optical access networks and optical energy-aware network modeling. He is currently involved in the COST 804 action. (german@ac.upc.edu)

Josep Solé-Pareta obtained his M.Sc. degree in Telecom Engineering in 1984, and his Ph.D. in Computer Science in 1991, both from the Technical University of Catalonia (UPC). In 1984 he joined the Computer Architecture Department of UPC. Currently he is Full Professor with this department. He did a Postdoc stage (summers of 1993 and 1994) at the Georgia Institute of Technology. He is co-founder of the UPC-CCABA, and UPC-N3cat. His publications include several book chapters and more than 150 papers in relevant research journals (>25), and refereed international conferences. His current research interests are in Nanonetworking Communications, Traffic Monitoring, Analysis and High Speed and Optical Networking and Energy Efficient Transport Networks, with emphasis on traffic engineering, traffic characterization, MAC protocols and QoS provisioning. He has participated in many European projects dealing with Computer Networking topics.

Ugo Fiore leads the Network Operations Center at the Federico II University, in Naples. He began his career with Italian National Council for Research and has also more than 10 years of experience in the industry, developing software support systems for telco operators. His research interests focus on optimization techniques and algorithms aiming at improving the performance of high-speed core networks. He is also actively pursuing two other research directions: the application of nonlinear techniques to the analysis and classification of traffic; security-related algorithms and protocols.

Francesco Palmieri is an assistant professor at the Engineering Faculty of the Second University of Napoli, Italy. His major research interests concern high performance and evolutionary networking protocols and architectures, routing algorithms and network security. Since 1989, he has worked for several international companies on networking-related projects and, starting from 1997, and until 2010 he has been the Director of the telecommunication and networking division of the Federico II University, in Napoli, Italy. He has been closely involved with the development of the Internet in Italy as a senior member of the Technical-Scientific Advisory Committee and of the CSIRT of the Italian NREN GARR. He has published a significant number of papers in leading technical journals and conferences and given many invited talks and keynote speeches.