



Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

A hybrid load-balancing and energy-aware RWA algorithm for telecommunication networks

Sergio Ricciardi^{a,*}, David Sembroiz-Ausejo^a, Francesco Palmieri^b, Germán Santos-Boada^a,
Jordi Perelló^a, Davide Careglio^a

^a Department of Computer Architecture, Technical University of Catalonia – BarcelonaTech (UPC), Barcelona Spain

^b Department of Computer Science, University of Salerno, Salerno, Italy

ARTICLE INFO

Article history:

Received 16 October 2014

Revised 26 May 2015

Accepted 16 June 2015

Available online xxx

Keywords:

Energy-efficiency

Load-balancing

RWA

ABSTRACT

In the last years, the power consumption of telecommunication networks has attracted the attention of both researchers and field experts in order to contain the associated energy bills and reduce their ecological impact. Many of the proposed solutions have been focused exclusively on the reduction of the power consumption, without adequately considering more traditional network engineering objectives such as balancing resource utilization, routing policy, or resilience schemes. As a consequence, network control plane strategies passed from one extreme to the other, from being totally energy-unaware to exclusively energy-efficient at the expenses of load-balancing, with obvious impacts on the power consumption in the former case and on the blocking rate in the latter one.

In this paper, we present a hybrid routing and wavelength assignment algorithm that, when the network is lightly loaded, operates in an energy-efficient way, by routing the connections on the paths requiring the lowest amount of energy, while, when the network load increases, it dynamically switches to a pure load-balancing scheme in order to best allocate the available communication resources. The switching decision among load-balancing and energy-awareness is taken dynamically, driven by a threshold on the number of new connections requests reaching the network during a prefixed time window.

Simulation results show the effectiveness of the hybrid algorithm, which achieves lower energy consumption than a pure load-balancing algorithm while keeping the network load fairly distributed on the available resources.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

In the last years the Internet traffic has been rising according to an exponential trend, leading to increased bandwidth demands which in turn has resulted in more sophisticated and energy-hungry communication equipment that affect significantly the networks' operational expenses. At the state of the art, telecommunications networks infrastructures require more than 1% of the worldwide production of electrical energy and the demand rate increases of about 12% per year [1]. These numbers give an immediate idea of the impact of the energy bills on the overall economy of most of the large-scale network communication provider organizations. The growth of energy demand together with the limited availability of new and more clean energy sources (where energy production is not based on

burning fossil fuel) also introduces significant issues in terms of ecological impact. Until most of the energy needed for network operations will not be drained from renewable, clean energy sources, a large amount of Greenhouse gases (GHG) will continue to be emitted into the atmosphere as a consequence of the direct and indirect usage of the communication equipment.

The only viable strategies for containing the power consumption in modern communication networks rely on the energy proportional behavior of most of the new generation equipment, that are able to adapt their own energy demand to the effective workload by dynamically switching between several operating states, each characterized by a higher or lower component performance (e.g., interfaces, memories, switching fabric, etc.). These strategies require the introduction of energy-awareness in the network control plane, providing visibility of the energy efficiency degree of all the network equipment in order to route communications on paths traversing switching devices and links characterized by a lower proportional absorption, and hence minimizing the load on the most energy-hungry devices.

* Corresponding author. Tel.: +34 93 40 17 182.

E-mail addresses: sergior@ac.upc.edu (S. Ricciardi), fpalmieri@unisa.it (F. Palmieri).

However, this may introduce several undesirable side effects on the overall network optimization policies and objectives. That is, while being effective in substantially reducing the energy costs and/or associated ecological footprint, routing strategies that strive to divert connections over the paths requiring minimum energy, without considering all the other more traditional traffic engineering objectives, may result in almost blind choices that tend to unbalance the network load and under-utilize many expensive communication links, with obvious consequences in terms of return of investment (ROI).

From the previous considerations, it is immediately evident how control-plane strategies aiming at balancing resource utilization or containing the energy consumption as well as GHG emission, can easily become mutually contradictory. To cope with this problem, we propose a hybrid load-balancing and energy-aware routing and wavelength assignment (RWA) algorithm (*HyLERA*) whose goal is to combine, according to a threshold-based scheme, the aforementioned energy demand optimization and resource usage balancing strategies in order to achieve more stable effects on the overall network engineering economy, by using *at any time* the most appropriate strategy. Thus it drives lightpath selection based exclusively on a load-balancing objective when the risk of request blocking grows over a certain threshold (i.e., the network is experiencing an overload condition), and it performs its connection routing and wavelength assignment choices in order to reduce the overall energy consumption when the network is unloaded.

The rest of the work is organized as follows. In the related work section we analyze the RWA experiences already available in literature that follow load-balancing or energy-aware approaches as well as those ones which try to unify or hybridize both of them. The next section defines the methodology proposed in *HyLERA* and the following one analyzes its performances through simulation and discusses the results obtained. Finally, the last section reports the conclusions and the future research plans.

2. Related work

In research efforts concerning RWA in optical telecommunications networks [2,3], two different lines, often providing very different results in terms of optimization objectives, have been developed.

At first, we can consider pure load-balancing strategies. This line of research focuses on developing traffic load-aware routing algorithms mainly pursuing the optimization of network resources usage together with the minimization of the service request rejection (connection blocking) in presence of specific QoS or impairment constraints. For example, in [4], a distributed algorithm for constraint-based path selection in WDM networks is presented, focusing on service-specific path quality attributes, traffic load and affinity conditions between previously routed connections. Analogously, the distributed dynamic routing approach proposed in [5] provides a common RWA framework considering traffic demands at the wavelength level and using additive and multiplicative metrics to achieve protection in case of single link failure. In [6] an integrated routing and grooming scheme for WDM networks based on a hierarchical network model using the concept of blocking islands to abstract network resources and keeping the integrity and load-balance of specific network sections (the above islands) as intact as possible. Also [7] proposes a routing optimization algorithm where connection requests characterized by the same source and destination are groomed and served according to a specific priority order. Another combined RWA/grooming approach implemented according to a two-stage path selection approach and driven by connection blocking minimization criteria is presented in [8], whereas a more sophisticated load-balancing approach based on maximum flow/minimum cut considerations on the network graph is presented in [9]. Finally, the work in [10] presents an off-line version of the problem that can

be used when the topology of the network, the number of available wavelengths and the traffic matrix are known in advance.

Secondly, there is another line in wavelength routing research, orthogonal to the previous one, that is based on optimizing the overall energy demand and the ecological footprint of the network. In [11] the authors envisioned a new management approach in communication networks oriented by energy efficiency and ecological awareness criteria. Both the above concepts have to be considered in modern energy-oriented RWA paradigms encompassing renewable energy sources and energy efficient devices in a systemic approach that spans the whole life-cycle assessment of the next generation networks. Accordingly, in [12,13] the reduction of both energy costs and GHG emission through use of renewable energy sources has been investigated. To achieve the needed optimizations new routing algorithms are envisioned, taking into consideration the availability of low-power or energy-proportional operating modes [14] as well as the type of power feed given to the wavelength switches (green, dirty), according to the information provided by the Smart electric grids [15]. Also in [16] a dynamic routing approach, with paths that are established on-demand between nodes powered with renewable energy, is proposed. Other authors have presented in [17] an algorithm that proposes a new eco-friendly distributed routing scheme, based on OSPF, that enables the traffic routing through a small subset of routers, allowing unneeded routers to pass into sleep mode by reducing energy consumption and hence GHG emissions. This work considers the network load as a determining factor in the decision-making. A similar eco-sustainable routing approach based on specific OSPF extensions has been presented in [18]. In [19] the authors consider the energy consumption of the dedicated path protection resources and investigate the savings achievable through the use of sleep mode, by discussing the trade-offs between energy saving and blocking probability. Similarly, in [20], the trade-off between energy consumption and modulation formats is discussed, showing that the most power efficient transponders do not always lead to the lowest overall network energy consumption, basically due to the intermediate 3R regeneration over long distances. An hybrid CDN-P2P system in an IP-over-WDM network is considered in [21], and the energy consumption is reduced by means of an energy-aware algorithm which besides reduces the load on the CDN. In [22], the authors compare the energy consumption of the OFDM and MLR optical transport technologies under different network traffic scenarios, and evaluate the results by employing a robust optimization technique.

In spite of the above advances in pure load-balancing or pure energy-efficient approaches, there are also some routing optimization research efforts considering a mixed situation in which both the balancing and the energy/ecological footprint are simultaneously used in the definition of hybrid algorithms. In [23] the authors present a novel integrated RWA framework (*GreenSpark*) using the traditional balancing objectives together with energy-awareness in its decision process, also making use of green energy sources wherever possible. The effectiveness of the above approach and its impact on the power consumption and carbon footprint of telecommunication networks have been evaluated, showing very interesting results at the expense of a small increase in the computational complexity. In this context, the authors of [24] proposed a *Weighted Power-Aware Lightpath Routing (WPA-LR)* scheme that leverages on a cost function that considers both the power state of the network and the information about wavelength usage and blocking performance. In this algorithm when computing each candidate path, a weight is assigned to fiber links in the network to be used by for selection of shortest paths that are sorted in ascending order based on their power consumption. We can consider *GreenSpark* and *WPA-LR* as an effort in the direction of unifying load-balancing and energy-awareness criteria by using a single hybrid algorithm characterized by a multi objective optimization behavior. However, an optimal strategy would be to operate the network exclusively targeting at a pure load-balancing

objective when the resources are limited, and at a pure energy-awareness objective when there is plenty of free resources and the objective is to minimize the power consumption and therefore the operational expenditures. Toward this goal, HyLERA, the algorithm proposed in this work, represents an important step forward, based on the adaptive adoption of different threshold-driven metrics leading to different behaviors, according to the current network traffic state, optimizing the routing decisions based on energy-awareness or unblocking conditions.

3. Energy-aware RWA in WDM networks

Network providers typically design and manage their own transport infrastructures in order to attract an ever increasing interest in their customers by providing high bandwidth, extremely reliable and quality-differentiated communication services. At the same time, they incur in significant initial investments and operating costs so that they are continuously faced with the challenge of using the most expensive equipment and communication resources as efficiently as possible, by keeping their power consumption costs at a minimum and using all the available optical links in a fairly balanced way. Unfortunately the last two objective are often in contrast, depending on the specific network topology and equipment used. However, energy expenses, due to their recurring nature become one of the most critical element in the overall operational costs. Consequently, deploying dynamic power management strategies that aim at decreasing the power demand in the operational phase and consequently reducing energy bills, becomes a fundamental prerequisite for maximizing their medium and long-term revenues. Stimulated by the above needs and by recent advancements in traffic engineering techniques, there is a growing interest in designing a unified control plane framework resulting in a smart integrated approach that combines the above resource and energy usage containment strategies into an hybrid optimization framework whose main aim is achieving the best compromise between the apparently disjoint (or, worse, conflicting) load-balancing and energy-containment objectives, by harmonizing them into a single revenue-maximization goal.

3.1. Circuit switching in multi-layer optical networks

A multi-layer optical network is a very complex mesh of variously interconnected heterogeneous sub-networks composed of an electronic IP layer (network edge), providing network access and connectivity distribution, built on top of an optical transport layer (OTN) playing the role of interconnection backbone (network core). Each sub-network consists of multiple heterogeneous switching devices, ideally operating according to a common control plane protocol and management policy, usually within the same Autonomous System (AS). In presence of very different types of devices, built by multiple vendors, all the switching decisions are based on a combination of packet, time slot, wavelength or interface depending on the location (edge or core) and role (intermediate or terminating) of the involved devices within the overall network topology.

Generalized multi-protocol label switching (GMPLS) is the control plane solution of choice that is used in most of the cases to automatically route the connections from source to destination, possibly crossing multiple fiber spans. When the connections are subject to some quality of service (QoS) requirement – e.g. required bandwidth, maximum latency, etc. – all the individual links involved in the generalized labeled switched path (LSP) – becoming a lightpath in the optical layer – have to satisfy the above requirements/constraints so that the whole path has to be determined according to a traffic or network engineering approach (this is usually referred to as MPLS-TE).

At the optical layer, wavelength division multiplexing (WDM) technology is used to multiplex several channels – each on an individual wavelength – on the same fiber. If no wavelength converters

are present in the network, a connection has to be allocated on the same wavelength in the optical domain (wavelength continuity constraint, WCC), and obviously two connections sharing the same fiber have to use different wavelengths (clash constraint). Since a single wavelength can normally carry a channel of 40 or even 100 Gbps, it is common to multiplex several tributary sub-channels in one wavelength (a process known as grooming); the wavelength resource is thus divided in time slots which will be assigned to the individual sub-channels, possibly with different occurrence distribution (in the case of statistical multiplexing).

The WCC can be relaxed by providing the optical cross connects (OXC) with wavelength converters, which are quite expensive devices; a more economically viable alternative is to extract the wavelength from the fiber, with a reconfigurable optical add and drop multiplexer (ROADM), and convert it in the electronic domain from which it can be reintroduced in the optical domain using a different wavelength – a process known as O-E-O conversion. However, due to the limiting processing speed of the electronics, such a process is usually avoided and limited only to the case in which a 3R regeneration (re-amplification, re-shaping and re-timing) of the optical signal is required in order to preserve the carried information with an acceptable bit error rate (BER).

When a connection request reaches a lambda edge router (LER), a feasible path composed of multiple communication links and nodes connecting the request source with its destination, together with an appropriate wavelength to be used on it, have to be selected in order to setup a dedicated end-to-end communication channel (circuit). Normally, there are several paths that can be selected with enough free bandwidth and satisfying the connection's QoS requirements. The final choice between the available options depends on the routing and wavelength assignment (RWA) algorithm and its optimization objectives.

3.2. RWA with hybrid optimization behavior: the HyLERA idea

Let us consider the network depicted in Fig. 1. Suppose that a request has been issued in order to connect two nodes at the IP layer. The source node, which has a complete view of the network provided by the internal routing protocol of the involved AS (such as OSPF-TE), looks for a feasible path at the optical transport layer. It is easy to see that two feasible paths are available, namely lightpath A and lightpath B (suppose there is a free wavelength or sufficient free bandwidth on both in case of grooming). Lightpath A is shorter (two hops at the optical layer), and traverses an unloaded, energy-inefficient node (i.e., currently switching few or no connections); lightpath B is longer (three hops at the optical layer), but traverses two energy-efficient nodes (e.g., modern/bigger nodes which consume less energy per bit than older/smaller ones). A pure load-balancing algorithm would select lightpath A, whereas a pure energy-aware algorithm would select lightpath B. The idea of the HyLERA algorithm is to select one or the other path according to the *current* network traffic state/operating conditions. If the network is experiencing a limited load (e.g., at night or during low load periods), it prefers optimizing energy-efficiency and hence selects paths requiring less power to be operated, even if longer in terms of traversed hops/nodes or quite unfair in terms of traffic load distribution (note that selecting a longer path will consume more resources for serving the same connection request compared to a shorter path). When the network becomes more loaded (i.e., more connections requests come into the network, for example during daytime or peak traffic hours), the algorithm automatically switches to load-balancing mode, preferring the more fairly-balanced lightpaths in order to save resources and serve the highest possible number of connection requests, eventually coming back again to energy-efficiency mode later, during lower load periods. The switch-over among the two criteria is driven by a threshold on the arrival rate of connections requests measured




Fig. 1. Path selection in a multilayer GMPLS/WDM network according to a pure load-balancing (lightpath A) or a pure energy-efficient (lightpath B) criteria; the optimal selection depends on the current network traffic state.

during a (parametric) sliding window. The rationale to use connections requests arrival rate and not directly the blocking rate is that it is difficult to find an appropriate value of the blocking rate that is effective regardless the network topology, the time of the day, the traffic pattern and the state of the network. The connections arrival rate is instead an easily tunable parameter for a network, regardless of the “accumulated traffic”, i.e. of the currently carried load of the network, and of the future incoming traffic. In this way, starting from the knowledge of the recent past (provided by the sliding window) and of the pseudo-sinusoidal traffic pattern common in this kind of networks (acquired through an initial baselining activity, analyzing the network usage trend on a daily and/or weekly basis) we can estimate the network behavior in the near future and accordingly operate it in the most efficient way.

It is worth to underline that an edge cost function that mixes energy and load in one function is not a good option due to the different optimization objectives that are in contrast; therefore, we prefer a hybrid solution that however is forced to act *either* load-balancing or energy-efficiently, according to the what is currently most convenient, given the current configuration of the network. This also provides simplicity to the HyLERA algorithm, though its effectiveness and ease of use.

Note also that the main objectives of the HyLERA algorithm cannot be achieved by using a traditional Integer Linear Programming (ILP) approach. In fact, the HyLERA algorithm has been explicitly conceived to solve the dynamic online RWA problem, in which connection requests dynamically arrive one-by-one, i.e. without prior knowledge of the involved nodes and requested bandwidth. On the contrary, an ILP formulation can only be applied to the static RWA problem, in

which the whole set of connection requests (and hence the whole traffic matrix) is completely known in advance. Therefore, due to the computational complexity of both the above problems, only (small) instances of the static RWA problem can be solved with ILPs (usually employed for network planning), while a heuristic-driven algorithm such as HyLERA is the only feasible way for solving the dynamic RWA problem by ensuring the performances needed for online network operations.

4. The energy model

In order to handle the energy-awareness objectives, we need the ability to estimate for each candidate connection, in addition to the more traditional link or node properties that lead to the determination of the shortest and/or less congested paths, also the specific metrics that describe their energy consumption. A linear energy model has been used to describe the energy consumption of routers at different loads. According to real router measurements [25], we consider that, when a router is turned on but *idle*, it consumes half of its total power consumption and, as the load increases, its power consumption linearly increases, up to its maximum value which is reached when the router is fully loaded. The slope with which the power consumption function increases with respect to the load is given by its energy *scaling factor*, measured in W/Gbps in the energy scaling index (ESI).¹ A scaling factor of x W/Gbps means that x W are required to route 1 Gbps of traffic. Typical values for the scaling factor vary

¹ The energy consumption rate (ECR) and the energy scaling index (ESI) are equivalent metrics used to measure the efficiency of routers, where the former uses

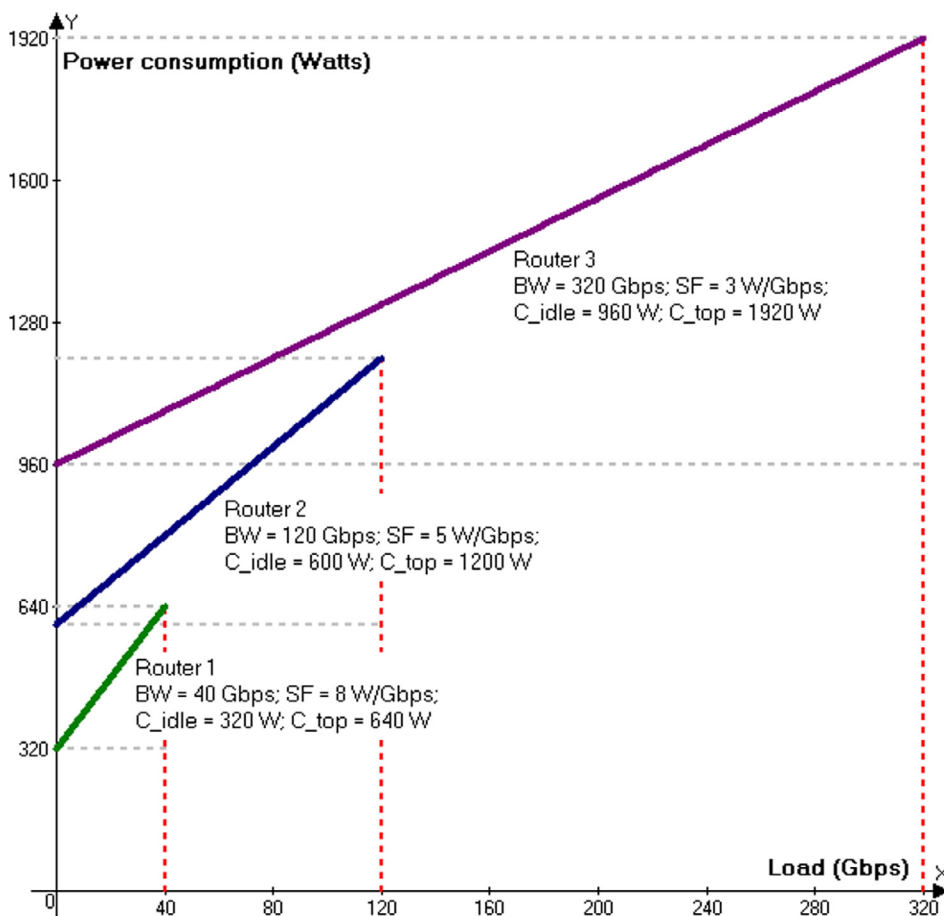


Fig. 2. The energy consumption of three routers with different aggregated bandwidth (BW) and scaling factors (SF) are represented.

from 1 to 10 W/Gbps [1], where small routers consume more energy per bit than larger routers, since the latter ones are typically more efficient and tend to be placed in the core of the network where the traffic is more aggregated [26].

In Fig. 2, the energy consumption of three different routers [1] has been drawn. Router 1 is a small router with a maximum capacity (aggregated bandwidth of all its interfaces) of 40 Gbps, and has a scaling factor of 8 W/Gbps; therefore, it consumes a maximum of 640 W (320 W just to stay ON, and 8×40 W when fully loaded). Similarly, Router 2 is a medium router (120 Gbps of aggregated bandwidth) with a scaling factor of 5 W/Gbps, and Router 3 is a big router (320 Gbps) with a scaling factor of 3 W/Gbps. It is worthwhile to note that the power consumption of idle routers is always present since we assume that no sleep mode is available at router level, i.e. an idle router cannot be put into sleep mode, as reported also in [13,27]. Therefore, any energy-related optimization relies on the variable power consumption of routers, selecting short routes passing through routers with low energy scaling factors.

Table 1 summarizes the parameters of the three routers and provides the equations describing their power consumption as a function of the load.

In general, the power consumption (y) of a router as a function of its load (x) is defined and takes value over the following domain and image:

$$y : [0, BW] \rightarrow [C_{idle}, C_{top}], \quad (1)$$

energy per bit (nJ/bit) and the latter uses Watts per Gbps (W/Gbps); in fact, it holds that $W/Gbps = (J/s)/(Gbit/s) = J/Gbit = nJ/bit$.

and it is given by the following equations:

$$y = SF \cdot x + C_{idle}, \quad (2)$$

where BW is the router aggregated bandwidth, C_{idle} and C_{top} are respectively the power consumption of the router when idle and when fully loaded and SF is its energy scaling factor. 3R regenerators have an energy scaling factor SF_{3R} of about 3 W/Gbps, as reported in [1].

5. The HyLERA cost function

The HyLERA algorithm is based on the Dijkstra shortest path algorithm, modified to operate in WDM networks and constrained on the availability of enough free bandwidth/wavelengths. The wavelength continuity constraint is imposed along a lightpath when no wavelength converters are present or no O-E-O conversion is performed.

The network is represented as a multigraph $G = (V, E)$, where V , $|V| = n$ is the set of nodes (either electronic routers or optical switches/cross-connects), E , $|E| = m$ is the set of edges modeling the links in the network. Note that, since WDM is deployed in the optical network, there can be more than one edge between a pair of nodes (therefore, G is a multigraph), each one representing a wavelength channel.

Each edge $(u, v) \in E$ in the network graph has nonnegative cost $c_{(u,v)} \in \mathbb{R}^+$ determined by the edge weighting function currently used by HyLERA (either load-balancing or energy-awareness, depending on the current network traffic state). The maximum capacity of each link $b_m(u, v) \in B$, where B is the set of possible bit-rates (310 Mbps, 622 Mbps, 2.5 Gbps, 10 Gbps, 40 Gbps, 100 Gbps), is given by the bit-rate of the interfaces operating on each wavelength; the current

Table 1

The values used for the three routers: aggregated bandwidth (BW), scaling factors (SF), power consumption when fully loaded (C_{top}) and in idle (C_{idle}), and finally the equations describing the power consumption (y) as a function of the load (x).

Router	Aggregated bandwidth (Gbps)	Scaling factor (W/Gbps)	Power consumption (W)	Power consumption (y) as a function of load (x) $y = f(x)$
Router 1	40	8	$C_{top} = 640$ W $C_{idle} = 320$	$y = 8x + 320$, $x \in [0, 40]$ $y \in [320, 640]$
Router 2	120	5	$C_{top} = 1200$ W $C_{idle} = 600$	$y = 5x + 600$, $x \in [0, 120]$ $y \in [600, 1200]$
Router 3	320	3	$C_{top} = 1920$ W $C_{idle} = 960$	$y = 3x + 960$, $x \in [0, 320]$ $y \in [960, 1920]$

available capacity $b_r(u, v)$ of a link is updated each time a connection is set up or torn down in the network.

Thus, network links may be weighted according to pure load-balancing or pure energy-awareness function, depending on the current traffic state of the network. In detail, when the network is unloaded, for example during the night, the links are weighted with the energy-aware function c_{EA} , which makes the Dijkstra algorithm select the lowest energy-consuming links:

$$c_{EA}(u, v) = SF(u) + SF(v) + SF_{3R}(u, v) \quad (3)$$

where $SF(i)$ is the energy scaling factor of node $i \in V$, and $SF_{3R}(u, v)$ is the scaling factor of 3R regeneration on edge (u, v) . Note that, since nodes in V can be either electronic routers equipped with energy-hungry switching fabrics or optical cross-connects, empowered by more energy-efficient matrices implemented by using MEMS or liquid-crystal on silicon (LCOS) devices, their energy scaling factors are different (we use the same values as [23]). Since no sleep mode is allowed for the optical amplifiers and they always operate on the entire band involved (e.g., the C-Band), they are not considered in the equation, being an invariant factor.

Alternatively, when the network is loaded, for example during peak day hours, the links are weighted according to the load-balancing function c_{LB} , which makes the Dijkstra algorithm prefer the less congested links in an effort to accommodate as much connections as possible.

$$c_{LB}(u, v) = [b_r(u, v) \cdot \log(b_m(u, v))]^{-1} \quad (4)$$

The more the available bandwidth and channel bandwidth are, the lower the load-balancing cost is (the maximum capacity is weighted with the logarithmic function, in order to lessen its importance with respect to the residual capacity). This is due to the assumption that the network revenue is directly proportional to the number of successfully routed connections (that is quite common for all the providers ensuring specific SLAs by relying on a traffic engineered circuit switching schema). Therefore, lowering the blocking probability as a consequence of a more balanced resource usage, achieved by always keeping enough residual capacity on each communication link, will increase the revenue for the operator. Note that the utilization of the nodes is considered only when pursuing load-balancing, since energy-efficiency is pursued only when the network traffic is light and therefore minimum blocking is experienced. It is worth to notice that we used a very simple naive load-balancing schema in order to control blocking probability also if more advanced and effective techniques are available in the literature, mostly based on distributing the load on the available communication resources in order to avoid the creation of bottlenecks over critical links. Thus, since the focus of this work is presenting an hybrid approach able to operate in a dynamic way, according to the network conditions, in order to clearly assess the effectiveness of the hybrid routing strategy-switching mechanism, two very simple routing strategies were employed, namely pure energy-efficiency and pure load-balancing.

In fact, it should be considered that the use of a more complex and sophisticated routing algorithm/strategy would have the effect of obfuscating the results, making more difficult to appreciate the effects of switching from a routing strategy to another one.

The *mode* in which the algorithm is working is identified by a parameter w which can be either *energy-awareness* or *load-balancing*. The switching between the former or the latter function is discriminated by a threshold evaluated in a sliding window. Specifically, the number ξ of connections requests that arrive in the network are monitored during the last k hours (sliding window) and, if they overcome the fixed threshold, the links cost function is changed. Despite a single threshold is used in each operating mode, in order to better bias the edge cost function to be applied at each time, two thresholds are defined (one for load-balancing and the other for energy-awareness): t_{high} and t_{low} . The values of these thresholds are highly dependent on the traffic pattern, the network topology and the load experienced in the different hours of the day. An initial traffic monitoring phase, aiming at acquiring information about the network baseline behavior, is needed in order to observe the traffic trend during low and high load periods of the day and properly tune the thresholds in such a way that the network is always operated in the desired mode. Starting from a void network (i.e., in the initial state), the algorithm works in energy-saving mode ($w = \text{energy-awareness}$), using the energy-aware function to weight the network links. As connections arrive at the network, they are continuously monitored and, if their number ξ exceeds the t_{high} value during the sliding window, the link cost function is switched to load-balancing ($w = \text{load-balancing}$) and stays until the connections arrival rate decreases to a value lower than the t_{low} threshold. When this occurs, the link cost function is switched back to energy-awareness to save energy during the low load period. Therefore, the link cost function is defined as:

$$c_{(u,v)} = \delta \cdot c_{EA} + (1 - \delta) \cdot c_{LB}, \quad (5)$$

where

$$\delta = \begin{cases} 1 & \text{if } \xi \leq t_{low} \text{ and } w = \text{load-balancing} \\ 0 & \text{if } \xi \geq t_{high} \text{ and } w = \text{energy-awareness} \end{cases} \quad (6)$$

5.1. OSPF-TE extensions

The hybrid behavior of HyLERA reaches its maximum effectiveness if all the routers in a network switch at the same time, i.e. in a coordinated, network-wide manner. Therefore, it is necessary to have a distributed mechanism to keep track of all the connections that arrive in the network during the sliding window, representing the instantaneous network load information needed to adaptively drive the switch-over between load-balancing and energy-aware behavior.

To this end, we extend the OSPF-TE Opaque Link State Advertisement (LSA) messages to include the arrival rate information during the sliding window, and spread this information with regular OSPF updates to all routers in the network.

Table 2
Sub-TLV for TE LSA.

Type	Length	Value
32768	4 octets	Connections arrival rate

In particular, we add a new Type-Length-Value (TLV) field to the Traffic Engineering extensions for OSPF-TE (Traffic Engineering LSA, opaque type=1) according to an agile implementation (the detailed explanation of each field of the TE LSA can be found in [28]). The *Type* field (16 bits) contains the ID of the newly defined entry (32768, which is the first available one) and the *Length* field (16-bits) specifies the extension of the *Value* field (in octets), which are contained in the payload of LSAs. As reported in Table 2 the Value field identifies the current connection arrival rate experienced by the involved device.

Note that the scaling factors (SF) used in the OSPF metric are statically defined at the network definition time, and it is not necessary to spread them periodically. However, it is also possible to configure OSPF to automatically spread them in case, for example, of changes in the topology; to this end, it is sufficient to add two other 4-octets sub-TLV fields, namely 32,769 and 32,770, containing respectively the node *id* and its SFs.

It is also worth to note that, even though OSPF is a link-state protocol (i.e., a flood only if a change), a link state refresh time (*LSRefreshTime*) is defined; when this time expires, a router floods a new LSA to all its neighbors, who will reset the age of the sending router's records to the new received age. OSPF sets the *LSRefreshTime* to 30 min [29], which is lower than the sliding window timespan (ranging from 1 to 6 h). Therefore, the periodic update done by the OSPF is directly usable to spread the number of connection requests arrived during the sliding window and the proposed TE LSAs will be flooded over the whole network on such a fixed time-basis, disseminating the network load information. The OSPF-TE convergence time on a typical carrier's transport network can be considered as negligible, being just few tens of milliseconds [30].

6. Time and space complexity analysis

The HyLERA algorithm is illustrated in Fig. 3. It takes as input the graph representing the network, the connection request with source and destination nodes and the required bandwidth, the high and low

Table 3
Parameters used in the simulation experiments.

Simulation parameters	Geant2 network
Simulated time	4 days
Time steps	2400 per day (one step every 36 s)
Number of generated connections	Varying from 0 to 263200
Connections lifetime	30 min
Connections bandwidth	310 Mbps
HyLERA sliding window k	1, 3, 6 h
HyLERA threshold t_{high}	Scenario 1: 9300; Scenario 2: 11,900; Scenario 3: 5400
HyLERA threshold t_{low}	Scenario 1: 8400; Scenario 2: 10,200; Scenario 3: 5200
Parameters Δ_{OA} , Δ_{BR}	80 km, 500 km
Distribution of nodes as source or destination of connections	Weighted according to the number of wavelengths of the routers (reported in Fig. 6)
RWA algorithms	HyLERA, pure load-balancing SPF, pure energy-aware SPF
Measurements	Blocked connections and power consumptions

thresholds and the connection requests count in the sliding window. The edge cost function (load-balancing/energy-awareness) initially is set to energy-awareness (since the network is void).

Lines 1–5 discriminate what is the edge cost function w that has to be used, given the current network traffic state, i.e., the number ξ of connections in the sliding window. This is performed by a simple check between ξ and the values of the thresholds and possibly the consequent assignment. These lines have a constant time complexity of $O(1)$.

The HyLERA algorithm is based on the Dijkstra's algorithm, modified to operate in WDM networks and constrained on the availability of enough free resources to serve the connections. The constrained-based routing is performed in line 6, checking, when a new node is discovered by the algorithm, if the available bandwidth on the link connecting that node is equal or greater than the connection request required bandwidth. The WDM network is represented as a multi-graph, in which there can be more than one edge between a pair of nodes, representing the different WDM channels. Therefore, when a new node at the lowest distance from the source is discovered, if there is enough free bandwidth, it is labeled not only with distance from the source and predecessor node, but also with the predecessor wavelength on which the node was reached. Both the bandwidth constraint check and the additional labeling have a constant

Algorithm 1 Hylera($G, c, t_{high}, t_{low}, \xi$)

Input: G : current network state $c = (s, d, b)$: connection request; s, d : source, destination nodes; b : required bandwidth t_{high} : ascending threshold to change from Energy-Awareness to Load-Balancing t_{low} : descending threshold to change from Load-Balancing to Energy-Awareness ξ : number of connection requests in the sliding window**Output:** G^* : new network state π^* : new lightpath

- 1: **if** $w = \text{energy-awareness}$ AND $\xi \geq t_{high}$ **then**
 - 2: $w \leftarrow \text{load-balancing}$
 - 3: **else if** $w = \text{load-balancing}$ AND $\xi \leq t_{low}$ **then**
 - 4: $w \leftarrow \text{energy-awareness}$
 - 5: **end if**
 - 6: $\pi^* \leftarrow \text{constrainedDijkstra}(G, c, w)$
 - 7: $G^* \leftarrow \text{Update the } w((u, v)_\lambda) \text{ costs of network edges } (u, v)_\lambda \text{ along the chosen path } \pi^*$
and increment ξ of one unit
 - 8: **return** (π^*, G^*)
-

Fig. 3. The HyLERA algorithm.

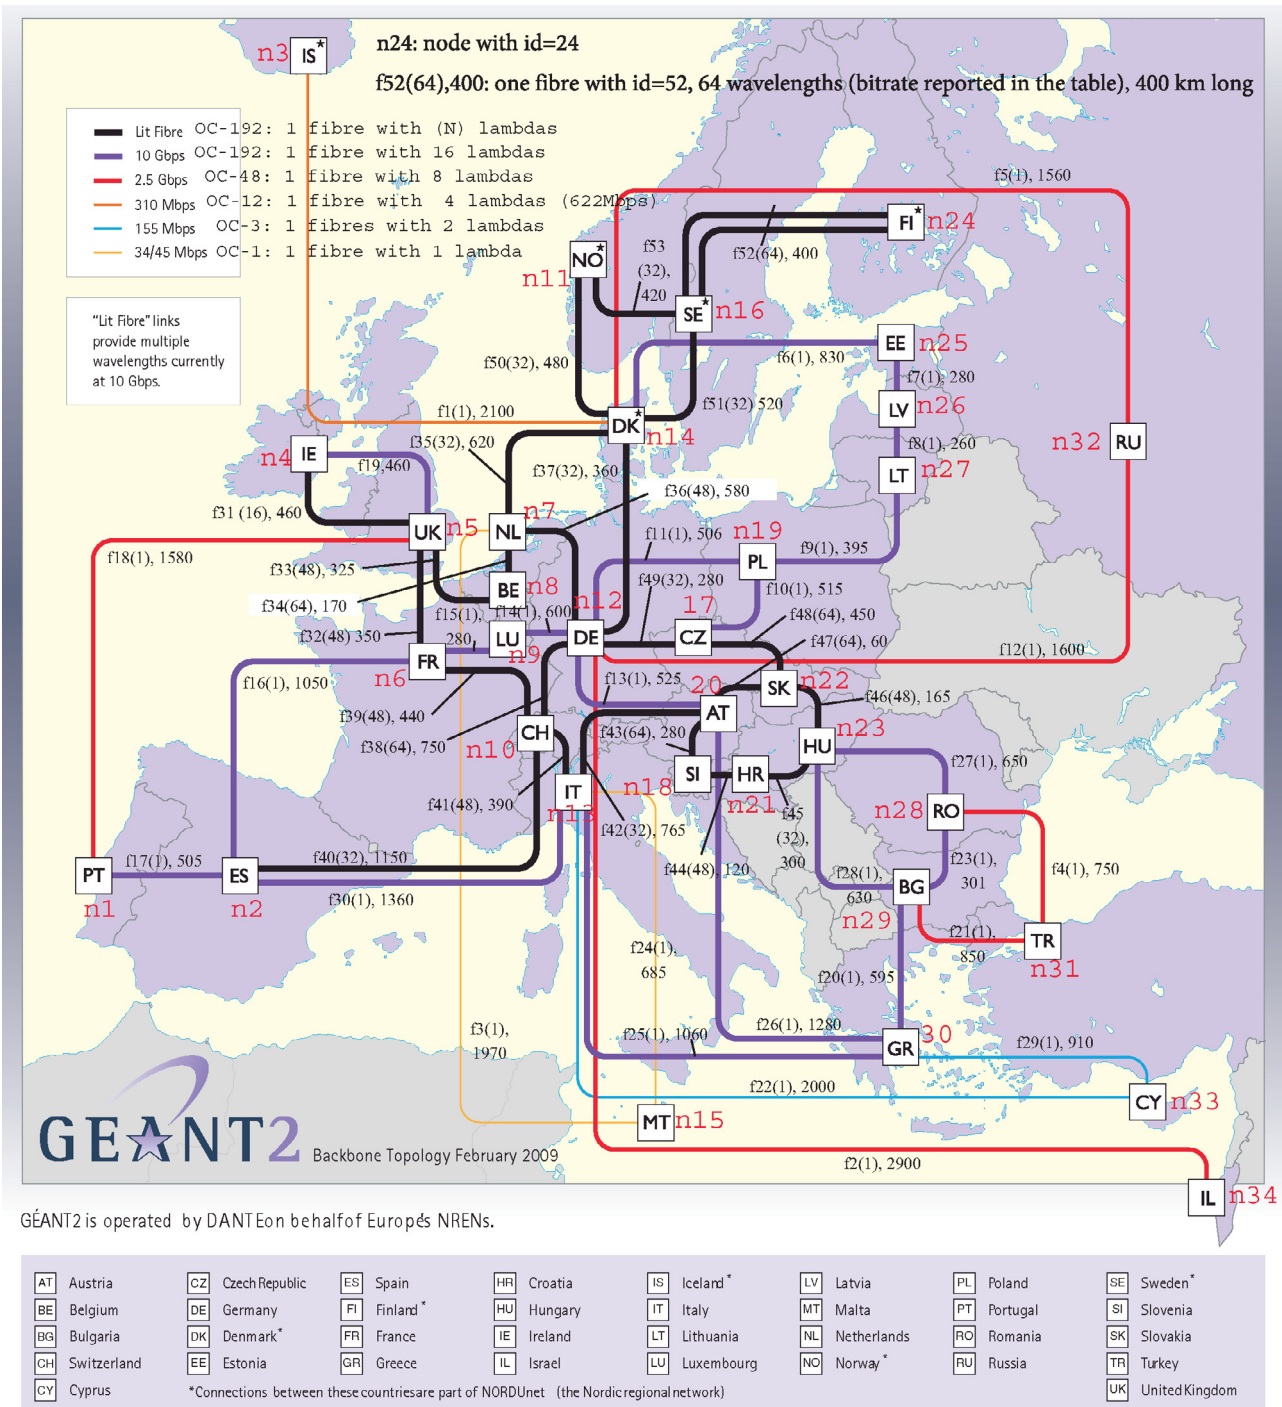


Fig. 4. The Geant2 network used in simulations.

time complexity of $O(1)$. The Dijkstra algorithm has a computational complexity of $O(m + n \log n)$, when improved by using a priority queue with a Fibonacci heap in the implementation [31].

In line 7, the new network state (i.e., the edges residual bandwidth and costs) is updated only for the network edges involved by the new path π^* , which, in a network with n nodes, has a maximum length of $n - 1$, therefore having a computational complexity of $O(n)$.

Finally, the new path π^* and the new network state G^* are returned in line 8.

Therefore, the HyLERA computational complexity in the worst case scenario is $O(m + n \log n)$ which is the same as the original Dijkstra algorithm.

As for space complexity, the multigraph network representation employed by HyLERA requires less space with respect to the layered graph approach conventionally used in dynamic RWA algorithms [23]. Using up to λ wavelengths on each edge, the layered representation with C converter nodes will require $\lambda n + 2$ nodes (λ layers, each dedicated to an individual wavelength, plus two additional nodes to serve as ingress and egress) and $\lambda m + 2\lambda + C \cdot (\lambda - 1)$ edges (converters can be modeled by cross-layer edges that connect each layer to the λ adjacent layer – a wavelength conversion spanning multiple frequencies will thus entail many such edges in sequence), while the equivalent multigraph representation will require only n nodes and λm edges, thus notably reducing the space complexity.




Fig. 5. The pseudo-sinusoidal traffic pattern.

Besides, in the layered graph, the ingress and egress nodes as well as the edges connecting them to the network have to be built each time a new connection arrives, while in the multigraph approach this preprocessing phase is not necessary thanks to its compact representation. Note that, even in absence of wavelength conversion, all the layers of the layered graph have to be explored, since the (first) wavelength of the lightpath may be any, which compensates the additional check needed in the multigraph approach to enforce the wavelength continuity constraint. Furthermore, the higher number of nodes and edges required by the layered graph with respect to the multigraph approach increases the time complexity which strictly depends on the n and m parameters.

7. Performance evaluation

Extensive simulation experiments have been performed in order to evaluate the efficiency of HyLERA under different operating conditions and scenarios. In order to resemble closest-to-reality scenarios, we made specific assumption on the traffic pattern, the distribution of connection requests and on the network topology and node design.

For this purpose, we modeled the well-known Geant2 pan-European research network, shown in Fig. 4. The network has 34 optical switches, each connected to an electrical router. The links between node pairs are modeled according to the real topology, with a WDM degree and interfaces bit-rates scaled up as shown in the figure. The size/class of the routing and switching equipment (and, thus, their energy scaling factor) is given by their aggregated bandwidth, i.e. the sum of the transport capacity of all their interfaces. Core nodes provide (sparse) wavelength conversion capability and a first-fit wavelength assignment algorithm (starting from the higher available wavelengths) is employed. The physical distance among nodes is reported in the figure and optical amplifiers (OA) and regenerators (3R) have been put accordingly on each link: an OA each 80 km and a 3R each 500 km. Simulations have been performed on an Intel® Core™ i7-950 CPU @ 3.07 GHz with 16 GB RAM and 64 bit operating system server equipped with the Sun® Java® Runtime Environment v.1.6. To perform the simulations, we used SimulNet [32], an ad-hoc optical network simulation environment that allows flexible modeling of network topologies as well as traffic load generation, data recording and post-processing. Simulation parameters are reported in Table 3.

We modeled the traffic pattern as pseudo-sinusoidal over the 24 h of a day, as shown in Fig. 5 and reported in [33].

In order to simulate real traffic matrices, the traffic pattern has not been uniformly distributed over the network nodes: bigger routers have been assigned higher probability to be selected as source or destination of connections requests than smaller ones. Such a probability is therefore linearly dependent on the aggregated number of wavelengths that a router manages, with the distribution probability shown in Fig. 6.

7.1. Simulations results

Due to space limitations, we only report the results associated to the sets of simulation experiments obtained in three *notable* scenarios, that show the flexibility of HyLERA in achieving different optimization objectives. In the first set of experiments (Scenario 1), HyLERA parameters (the high and low thresholds) have been set up in order to obtain the optimal behavior, given the current conditions, i.e. for the given network topology, traffic pattern and distribution, etc. Since the behavior of HyLERA is determined by the thresholds, we can force HyLERA to work by using one or the other cost function, either trying to save more energy or trying to block less connections. Therefore, the tests associated to the second and third scenarios have been performed by setting the parameters in order to make HyLERA behave in a *more* energy-aware or load-balanced way, respectively. In order to evaluate the performance of HyLERA, we compared its results with the ones associated to pure load-balancing and pure energy-aware algorithms, which are shortest path first (SPF) algorithms in which edges are labeled exclusively with the load-balancing (Eq. (3)) or energy-aware (Eq. (4)) cost functions, respectively. These algorithms provide the lower and upper bounds reference values of the achievable savings and blocking, and therefore the pure SPF algorithms are the same in the three sets of simulation experiments and are reported as a comparison. Note, however, that HyLERA is a hybrid solution which switches on-the-fly from a routing strategy to another one, when the network is already being operated. As a consequence, it takes some time to perform operating mode switching, that implies some progressive change/adaptation in its behavior (temporarily producing hybrid results) until it returns to operate according to only a single strategy characterizing the current operating mode. In absence of re-optimization, in fact, all the connections previously routed with a different strategy need to be terminated. This explains why HyLERA performance metrics are not always contained between the lower and upper bounds.

The window size k is a key parameter which depends on the network topology and traffic distribution. In our simulations, several time windows have been tested (1, 3 and 6 h), but only the results for $k = 3$ h are reported, since they give the best trade-off between the past traffic and the reactivity of the algorithm to changes. However, the best value for k has to be set experimentally, by also considering that a too short time window would increase the fluctuations in the change of the cost functions to use, and a too long window size would generate too much inertia and decrease the reactivity of the algorithm.

In order to show in a clear manner the results of the simulation experiments, the graphics contain extra axes distinguished by different colors. Blue and pink axes represent a change in the edge cost function used by HyLERA, being blue the change from energy-awareness to load-balancing, and pink the change from load-balancing to energy-awareness. White axes are used to split between simulation days: each simulation has been performed during 4 days (a sufficient time to reach an equilibrium between incoming and outgoing connections in the network, given the traffic pattern of Fig. 5).

The partial block and partial consumption represent the accumulated blocking and power consumption within the sliding window ($k = 3$ h); during the simulation, these values are calculated (every 36 s, as reported in Table 3) and shown in the result figures (one point every 6 min, i.e., every 10 update steps, to improve readability of the graphics).

7.2. Scenario 1: achieving the best-balance between energy-efficiency and load-balancing

The aim of this set of experiments is to show the flexibility of HyLERA to route as many connections as possible during the

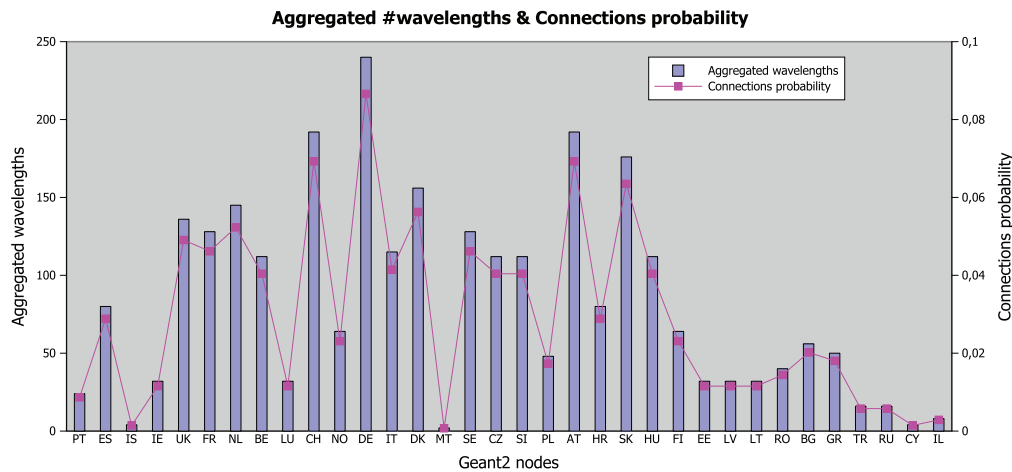


Fig. 6. The number of wavelengths of each network node is reported (columns) together with its probability to appear as either source or destination in connection request (square boxes).

highest congestion period (central 12 h of the day) and to save as much energy as possible during the remaining night hours.

Fig. 7 (a) shows the connections blocking probability measured during the sliding window (3 h) along the 4 days of simulation. In the first day of simulation, the three algorithms (HyLERA, the pure load-balancing SPF and the pure energy-awareness SPF) show an almost equal behavior, since the network starts empty and the only blocked connections are the ones that have to be established between a source or destination node with very few wavelengths (such as nodes 15 and 34 of Fig. 4) which are already used by previous connections that are still alive in the network (connections have a mean lifetime of 30 min).

In day 2, the network is not empty due to some connections that are still alive from day 1 and, when the load increases (at around 8 a.m.), the performance of the three algorithms begins to diverge. The pure energy-aware algorithm will select routes that pass through nodes with lower energy scaling factor (cf. Eq. (3)), which will result in longer, but more efficient, routes. However, selecting longer routes will also consume more resources, since the requested bandwidth is occupied on each link of the route; as a consequence, the connection blocking will significantly increase during central day hours.

On the other hand, the pure load-balancing algorithm keeps the connections blocking percentage at much lower rates, since it prefers the routes that pass through nodes with higher available and maximum bandwidths (cf. Eq. (4)). However, these routes are not the most energy-efficient ones, since the pure load-balancing SPF does not consider the nodes scaling factor in its decision process.

HyLERA, instead, changes dynamically the edge cost function according to the current network traffic state. It starts with energy-efficiency, since at the beginning of day 1 the network is empty. Then, the first change occurs at the middle of day 1, when the traffic increases during the peak hours. Starting from 4 p.m., the traffic decreases, and at 9 p.m. HyLERA switches back to the energy-saving modality. It is worth to note that, even if in day 1 the network starts empty and therefore there is no appreciable difference in the blocking connections, HyLERA is “working behind the scene”, saving energy already in day 1 (it will be evident when looking at the energy consumption shown in Fig. 7(c)). As for the blocking, the better behavior of HyLERA becomes evident starting from day 2 on. When the traffic load begins to increase (at approx. hour 35), HyLERA switches to the load-balancing cost function, keeping the blocking connections at values close to the pure load-balancing SPF, while the pure energy-efficient SPF increases notably the blocking. When the traffic decreases (at approx. hour 45), the edge cost function is changed to energy-efficiency to save energy during the night hours. During the


rest of the simulation (days 3 and 4), HyLERA keeps its blocking probability very low and close to the pure load-balancing SPF, avoiding the blocking peaks experienced by the pure energy-efficient SPF.

It is worth to note that, when switching to load-balancing mode, HyLERA does not present equal results than the pure load-balancing algorithm even if they are actually using the same cost function. This happens since the network has been operated using different routing schemes from the beginning of the simulation, thus creating completely different network states (paths of the connections and free bandwidth on the links) when the change in the HyLERA cost function occurs. The same holds for the energy-efficiency modality.


Fig. 7 (b) shows the total blocking during the simulation. The first day of the simulation presents a similar behavior to the partial block graph shown in Fig. 7(a). The difference between the algorithms can be clearly seen in the second day, where the network is fully loaded and the repetitive traffic pattern starts over. The pure load-balancing algorithm maintains a low blocking profile, presenting a very low variation in the values obtained: these are actually the lower bound values of the achievable blocking. The pure energy-aware algorithm notably increases its blocking rate every day until it reaches a stable value on day 4, which is substantially higher than the pure load-balancing one. HyLERA algorithm has a behavior similar to the pure load-balancing SPF. At the end of day 2 and for the rest of the days of the simulation, its blocking rates are always within a short range of values from the pure load-balancing ones.

The energy consumption of the three algorithms in the sliding window is reported in Fig. 7(c). Here we can see that HyLERA actually starts acting since the very first day of simulation as for the energy consumption, even if the results were not evident as for the blocking in day 1.

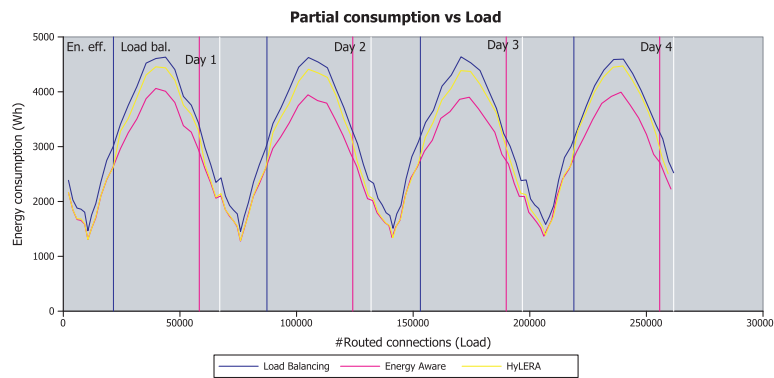
Since the first day, the energy consumption exhibits a pattern that is repeated until the end of the simulation. As HyLERA starts using the energy-aware function, its energy consumption values are almost identical to the pure energy-aware algorithm and, when the function change occurs, it can be clearly seen that the energy consumption increases up to values close to the ones of the pure load-balancing algorithm, but still lower. Note that, as before, HyLERA does not consume the same amount of energy as the pure load-balancing algorithm when it is working in load-balancing modality, since the network has been operated in a different way in the close past and is therefore in a different state. When the second function change occurs in day 1, HyLERA starts to behave as a pure energy-aware algorithm again, the energy consumption drastically decreases and the values obtained are very close to the pure energy-aware algorithm.



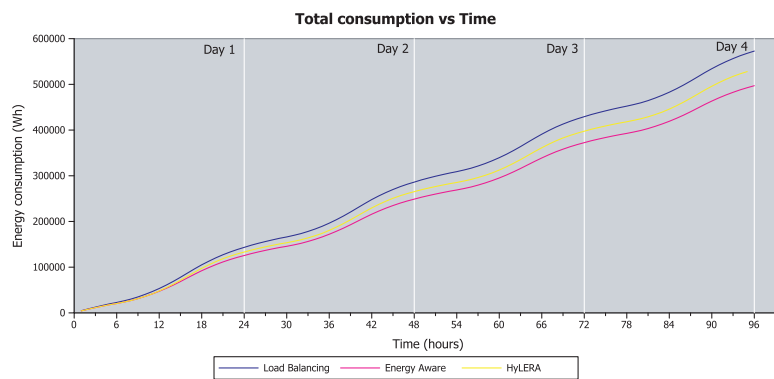
(a) Partial block vs Time



(b) Total blocking vs Time



(c) Partial energy consumption vs Load



(d) Total energy consumption vs Load

Fig. 7. Simulation results for Scenario 1. (For interpretation of the references to color in the text, the reader is referred to the web version of this article.)

Table 4
Final simulation values for the three set of simulations.

Algorithm	Total connections	Blocked connections	Blocking percentage	Consumed energy (kWh)
Load balancing SPF	263,200	1503	0.5710	572.85
Energy aware SPF	263,200	2671	1.0148	497.06
HyLERA Scenario 1	263,200	1702	0.6466	532.35
HyLERA Scenario 2	263,200	2510	0.9536	513.18
HyLERA Scenario 3	263,200	1652	0.6277	543.99

Finally, we show in Fig. 7(d) the total energy consumption of the network, taking into account all the connections since the beginning of the simulation. This graphic shows that the HyLERA energy consumption lays almost exactly in between the two pure SPF algorithms, since HyLERA uses each cost function during half a day.

Table 4 shows the final numerical results of the simulation. Comparing the final blocking of the network for the three different algorithms, we see that the pure energy-aware algorithm has almost twice the blocking percentage of the pure load-balancing algorithm, while HyLERA (Scenario 1) exhibits a slightly higher value than the pure load-balancing algorithm very far from the pure energy-aware SPF. About the energy consumption, HyLERA consumes around 40 kWh less (from the 572 kWh in total) than the pure load-balancing algorithm maintaining almost the same blocking percentage.

7.3. Scenario 2: achieving the highest energy-efficiency

The main objective of this set of experiments is to show the flexibility of HyLERA to operate in high energy saving mode, leaving only few peak hours in which HyLERA operates balancing the load of the network. Therefore, in order to obtain such a behavior, the load-balancing threshold (t_{high}) and the energy threshold (t_{low}) have been increased (to make it more difficult to switch to load-balancing and more easy to switch back to energy-efficiency).

Fig. 8 (a) shows the blocking probability of the network during the four simulation days (as observed in the sliding window). As in Scenario 1, during the first day no appreciable difference in the three algorithms is found. When the second day starts, it can be seen that HyLERA remains in energy-awareness and stays in that modality much longer than before, just switching back to load-balancing during few peak hours. As a consequence, HyLERA performs very close to the pure energy-aware algorithm since they are actually using the same cost function and the initial conditions of the network were the same (empty network). When the cost function changes to load-balancing, HyLERA continues to obtain values close to the pure energy-awareness algorithm for a while, reducing them little by little as the sliding window is being filled with connections routed with the load-balancing modality. This little reduction is due to the tiny contribution of the pure load-balancing cost function since it is only being used for 5 h every day.

Fig. 8 (b) shows the total blocking during the simulation. As it happened on the previous scenario, during the first day the three algorithms have almost the same blocking, being the difference negligible. When the second day starts, and the network gets more congested, since HyLERA is using most of the time the energy-aware function, its behavior is closer to the pure energy-aware algorithm, but the short range of hours where HyLERA uses the load-balancing function positively affects its blocking percentage, obtaining values that are always lower than the pure energy-aware SPF ones.

In Fig. 8(c), the evolution of the partial consumption is illustrated. As it happened in Scenario 1, the behavior is a daily pattern that is repeated since day 1. HyLERA starts using the energy-aware cost function, and that is why its consumption starts overlapped with it. It can be noted that, when the change to load-balancing occurs (blue line), an increase in the consumption is recorded, up to intermediate values between the two pure SPF algorithms. When HyLERA switches back

to energy-awareness, it quickly decreases its energy consumption to the same values of the pure energy-aware SPF algorithm.

Fig. 8 (d) shows the total energy consumption of the three algorithms. Focusing on HyLERA, it can be seen that its tendency has decreased compared with the first simulation, and its cost is closer to the pure energy-awareness algorithm.

In Table 4, we can observe that the blocking percentage of HyLERA (Scenario 2) is now substantially higher than in Scenario 1 (but still lower than pure energy-aware SPF), and the consumed energy has significantly decreased, saving 60 kWh (from the 572 kWh in total).

This configuration can be used by network operators with relative low traffic who want to save as much energy as possible to reduce their CAPEX expenses as much as possible, while preserving peak hours resource provisioning.

7.4. Scenario 3: achieving the best load-balancing

The simulation experiments in this scenario try to achieve the opposite behavior than the Scenario 2. In this case, the thresholds have been decreased in order to force HyLERA to use the load-balancing cost function during almost all the day time, and thus preserving the connectivity provisioning while not totally discarding energy-efficiency.

Fig. 9 (a) shows the blocking of the algorithms during the sliding window. As in the previous scenarios, the behavior of day 1 is the same for the three algorithms since the network is still being loaded with connections. Here we can see that the change to load-balancing closely follows the change to energy-efficiency, achieving the required more load-balanced behavior. During the 4 days of simulation, HyLERA maintains a performance very close to the pure load-balancing SPF algorithm, slightly changing toward the energy-aware SPF during few night hours.


Fig. 9 (b) shows the total blocking of the three algorithms, in which it can be seen how HyLERA performs very well in terms of blocking probability.

Fig. 9 (c) shows the energy consumption of the algorithms during the sliding window. As HyLERA is behaving almost all the day time as the pure load-balancing algorithm, their consumptions are very close. Only when HyLERA uses the energy-aware cost function it can be seen how its consumption decreases, but for the rest of the time its values are the closest ones with the pure load-balancing algorithm that can be observed among all three simulation scenarios.

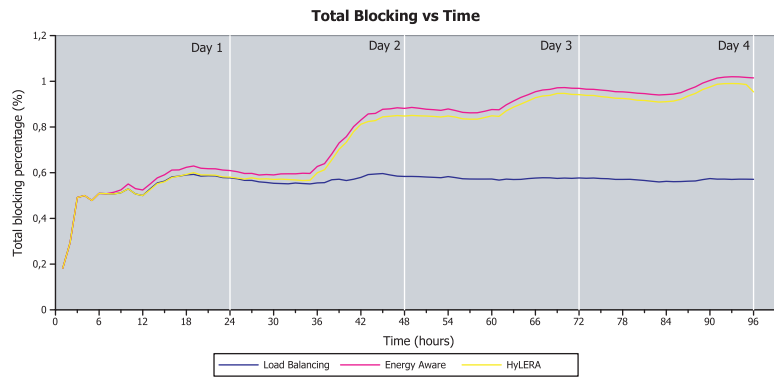
Fig. 9 (d) shows the total energy consumption. As expected, reducing the blocking percentage makes HyLERA consume more energy, even if its energy consumption stays lower than the pure load-balancing SPF algorithm.

In Table 4, we can observe that the total energy consumption of HyLERA (Scenario 3) is higher than the pure energy-aware SPF but still lower than the pure load-balancing algorithm, saving energy (28 kWh from the total of 572 kWh) while maintaining a total blocking very close to the pure load-balancing algorithm.


Considering an average energy cost of 0.12 € /kWh [34][35], HyLERA saves 443 € /year in Scenario 1 (7% savings with respect to pure LB over a year), 653 € in Scenario 2 (10% savings) and 316 € in Scenario 3 (5% savings). Note that the economic benefits of HyLERA have been calculated within the considered simulation environment, which is an IP/WDM optical network in which every node has its




(a) Partial block vs Time



(b) Total blocking vs Time




(c) Partial energy consumption vs Load




(d) Total energy consumption vs Load

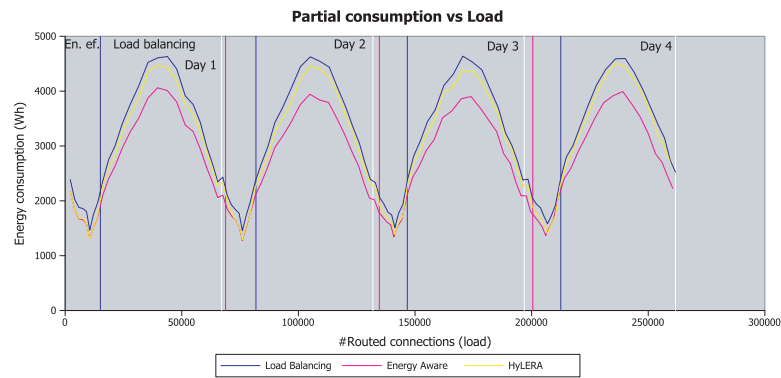
Fig. 8. Simulation results for Scenario 2. (For interpretation of the references to color in the text, the reader is referred to the web version of this article.)




(a) Partial block vs Time



(b) Total blocking vs Time



(c) Partial energy consumption vs Load



(d) Total energy consumption vs Load

Fig. 9. Simulation results for Scenario 3. (For interpretation of the references to color in the text, the reader is referred to the web version of this article.)

counterpart in the optical domain (i.e. a full optical layer is present and electronic processing only occurs at network edges). This can be considered as the best-case scenario in terms of energy demand, since the use of optical technology introduces significant energy savings. Nevertheless, the HyLERA algorithm can be seamlessly applied to electronic networks in which there are no optical nodes or the optical elements are limited to the core segment. Since the power consumption of electronic devices is as much as 100 times higher than the optical components' [1], the economic benefits of operating HyLERA in such networks easily rises to many thousands of Euros per year.

As it has been shown, the relation between blocked connections and energy consumption is a trade-off, and if we want to reduce one of them, we have to necessarily increase the other, but HyLERA achieves a substantial energy reduction at the expense of a very limited increase in the blocking percentage.

8. Conclusions and future work

In this paper we presented HyLERA, a simple though effective RWA algorithm for WDM networks that dynamically adapts its behavior according to the current traffic state of the network in order to achieve more effective service provisioning during peak traffic hours and higher energy efficiency during low load periods. HyLERA is based on an ad-hoc version of the Dijkstra algorithm modified in order to work with WDM network and in presence of a QoS requirement on the provisioned bandwidth. Simulation results show that HyLERA achieves very low connection blocking probability during peak hours and very low energy consumption during offload periods. It has also been showed that HyLERA can be configured to work by privileging the service provisioning or the energy efficiency objective, depending on the network operator's policy. The very low computational complexity of HyLERA and its simplicity make it a valuable candidate for considering the energy efficiency in the modern network infrastructures without affecting (or controlling) the service provisioning level.

Acknowledgments

This work was partially supported by the *Ministerio de Economía y Competitividad* of the Spanish Government under the SUNSET project TEC2014-59583-C2-2-R and the *Departament d'Universitats, Recerca i Societat de la Informació de la Generalitat de Catalunya* under the contracts 2014SGR-1427 and 2015FI_B00931.

References

- [1] BONE project, WP 21 TP green optical networks, D21.2b report on Y1 and updated plan for activities, 2009.
- [2] H. Zang, J.P. Jue, B. Mukherjee, et al., A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks, *Opt. Netw. Mag.* 1 (1) (2000) 47–60.
- [3] R. Ramaswami, K.N. Sivarajan, Routing and wavelength assignment in all-optical networks, *IEEE/ACM Trans. Netw. (TON)* 3 (5) (1995) 489–500.
- [4] F. Palmieri, U. Fiore, S. Ricciardi, Constrained minimum lightpath affinity routing in multi-layer optical transport networks, *J. High Speed Netw.* 17 (4) (2010) 185–205.
- [5] S.D. Rao, C. Murthy, Distributed dynamic qos-aware routing in WDM optical networks, *Comput. Netw.* 48 (4) (2005) 585–604.
- [6] D. Zhemine, M. Hamdi, J.Y. Lee, V.O. Li, Integrated routing and grooming in GMPLS-based optical networks, in: 2004 IEEE International Conference on Communications, 3, IEEE, 2004, pp. 1584–1588.
- [7] B. Chatterjee, N. Sarma, P. Sahu, Priority based routing and wavelength assignment with traffic grooming for optical networks, *IEEE/OSA J. Opt. Commun. Netw.* 4 (6) (2012) 480–489, doi:10.1364/JOCN.4.000480.
- [8] F. Palmieri, U. Fiore, S. Ricciardi, SPARK: a smart parametric online RWA algorithm, *J. Commun. Netw.* 9 (4) (2007) 368–376.
- [9] F. Palmieri, U. Fiore, S. Ricciardi, A minimum cut interference-based integrated RWA algorithm for multi-constrained optical transport networks, *J. Netw. Syst. Manage.* 16 (4) (2008) 421–448.
- [10] K. Manousakis, K. Christodouloupoulos, E. Kamitsas, I. Tomkos, E. Varvarigos, Offline impairment-aware routing and wavelength assignment algorithms in translucent WDM optical networks, *J. Lightwave Technol.* 27 (12) (2009) 1866–1877, doi:10.1109/JLT.2009.2021534.
- [11] S. Ricciardi, F. Palmieri, U. Fiore, D. Careglio, G. Santos-Boada, J. Solé-Pareta, Chapter 19 – Towards energy-oriented telecommunication networks, in: M.S. Obaidat, A. Anpalagan, I. Woungang (Eds.), *Handbook of Green Information and Communication Systems*, Academic Press, 2013, pp. 491–512. <http://dx.doi.org/10.1016/B978-0-12-415844-3.00019-X>.
- [12] S. Ricciardi, D. Careglio, F. Palmieri, U. Fiore, G. Santos-Boada, J. Solé-Pareta, Energy-aware RWA for WDM networks with dual power sources, in: 2011 IEEE International Conference on Communications (ICC), 2011a, pp. 1–6.
- [13] S. Ricciardi, D. Careglio, G. Santos-Boada, J. Solé-Pareta, U. Fiore, F. Palmieri, Towards an energy-aware internet: modeling a cross-layer optimization approach, *Telecommun. Syst.* (2011b) 1–22.
- [14] S. Ricciardi, D. Careglio, U. Fiore, F. Palmieri, G. Santos-Boada, J. Solé-Pareta, Analyzing local strategies for energy-efficient networking, in: V. Casares-Giner, P. Manzoni, A. Pont (Eds.), *NETWORKING 2011 Workshops, Lecture Notes in Computer Science*, 6827, Springer, Berlin/Heidelberg, 2011c, pp. 291–300.
- [15] S. Ricciardi, G. Santos-Boada, M. Klinkowski, D. Careglio, F. Palmieri, Towards service orchestration between smart grids and telecom networks, in: J.-M. Pierson, G. Da Costa, L. Dittmann (Eds.), *Energy Efficiency in Large Scale Distributed Systems, Lecture Notes in Computer Science*, Springer, Berlin/Heidelberg, 2013, pp. 300–309, doi:10.1007/978-3-642-40517-4_25.
- [16] M. Gattulli, M. Tornatore, R. Fiandra, A. Pattavina, Low-carbon routing algorithms for cloud computing services in IP-over-WDM networks, in: 2012 IEEE International Conference on Communications (ICC), 2012, pp. 2999–3003, doi:10.1109/ICC.2012.6364347.
- [17] D. Arai, K. Yoshihara, Eco-friendly distributed routing protocol for reducing network energy consumption, in: 2010 International Conference on Network and Service Management (CNSM), 2010, pp. 104–111, doi:10.1109/CNSM.2010.5691325.
- [18] S. Ricciardi, J. Wang, F. Palmieri, D. Careglio, A. Manolova, G. Santos-Boada, Eco-sustainable routing in optical networks, *Photonic Netw. Commun.* 26 (2–3) (2013) 140–149.
- [19] A. Jirattigalachote, C. Cavdar, P. Monti, L. Wosinska, A. Tzanakaki, Dynamic provisioning strategies for energy efficient (WDM) networks with dedicated path protection, *Opt. Switch. Netw.* 8 (3) (2011) 201–213. <http://dx.doi.org/10.1016/j.osn.2011.03.008>. Special Issue on Green Communications and Networking
- [20] A. Udalcovs, P. Monti, V. Bobrovs, R. Schatz, L. Wosinska, Power efficiency of (WDM) networks using various modulation formats with spectral efficiency limited by linear crosstalk, *Opt. Commun.* 318 (2014) 31–36. <http://dx.doi.org/10.1016/j.optcom.2013.12.033>.
- [21] U. Mandal, M.F. Habib, S. Zhang, C. Lange, A. Gladisch, B. Mukherjee, Adopting hybrid CDN-P2P in IP-over-WDM networks: An energy-efficiency perspective, *J. Opt. Commun. Netw.* 6 (3) (2014) 303–314, doi:10.1364/JOCN.6.000303.
- [22] A. Nag, T. Wang, B. Mukherjee, Robust design of spectrum-efficient green optical backbone networks, *J. Lightwave Technol.* 31 (7) (2013) 1138–1144, doi:10.1109/JLT.2013.2245301.
- [23] S. Ricciardi, F. Palmieri, U. Fiore, D. Careglio, G. Santos-Boada, J. Solé-Pareta, An energy-aware dynamic (RWA) framework for next-generation wavelength-routed networks, *Comput. Netw.* 56 (10) (2012) 2420–2442. <http://dx.doi.org/10.1016/j.comnet.2012.03.016>.
- [24] P. Wiart, P. Monti, L. Wosinska, Power savings versus network performance in dynamically provisioned WDM networks, *IEEE Commun. Mag.* 50 (5) (2012) 48–55, doi:10.1109/MCOM.2012.6194382.
- [25] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsang, S. Wright, Power awareness in network design and routing, in: INFOCOM 2008, The 27th Conference on Computer Communications. IEEE, 2008, pp. 457–465, doi:10.1109/INFOCOM.2008.93.
- [26] W. Vereecken, W. Van Heddeghem, D. Colle, M. Pickavet, P. Demeester, Overall ICT footprint and green communication technologies, in: 2010 4th International Symposium on Communications, Control and Signal Processing (ISCCSP), 2010, pp. 1–6, doi:10.1109/ISCCSP.2010.5463327.
- [27] A. Muhammad, P. Monti, I. Cerutti, L. Wosinska, P. Castoldi, A. Tzanakaki, Energy-efficient WDM network planning with dedicated protection resources in sleep mode, in: 2010 IEEE Global Telecommunications Conference. GLOBECOM 2010, 2010, pp. 1–5, doi:10.1109/GLOBECOM.2010.5683205.
- [28] J. Moy, OSPF version 2, Technical Report, IETF RFC 2178, 1997.
- [29] J. Doyle, J.X. Carroll, *Routing TCP/IP*, vol. 1, 2nd ed., Cisco Press, 2005.
- [30] J. Wang, S. Ricciardi, A. Manolova, S. Ruepp, D. Careglio, L. Dittmann, OSPF-TE extensions for green routing in optical networks, in: 2012 17th Opto-Electronics and Communications Conference (OECC), 2012, pp. 411–412, doi:10.1109/OECC.2012.6276497.
- [31] R. Ramamurthy, B. Mukherjee, Fixed-alternate routing and wavelength conversion in wavelength-routed optical networks, *IEEE/ACM Transactions on Netw.* 10 (3) (2002) 351–367, doi:10.1109/TNET.2002.1012367.
- [32] F. Palmieri, U. Fiore, S. Ricciardi, Simulnet: a wavelength-routed optical network simulation framework, in: IEEE Symposium on Computers and Communications, 2009. ISCC 2009, 2009, pp. 281–286, doi:10.1109/ISCC.2009.5202259.
- [33] L. Chiaraviglio, M. Mellia, F. Neri, Energy-aware backbone networks: a case study, in: IEEE International Conference on Communications Workshops, 2009. ICC Workshops 2009, 2009, pp. 1–5, doi:10.1109/ICCW.2009.5208038.
- [34] Europe electricity prices: http://ec.europa.eu/eurostat/statistics-explained/index.php/Energy_price_statistics, 2015.
- [35] United states energy information administration, state electricity prices, <http://www.eia.doe.gov/ncic/rankings/stateelectricityprice.htm>, 2014.