This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

A Multiobjective Wavelength Routing Approach Combining Network and Traffic Engineering With Energy Awareness

Francesco Palmieri, Sergio Ricciardi, Ugo Fiore, Aniello Castiglione, *Member, IEEE*, and Davide Careglio

***Abstract*—In modern optical networks, infrastructure management is faced with the challenge of using expensive equipment and communication resources as efficiently as possible. This now includes keeping power consumption costs at a minimum and using the available optical links in a balanced way, in addition to the traditional goals of providing the best possible performance to the end customers while meeting their quality requirements. Accordingly, this paper presents a heuristic single-step lightpath routing and wavelength assignment algorithm, handling online dynamic connection requests within a fully distributed network control plane. By using shortest path routing, the presented scheme determines the best compromise solution between the users' and carrier's objectives. The former can be mainly expressed in terms of connection QoS requirements, while the latter comprises network engineering (distributing the load in order to achieve near-optimum resource usage) and containing energy consumption. This approach is able to find, in a polynomial computing time, a multiobjective optimization solution that maximizes the carriers' return of investment and supports a high number of users' request while drastically reducing the network operational expenditures, as extensively demonstrated through a simulation.**

***Index Terms*—Algorithms, computer networks, energy efficiency, routing, telecommunications.**

I. Introduction

LARGE-SCALE wide-area transport networks are a strategic component of today's global communication infrastructure. Wavelength division multiplexing (WDM) and wavelength routing are among the best available technological options enabling these networks to offer highly scalable transmission capacity, protocol transparency, and simplified management, satisfying at the same time the growing demand on energy efficiency. In such networks, two adjacent nodes are connected by one or multiple fibers, each carrying multiple wavelengths or channels. Each node consists of a dynamically configurable optical switch that supports fiber switching and wavelength switching, i.e., the data on a specified input fiber and wavelength can be switched to a specified output fiber on the same wavelength. In order to transfer data between generic source–destination node pairs, a dedicated optical channel or *lightpath* has to be established by allocating an available wavelength throughout the entire route of the transmitted data. The problem of finding suitable paths and optical channels for a set of connections is known as routing and wavelength assignment (RWA) problem, which is known to be NP-complete [1]. Allocation takes place according to a circuit-switched model where an end-to-end connection is assigned the same wavelength resource for its entire duration. Benefiting from optical amplifiers and transparent optical switches, lightpaths can span more than one fiber link and remain entirely optical from source to destination, limiting the use of expensive energy-hungry signal regeneration equipment and other intermediate devices converting optical signals into the electronic domain and back.

Network providers aim at using their expensive connections and equipment as efficiently as possible to maximize their medium and long-term revenues. This encompasses three fundamental tasks. The first one, referred to as *network engineering*, corresponds to designing the network to achieve a correct dimensioning of communication resources and switching equipment, keeping them continuously up-to-date with respect to the state-of-the-art technologies and the expected growth trends. Since the network should be overprovisioned—to sustain traffic peaks—providers also strive to take the most from their investments by distributing the traffic load on all of the available resources in order to use them in a more balanced way, maximize the average capacity available to customers, and avoid the creation of bottlenecks.

The second task, commonly known as *energy awareness* (EA), entails dynamic power management throughout the network. The objective is to decrease power consumption as well as to reduce energy costs by exploiting the energy-proportional features of new-generation network equipment (adapting power consumption to traffic load), privileging paths through elements powered by renewable (and possibly cheaper) energy sources, or taking advantage from time- or location-dependent fluctuations in electricity costs. Achieving energy efficiency in network elements, usually through energy-proportional behavior, is an objective strongly pursued by the most recent standardization efforts, and technologies are available off-the-shelf [2]. The latter task, often referred to as *traffic engineering*, involves optimizing the routing control logic responsible for

Manuscript received October 21, 2014; revised February 20, 2015; accepted May 1, 2015. This work was supported in part by the Ministerio de Economía y Competitividad of the Spanish Government (SUNSET project TEC2014-59583-C2-2-R) and the Departament d'Universitats, Recerca i Societat de la Informació de la Generalitat de Catalunya (2014SGR-1427).

F. Palmieri and A. Castiglione are with the Department of Computer Science, University of Salerno, 84084 Fisciano, Italy (e-mail: fpalmieri@unisa.it; castiglione@ieee.org).

S. Ricciardi and D. Careglio are with the Department of Computer Architecture, Technical University of Catalonia, 08034 Barcelona, Spain (e-mail: sergio.ricciardi@ac.upc.edu; careglio@ac.upc.edu).

U. Fiore is with the Information Services Center, University of Naples Federico II, 80138 Naples, Italy (e-mail: ufiore@unina.it).

Digital Object Identifier 10.1109/JSYST.2015.2433671

dynamically routing traffic flows in the network according to their performance requirements.

In modern carrier networks, a smart integrated network engineering approach can be developed to combine the aforementioned three tasks into a common multiobjective optimization framework, aimed at harmonizing the apparently disjoint (or, worse, contrasting) energy and traffic-related objectives. In particular, according to our vision, energy efficiency becomes a first-class objective in the wavelength routing scenario, together with the more traditional network-related ones.

Since the optimal selection of paths satisfying multiple independent requirements, objectives, and constraints is a computationally intractable problem, incompatible with a dynamic online routing environment, in this paper, we propose a heuristic approach to effectively find feasible paths, leading to suboptimal solutions, in a bounded time. For this purpose, we developed a single-step shortest path routing-based dynamic RWA scheme, using the Dijkstra algorithm, suitably modified to properly work into the wavelength switching environment, with a combined network/traffic engineering and energy-aware cost functions whose main goal is finding the best compromise solution among the three aforementioned carriers' and users' objectives in an extremely simple and performance-effective way. It performs dynamic online constrained shortest-path-first selections by considering multiple weighting factors such as optical link transmission properties (e.g., wavelengths per link, distance/delay, total link capacity, etc.) and impairments (e.g., bit error rate and intermediate amplification steps) as well as energy-related ones (e.g., power consumption of intermediate elements, such as routers, switches, and optical amplifiers). The proposed framework can accommodate for particular operating conditions, hardware characteristics, or carrier necessities. Other parameters related to energy consumption can be added and integrated, and the choice of the objective weighting function may be modified to tailor specific needs. The presented RWA scheme relies on the presence of an underlying fully distributed network control plane, implying cooperation between the nodes concurring to the RWA problem solution and providing a link-state advertisement protocol to synchronize the nodes' network views by conveying all of the link status information (including energy-related ones) to every participant, as well as a signaling mechanism to be used for the reservation and establishment of paths (e.g., [3]). The presence of such a common control plane guarantees the convergence of all of the independent decisions taken by the nodes that behave according to a fully decentralized scheme but share the same network topology and status view, resulting in identical routing plans when running the same shortest-path-first algorithm at the same time.

II. Related Work

Several RWA approaches available in literature focused their attention on the use of shortest path routing mechanisms, to keep the implementation simple and limit the overall problem computational complexity. For example, the work in [4] discusses some structural properties of unsplittable shortest path routing by developing several ILP models to find lengths that induce a prescribed set of shortest paths. In addition, in [5], it has been shown that the problem of finding simple routing weights that uniquely induce a prescribed set of shortest paths is computationally hard. In [6], the authors analyzed the signaling mechanisms supporting fully adaptive shortest path routing in wavelength-routed networks. The proposal in [7] used the Bellman–Ford algorithm for shortest path routing on each individual wavelength plane, which, however, suffers from several well-known convergence problems typical of distance vector routing mechanisms. In the works presented in [8]–[10], the Dijkstra's shortest path algorithm has been used to determine the shortest lightpath or semilightpath in a polynomial time; such an algorithm employs a uniform search strategy which implies that, in large networks, it may unnecessarily visit many nodes and take much time before the shortest path is identified. Following the seminal work of [11], which first envisioned the idea of energy conservation in Internet-based infrastructures, the concept of EA has been introduced in shortest path routing by [12] and [13]; in order to contain energy consumption, a Dijkstra scheme modified to reduce the number of links and share the best paths under light loads has been proposed. Another approach managing EA at the subwavelength level has been presented in [14], where a modular physical architecture is discussed for the optical multiplexers and the advantages of EA in individual components of such an architecture are evaluated. On the other hand, some greedy heuristics to contain network power consumption, based on ranking nodes and links with respect to the amount of traffic that they would carry in the context of an energy-agnostic configuration, have been proposed in [15]. A two-stage approach, based on performing energy-aware selection on the K-shortest paths determined according to traditional network engineering criteria, has been presented in [16]. The great advantage of this work on the aforementioned ones is that it handles an inherently multiobjective optimization problem into a unique two-stage routing scheme. Nonetheless, its two-stage selection process (originally developed in [17]), whose complexity linearly depends on the number of K alternative routes considered, can introduce additional effort in the computation of the "best" path. In this sense, the routing scheme here proposed has the advantage of being a single-stage Dijkstra-based algorithm, whose computational complexity is polynomial in time and does not depend on any of the parameters of the algorithm. Finally, a single-stage shortest path based scheme has also been proposed in [18], mainly targeted, however, at the seamless introduction of energy-related information into the standard GMPLS framework, with minimal modification efforts in the control plane protocols (e.g., OSPF with traffic engineering extensions), and providing backward compatibility with all of the existing implementations.

While the aforementioned contribution achieves good results in terms of energy demand and greenhouse gases emission containment, the scheme proposed in this work, being based on a more sophisticated paradigm for combining and weighting multiple optimization objectives (QoS/SLA, network management, and energy-related ones), is able to achieve significantly better tradeoffs between these objectives and hence more satisfactory results for all of the involved actors (i.e., end users and network providers).

III. RWA Framework

An effective RWA solution applies methodological and technological considerations and principles to the modeling, characterization, control, and performance optimization of the network behavior. Optimization in this context refers to a more complex joint criterion combining the traditional traffic/network engineering goals with the new energy-related ones in a common multiobjective optimization framework, searching from the best compromise between the aforementioned goals. We structured our solution according to an adaptive shortest path routing scheme, which dynamically selects the minimum cost path between each pair of source and destination nodes based on the up-to-date global network status and specifically on the individual costs assigned to the underlying communication links. The motivations beyond this choice arise from the evidence that shortest path routing is one of the most commonly used strategies in wavelength-routed optical networks, since, while being really easy to implement and effective, it is known to be asymptotically cost-optimal in heavily loaded networks and asymptotically near-optimal in large sparse networks supporting any-to-any communication [19].

A. Network Model

In the proposed scheme, the network topology can be modeled as a graph $G(N, E)$, where N denotes the set of vertices (the wavelength switching devices) and E denotes the set of edges (the optical fiber links). A weighting function $w_t(u, v)$, whose value is dynamically determined at each time t, is associated to each link $(u, v) \in E$, representing the cost of using the link at the time t. Every edge (u, v) represents an optical connection between two sites $u, v \in N$, composed by $f_{(u,v)}$ independent fibers, each one with the same length $l_{(u,v)}$, delay $d_{(u,v)}$, bit error rate $e_{(u,v)}$, number of intermediate optical amplification stages $a_{(u,v)}$, and/or number of regeneration stages $r_{(u,v)}$. All of the $f_{(u,v)}$ fibers associated to an edge (u, v) support the same number $\lambda^c_{(u,v)}$ of wavelength channels, with maximum nominal bandwidth $b_{(u,v)}$, where $\lambda^{a(t)}_{(u,v)}$ of them are available at the time t. We also model a set of service requests $R \subseteq N^2$, where each request $r = (s, d) \in R$ is characterized by its time of arrival t and by a set Q_r of QoS constraints comprising the minimum requested bandwidth b_r, the maximum acceptable BER e_r or delay d_r so that $Q_r = \{b_r, e_r, d_r, \ldots\}$. When a new service request $r = (s, d)$ arrives, a dedicated lightpath $p_\lambda = \{(s, x_1), \ldots, (x_n, d)\}$ defined as a unique wavelength channel λ, carved onto a set of optical links physically connecting the source and destination nodes s, d through the intermediate switching devices $\{x_1, \ldots, x_n\}$, should be created. For each optical link $(u, v) \in p_\lambda$, it must hold that $b_{(u,v)} \geq b_r \wedge e_{(u,v)} \leq e_r \wedge d_{(u,v)} \leq d_r$. Given the set R of service requests, our main goal is to allocate the optical wavelength channels so that the maximum number of requests can be simultaneously satisfied, exhausting the capacity of the minimum number of fibers and reducing as much as possible both the energy and network-related costs. At first, in the so-called *routing* phase, the optical path for a generic request r is determined by using a traditional constrained shortest path routing algorithm such as Dijkstra's, working at each time t of invocation on the graph G in which the edges in E are dynamically assigned a cost value given by the link weighting function $w_t(u, v)$. The resulting path is the minimum cost path satisfying the Q_r requirements, according to the aforementioned costs calculated at the time of invocation t, considering that the cumulative cost associated with a lightpath p_λ on the wavelength λ is defined as the sum of its constituent link $\{(s, x_1), \ldots, (x_n, d)\}$ costs. In other words, the shortest path routing algorithm is constrained by operating on a graph whose links or nodes are restricted to those ones in the original topology that satisfy the Q_r requirements. Successively, in the *wavelength assignment* phase, a wavelength reservation request is propagated to all of the intermediate devices along the path, to provisionally reserve and then allocate a dedicated wavelength on each involved fiber link. In a pure wavelength routing environment where conversion capability either in the optical or in the electric domain is not provided, the same wavelength must be used on all links, by enforcing the *continuity constraint*. Such a unique wavelength is selected by using a traditional *best-fit* scheme between the available ones. Clearly, the choice of avoiding wavelength conversion on intermediate nodes can adversely affect the blocking probability experienced by connection requests. However, it avoids long delays and ensures end-to-end optical transparency to the signal, which is a very desirable property in modern protocol-independent transport networks. To improve the overall efficiency of the whole process, the aforementioned logical phases can be combined into a dynamic single-step implementation, resulting in an integrated RWA framework, where the use of a unique wavelength must be introduced as an additional constraint to the routing algorithm. The aforementioned single-step integrated RWA framework based on shortest path routing offers many practical advantages, starting from the fully decentralized and distributed routing architecture, providing excellent scaling properties with growing network dimensions (bounded by the polynomial Dijkstra algorithm's complexity—$O(N^2)$ or $O(E + N \log N)$ if a Fibonacci heap is used [20]) at the expense of a very limited administrative overhead.

However, some less obvious side effects, emerging behind the aforementioned advantages, require great attention in the design of a really effective RWA paradigm based on shortest path routing. In order to prevent unwanted reordering and other undesirable effects of multipath, particularly critical in the optical domain, a discrete *unsplittable* routing model has to be adopted so that the traffic demand from a source/destination pair must be satisfied by choosing and using only a single path between them. While preventing a large number of problems, the unsplittable nature of the routing model introduces additional difficulties, especially from the network engineering perspective. As links with lower costs are preferred by all communication demands, unsplittable shortest path routing protocols potentially lead to localized congestion phenomena (traffic concentrates on the lowest cost links) and unbalanced load distribution in the network. This may have severe effects on the number of service requests that can be satisfied at any time (for, as congestion increases, the number of rejected connections grows) and hence on the overall network providers'

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

revenue. Furthermore, the shortest path routing model implies the existence of some complex and subtle interdependences among the paths determined as the service demand evolves. More precisely, the choice of the end-to-end routing paths that constitute a valid solution can be only controlled in an indirect way by changing the costs assigned to the individual links. In addition, the associated weighting functions jointly influence all of the paths together, without any granular control on specific paths and their demands or service classes. Thus, designing weighting functions that cause the selection of globally efficient (both in terms of energy and traffic-related goals) end-to-end paths is the major challenge in modeling such a routing scheme.

IV. Multiobjective Weighting Function

The partially conflicting goals of serving the maximum possible number of users' requests, characterized by specific QoS constraints (concerning requested bandwidth, minimum acceptable link quality, etc.), while keeping the network resource usage fairly balanced, and optimizing the overall power consumption by reusing, as possible, energy-efficient paths across the network, give origin to a multivariate and multiobjective optimization problem, which is known to be NP-hard [21]. In the proposed approach, such a problem can be heuristically coped with by designing a composite scalar weighting function simultaneously combining the impact on the final solution of the different network, traffic engineering and energy-related objectives. This technique, commonly referred to as *weighted-sum* or *scalarization*, aggregates together n objectives by assigning a specific weight ω_i to each of them, according to the relative importance of the individual objective function $o_i(x)$ in the cumulative goal, resulting into a linear combination representing the whole optimization problem as

$$\max \sum_{i=1}^{n} \omega_i \cdot o_i(x) \tag{1}$$

subject to

$$\omega_i > 0, \qquad \forall\, i \in \{1,\ldots,n\} \ \sum_{i=1}^{n} \omega_i = 1. \tag{2}$$

It can be shown that the optimization of such single-objective convex sum is an efficient solution for the original multiobjective problem [22], i.e., its image belongs to the associated Pareto curve. The Pareto curve is the set of all efficient feasible solutions, i.e., the solutions whose objective vector is not dominated by any other solutions.[1] The shape of the Pareto curve sketches the tradeoff between the different objective functions $o_i(x)$. Clearly, modifying the weights ω_i may lead to different points of the curve, even if a uniform spread of the assigned weights does not lead to a uniform spread of points on the Pareto front, i.e., all solutions are clustered only in certain areas of the front. By slightly relaxing the convexity constraint, some objectives may be privileged over the others so that suboptimal solutions to the aggregate problem can be found.

A. Shortest Path Wavelength Routing in a Multiobjective Scenario

When formally defining the shortest path routing problem within a multiobjective optimization scenario, it must be considered that the meaning of the term "shortest" should be simultaneously associated to the different objectives involved. Therefore, the cost corresponding to each edge, which is the real decision maker in all of the available formulations, must result from the composition of multiple edge features, such as channel capacity, available resources, and power consumption. In order to model such behavior, a vector $\vec{w}_t(u,v) = (w_1(u,v),\ldots,w_n(u,v))$ of n different weights/costs must be associated at the time t to any edge (u,v) in the network graph G. Accordingly, each lightpath p_λ on G can be weighted by means of a vector $\vec{\Pi}_{p_\lambda} = (\pi^{(1)}_{p_\lambda},\ldots,\pi^{(n)}_{p_\lambda})$ where

$$\pi^{(i)}_{p_\lambda} = \sum_{(u,v)\in p_\lambda} w_i(u,v) \qquad \forall\, i \in \{1,\ldots,n\}. \tag{3}$$

We define a set of binary variables $x_{u,v}$ so that

$$x_{u,v} = \begin{cases} 1 & \text{if } (u,v) \in p_\lambda, \quad \forall\,\lambda \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

Hence, the problem of determining the multiobjective shortest path from an origin s to a destination d, in the presence of n different objectives [22], can be formalized as

$$\min_{k\in\{1,\ldots,n\}} \sum_{(u,v)\in E} w_k(u,v)\cdot x_{u,v} \tag{5}$$

subject to

$$\sum_{u:(u,v)\in E} x_{u,v} - \sum_{v:(v,u)\in E} x_{v,u} = \begin{cases} 1 & \text{if } u = s \\ 0 & \forall\, u \in N \setminus \{s,d\} \\ -1 & \text{if } u = d \end{cases}$$

$$x_{u,v} \geq 0 \qquad \forall\,(u,v) \in E. \tag{6}$$

The lightpath p_λ joining the endpoints (s,d) is an efficient solution to the aforementioned problem if it does not exist another lightpath q_μ between s and d such that

$$\pi^{(i)}_{q_\mu} < \pi^{(i)}_{p_\lambda} \qquad \forall\, i \in \{1,\ldots,n\}. \tag{7}$$

B. Weighting the Individual Objectives

An effective heuristic-based RWA scheme relying on the aforementioned multiobjective shortest path routing optimization model can be implemented by using the traditional Dijkstra algorithm, almost totally driven by an edge weighting function that dynamically associates a specific *cost* value to each edge in the network graph, with the effect of combining the $\vec{w}_t(u,v)$ vectors into a single scalar value $w_t(u,v)$. Hence, a correct choice of the weighting function is of fundamental importance for the overall success of the RWA framework. It should

[1] An objective vector *dominates* another objective vector if it is at least as good in all of the objectives; domination is strict if at least one inequality is strict.

condition the edge selection according to the best compromise between the following classes of objectives.

i) First, *traffic engineering objectives* essentially concern the ability to place the traffic associated to new incoming end-to-end connection requests characterized by specific QoS requirements (bandwidth, latency, BER, etc.) when sufficient capacity exists to accommodate the connection and to discard the associated request when such a capacity is not available (no end-to-end lightpath solution on the network is able to fully support the aforementioned requirements).
ii) On the other hand, *network engineering objectives* are associated to the ability of using at best the available capacity in order to accommodate as much connection requests as possible.
iii) Finally, *EA objectives* are related with the purpose of placing the traffic on the network so that the communication resources (circuits and nodes) that minimize the overall energy consumption are privileged, thus containing the energy-related expenses.

The traffic engineering objectives have an acceptance threshold: either a path can accommodate a connection request satisfying its QoS requirements, or it cannot. The weighting function must then adapt its behavior to each specific end-to-end connection request enabling the selective discard of all of the edges (i.e., the communication links) that do not satisfy the involved QoS requirements (admission policy). Accordingly, the cost of all of the edges corresponding to optical links that are not fully compliant to the aforementioned requirements is set to infinity, so that such edges are logically removed from the graph and cannot be selected in any search for feasible paths. Analogously, the edges on which all of the available wavelengths are currently utilized are also marked as unavailable by setting their cost to infinity until the next routing update. Furthermore, also in case of QoS compliance, the link parameters associated to potential QoS requirements are used to proportionally increase the edge cost in order to influence the path selection according to a rigid *best-fit* model, so that the selected paths will be preferentially composed by the edges that present the minimum gap between the requested amount of resource quality (e.g., the free bandwidth or the minimum latency or BER) and the available ones. In such a way, the risk of "over-provisioning" (routing a connection onto a path that is "too good" for it) is avoided, and the number of future requests that can potentially be accommodated is maximized. To this end, the aforementioned QoS parameters should be properly weighted to result in the desired impact on the edge cost, as described in Section IV-C.

In heavily loaded network scenarios, end-to-end connection requests cannot be satisfied because there are no wavelengths available on any of the links along all of the feasible paths. This phenomenon is commonly known as *connection blocking* (or rejection). Network engineering objectives essentially aim at reducing the blocking probability, ensuring that a maximal number of requests are accepted, thereby minimizing congestion and ensuring that the network resources are not over- or underutilized (unbalanced traffic loads). This can be intuitively achieved by privileging the selection of edges which guarantee that the maximum aggregated available *flow* between all of the source and destination node pairs is kept at the highest possible value. The larger the available maximum aggregated flow between a specific (source and destination) pair is, the smaller the blocking probability of connection requests between the involved pair will be, considering that the flow between all node pairs is a rough measure of how many routing options will be open when the (unknown) upcoming requests will be served. Aggregated flow information in the considered atomic unsplittable problem can be approximated by considering the *hit ratio* $\sigma_{(u,v)}$ of each edge (u, v), defined as the ratio between the number of times such edge has been selected and the total number of requests. The higher the hit ratio is, the greater the likelihood of an edge to be selected again in the future is; hence, edges with high hit ratio values have greater probability to become bottlenecks in the maximum end-to-end flow perspective. Therefore, the hit ratio provides an indication about the "criticality" of each edge for the overall network economy. Thus, in order to avoid as much as possible a reduction of the maximum source–destination flow, the weighting function has to assign to all of the most critical edges (the ones whose hit ratio exceeds a specific threshold) a cost value that is inversely proportional to the criticality measure, whose value is further amplified by using the residual capacity as an inverse multiplicative factor.

Finally, to handle EA objectives, the weighting function should provide the ability to consider as candidate paths for connections, in addition to the shortest and/or less congested paths, also the paths which minimize the energy consumption or the overall energy costs/bills. Accordingly, the weighting function w_t should properly condition the cost of each edge (u, v) by considering the fixed and variable energy consumption associated to the involved end-to-end interfaces and to the intermediate optical amplification $(a_{(u,v)})$ or 3R regeneration stages $(r_{(u,v)})$.

Clearly, all of the aforementioned strategies for dynamically determining the edge costs can be implemented by properly combining several per-link parameters to be weighted according to their relative importance with respect to both the individual (traffic, network, or energy related) objectives and the aggregated one. Therefore, we map the parameters into the following classes, in correspondence to the three aforementioned objectives.

i) **QoS**-related parameters, directly affecting the suitability of an optical link to carry a connection: $b_{(u,v)}$, $d_{(u,v)}$, and $e_{(u,v)}$;
ii) **network**-related parameters, directly affecting the blocking rate: $\lambda^{a(t)}_{(u,v)}$, $\lambda^{c}_{(u,v)}$, $f_{(u,v)}$, and $\sigma_{(u,v)}$;
iii) **energy-consumption**-related parameters, directly affecting the power draw: $a_{(u,v)}$, $r_{(u,v)}$, and $l_{(u,v)}$.

The parameters in the first group are the *threshold* parameters, which means that they indicate the requirement thresholds that must be met by an optical connection in order to support a connection request. Note that the value of a QoS parameter

TABLE I
SENSITIVITY OF THE PARAMETER CLASSES FOR THE DIFFERENT TRAFFIC OBJECTIVES

	QoS	Network	Energy
Traffic Engineering	high	medium	low
Network Engineering	medium	high	low
Energy-Related	low	medium	high

for a path p_λ is the minimum value of the parameter across all of the links comprising the path. Nevertheless, when many paths meet the threshold requirements for a connection request, one of these paths must be chosen somehow. In this case, it is reasonable to select the path that has the *lowest* QoS values, so that the costly high-performance links will be spared for use with more demanding connection requests (in accordance with a best-fit allocation strategy). Energy consumption-related parameters are clearly conditioned by specific interface and equipment-level power consumption characteristic, as will be described in detail in Section IV-E. Parameters take a wide range of different values. In order to be able to combine them in a significant way and avoid dominant effect, the possible values of the different parameters should be rescaled to a common interval. Thus, in order to be comparable, parameters are normalized into the interval [0, 1]. The relative importance of the parameter classes and their different sensitivity within the aforementioned three classes of objectives can be expressed by differentiating the growth rate of the weighting function with respect to the parameters. This has the effect of biasing the overall multiobjective optimization problem toward suboptimal solutions that privilege an objective class over the others but still keeping the other optimization tasks into an acceptable success range. The mapping between parameter and objective classes can be specified as in Table I. According to the classic sensitivity analysis theory, we can obtain a simple and effective measure of the sensitivity of the cost function C with respect to a specific parameter by estimating the value of the second-order partial derivative of the function C with regard to that parameter. With this approach, any change observed in the cost function will unambiguously be due to the specific parameter changed. For a high-sensitivity parameter x, $(\partial^2 C/\partial x^2)$ should be negative in the interval [0, 1], whereas for a low-sensitivity parameter y, $(\partial^2 C/\partial y^2)$ should be positive. In the former case, small increments in the parameter will quickly lead to saturation, whereas only values close to the maximum will have an effect in the latter case.

Let χ be a parameter class (QoS, network or energy related), and let S_χ denote the set of parameters in the class. Then

$$C = \sum_{\chi} \sum_{x \in S_\chi} x^{\alpha_\chi} \tag{8}$$

where α_χ is an assigned (tunable) constant for the class χ. We can assume, without loss of generality, that $\alpha_\chi > 0$. Clearly, the sign of $(\partial^2 C/\partial x^2)$ will be determined by $\alpha_\chi(\alpha_\chi - 1)$. Then, a reasonable starting point can be

$$\alpha_\chi = \begin{cases} \frac{1}{2} & \text{for high-sensitivity parameters} \\ 1 & \text{for medium-sensitivity parameters} \\ 2 & \text{for low-sensitivity parameters.} \end{cases} \tag{9}$$

Since we can directly associate the three aforementioned parameter classes with the individual objective functions $o_i(x)$ in (1), this implies that the convexity constraint in the scalarization weighted sum is slightly relaxed only for the weight associated to high-sensitivity parameters, in order to give more importance to the corresponding objective in the global multiobjective optimization problem. The effect of such choice is creating a perturbation effect in the problem optimality that results in the creation of three service classes whose expected treatment in terms of balancing of the individual objectives corresponds to the schema reported in Table I. The association of the individual service request to these classes is up to the carrier, according to specific agreements, economic conditions, policies, or internal strategic considerations.

C. QoS-Related Parameters

The QoS service level agreements (SLAs) of the connection requests $r = (s, d)$ have to be enforced on the lightpaths that are being established; therefore, the QoS-related parameters represent thresholds that admit or not paths in the network to be eligible routes for accommodating the incoming connection requests. When no routes satisfying the QoS SLAs are available from the source node s to the destination node d, the connection request has to be blocked. On the other hand, when more than one route connecting the involved source and destination nodes satisfies the SLAs, a selection criterion has to be employed in order to minimize the use of expensive network resources. In such a scenario, a best-fit policy is chosen, which reserves the least performing resources that are still able to satisfy the requested SLAs while leaving more expensive resources available for future highly demanding connection requests. We model such a best-fit criterion in (10), which assigns lower cost to links that best satisfy the SLA requirements (on bandwidth b_r, BER e_r, and delay d_r) and infinite cost to those links which do not comply with the specified requisites

$$C^{\mathrm{QoS}}_{(u,v)}(b_r, e_r, d_r) = \frac{1}{3}\left(T\left(\frac{b_{(u,v)} - b_r}{b_{(u,v)}}\right) + T\left(\frac{e_r - e_{(u,v)}}{e_r}\right) + T\left(\frac{d_r - d_{(u,v)}}{d_r}\right)\right) \tag{10}$$

where T is the threshold function discriminating between eligible and not eligible links

$$T(x) = \begin{cases} x & \text{if } x \geq 0 \\ \infty & \text{otherwise.} \end{cases} \tag{11}$$

D. Network-Related Parameters

The aim of the network parameters is to lower the connection blocking probability. The idea is to take traffic away from the most congested paths by using the hit ratio $\sigma_{(u,v)}$ of each edge (u, v) as a measure of its criticality and routing connections over the set of links that are currently underutilized, in a more traditional load-balancing (LB) fashion. The higher the number of available wavelengths, the lower will be the link cost. However, also the global link capacity and the number of physical fibers available on the link positively drive the link selection but in a more attenuated way, properly conditioned by a logarithmic trend. The resulting cost function involving the network-related parameters is reported in (12), in which edges are assigned a cost that is proportional to its actual congestion and hit ratio

$$C^N_{(u,v)} = \frac{1}{2}\left(\sigma_{(u,v)} + \frac{1}{\lambda^{a(t)}_{(u,v)} \cdot \log_\beta\left(\beta \cdot \lambda^c_{(u,v)} \cdot f_{(u,v)}\right)}\right) \quad (12)$$

with β being the base of the logarithm, a tunable parameter that characterizes the dampening effect of fiber and wavelength capacities on the link cost. In the presence of no available wavelength ($\lambda^{a(t)}_{(u,v)} = 0$), the link cost in (12) goes to infinity, and hence, the link is kept off from the graph in all of the shortest path calculations. Furthermore, the cost function explicitly considers the very special case in which we have only a single channel on a single fiber link ($\lambda^c_{(u,v)} \cdot f_{(u,v)} = 1$) that can be used for modeling non-WDM links that have to be selected only as a worst case alternative. In this case, the highest cost value (1) is assigned. Noting that $\log_\beta(\beta \cdot \lambda^c_{(u,v)} \cdot f_{(u,v)}) = 1 + \log_\beta(\lambda^c_{(u,v)} \cdot f_{(u,v)})$, it can be seen that a slightly better cost is assigned when a very limited number of fibers and wavelengths are available on the edge (u, v), with β being the threshold discriminating such bottleneck links. Finally, the logarithmic function assigns a low cost to the edge (u, v) when a high number of fibers and wavelengths are available on it.

E. Energy-Related Parameters

According to [23] and [24], we assume that the energy demand of a communication link is characterized by two fundamental components respectively associated to "fixed" and "variable" power absorptions. The fixed component is needed to keep the communication link "on," while the variable one depends on the traffic load that is currently carried by the link. Starting from these considerations, a sufficiently general per-link energy consumption model can be built, expressing the power consumption of any kind of communication circuit as a linear combination (according to [25]) of its static and traffic-dependent characteristics, such as the presence of intermediate amplification or regeneration stages, the involved endpoint interfaces, and their aggregated bandwidth in gigabit per second. More specifically, we define a power consumption function $P_{(u,v)}(x)$ expressing the power requirements of the link (u, v), characterized by an aggregated per-endpoint interface consumption $P_i(x)$ and by the number of amplification ($a_{(u,v)}$) and regeneration ($r_{(u,v)}$) devices, variably conditioned by a traversing traffic load x

$$P_{(u,v)}(x) = \underbrace{P_u(x) + P_v(x)}_{\text{link } (u,v) \text{ interfaces}} + \underbrace{\xi_{(u,v)} \cdot \varphi_{(u,v)} \cdot a_{(u,v)}}_{\text{Optical Amplification}} + \underbrace{\rho_{(u,v)} \cdot x \cdot r_{(u,v)}}_{3R \text{ Regeneration}} \quad (13)$$

where $\varphi_{(u,v)}$ is the power consumption value (measured in watts) for an individual optical amplifier used on the link (u, v), whereas $\rho_{(u,v)}$ (expressed in watts per gigabit per second) is the power required for regenerating a 1-Gb/s flow according to the regeneration technology used on (u, v). We assume, for simplicity sake, that all of the amplifiers and regenerators used on a single link (u, v) are of the same type. Furthermore, since optical amplifiers work simultaneously on the entire C-band, the contribution of $\varphi_{(u,v)}$ must be considered once for each link when determining the incremental per-link power-related cost at the time t, i.e., when the first wavelength of the link (u, v) is allocated to a lightpath, then the power consumption $\varphi_{(u,v)}$ of each amplifier activated on the link has to be added to $P_{(u,v)}(x)$. Otherwise, if at least a lightpath traverses (u, v), then the entire term can be zeroed to calculate the power increment of the new connection as amplifiers are already active. This is accomplished by using the binary variable $\xi_{(u,v)}$, defined as

$$\xi_{u,v} = \begin{cases} 0 & \text{if} \quad \exists\, p_\lambda \in \Lambda | (u, v) \in p_\lambda \\ 1 & \text{otherwise} \end{cases} \quad (14)$$

where Λ is the set of all of the active lightpaths on the network graph G. The aggregated per-endpoint interface consumption $P_i(x)$ can be modeled a linear function of its current load x

$$P_i(x) = \theta_i + x \cdot \vartheta_i \quad \text{with } \theta_i \le P_i(x) \le 2 \cdot \theta_i \quad (15)$$

so that, when an interface on the endpoint i is totally unloaded ($x = 0$), it is characterized by a fixed power consumption θ_i that is half of its maximum power demand [23] and, as the load increases, its power consumption linearly increases, up to its maximum value which is reached when the interface is fully loaded. Accordingly, also totally idle nodes are characterized by a (minimum) fixed power consumption since we assume that no sleep mode is available at the node level, to avoid wasting previous infrastructural investments, as reported also in [25]. The slope according to which the power consumption grows together with the load depends on a specific *scaling factor* ϑ_i, measured in watts per gigabit per second, representing the number of watts needed to route 1 Gb/s of traffic. The values for ϑ_i may usually range from 1 to 10 W/Gb/s [26] depending on the endpoint node features, where small-sized nodes require more energy per bit than bigger ones, which are characterized by the use of more energy-efficient technologies and usually are designed to aggregate large volumes of traffic [27]. Finally, we can define an energy cost function $C^E_{(u,v)}(x)$ for the link (u, v) as

$$C^E_{(u,v)}(x) = \frac{P_{(u,v)}(x)}{P_{\max}} \quad (16)$$

where

$$P_{\max} = \max_{(u,v)\in E, x=b_{(u,v)}} P_{(u,v)}(x) \quad (17)$$

is the maximum power consumption that can be experienced on any feasible end-to-end connection at its maximum load.

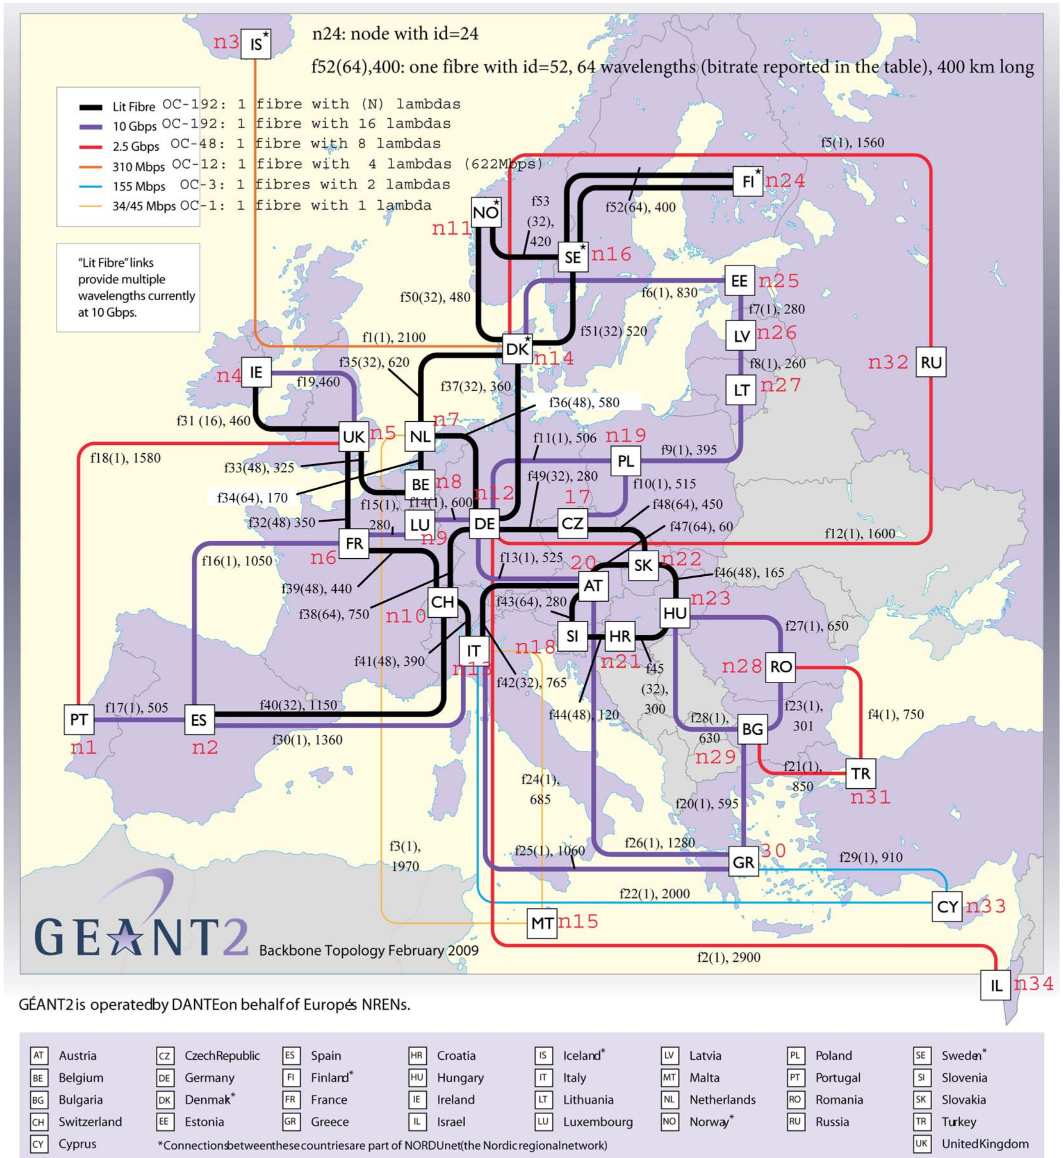

Fig. 1. GÉANT2 network topology [28] used in the simulations.

V. Performance Analysis

In order to evaluate the effectiveness of the proposed RWA framework according to the traditional carriers' goals, as well as its impact on the infrastructure-level energy consumption, we conducted extensive simulation studies on the well-known GÉANT2 Pan-European research and education network [28], modeled as an undirected graph in which each link has multiple fibers with a nonnegative capacity and a specific power demand depending from both its physical and technological features. The specific 34-node GÉANT2 topology used in our experiments is reported in Fig. 1, where only optical nodes are represented: each optical node (indicated as n-id_O) is connected to an electrical router $(id_E = id_O + 34)$ with one fiber link 1 km long with 32 λ, each with a capacity of 48 OC-units. We used in our analysis an ad hoc optical network simulation environment [29], allowing flexible and effective modeling of network topologies as well as traffic load generation, data recording, and postprocessing, running on an Intel Core i7-950 CPU at 3.07 GHz with 16-GB RAM and 64-b operating system server. In order to improve the significance of the obtained results and make them more easily comparable with the other experiences available in literature, we spent a significant effort on the use of realistic data in all of our experiments (network topology, traffic demands, cost, and power consumption models). The connection requests, bidirectional and satisfied by using the same wavelength in both directions, have been modeled by using different randomly generated or static [6], [30] traffic matrices. In the former case, the connections, generated according to a dynamic traffic scenario characterized by Poissonian arrivals, have been distributed uniformly among all of the network nodes, whereas in the latter one, the traffic volumes have been scaled proportionally to the reported traffic distributions. Each connection was characterized by a random bandwidth demand ranging from OC-3 to OC-48 units (i.e., from 155 Mb/s up to 2.5 Gb/s), a random delay ranging from 2 to 100 ms, and a BER ranging from 0% to 10%. The energy consumption data for

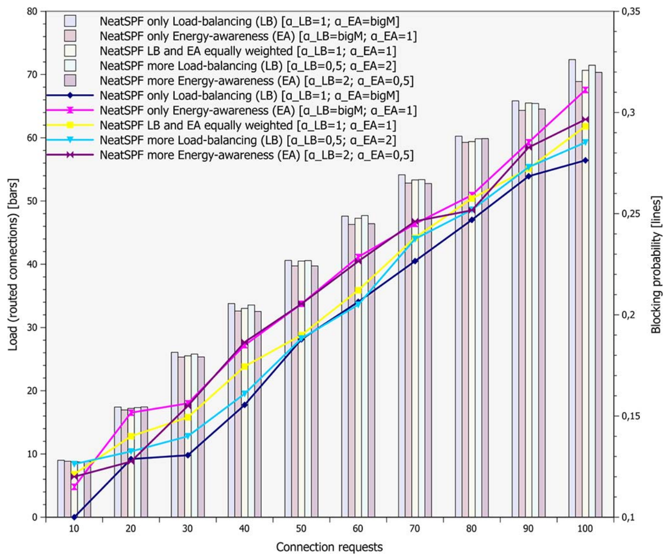

Fig. 2. Blocking probability (lines) and the load (bars) of the NeatSPF RWA scheme with different values of the α_χ parameters versus the connection requests.

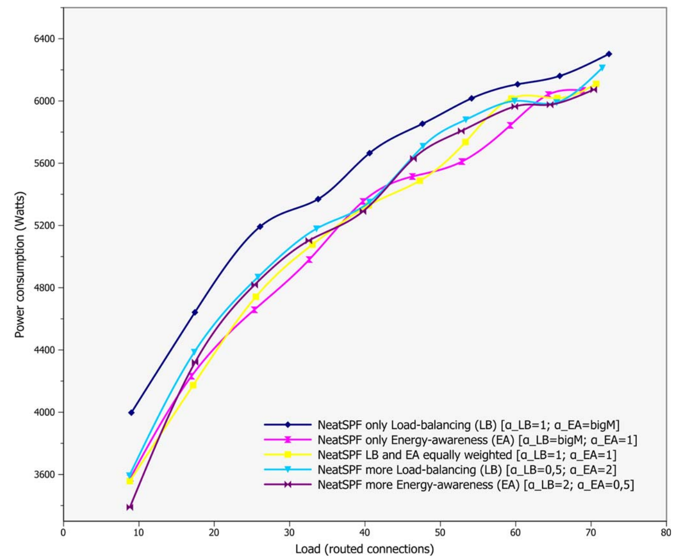

Fig. 3. Power consumption of the NeatSPF RWA scheme with different values of the α_χ parameters versus the load (routed connections).

each link have been populated with the real power consumption values taken from [25], [26], and [31]. Results have been determined with a 95% confidence interval not exceeding 6% of the indicated values, estimated by using the batch means method with at least 40 batches. As the network load grows, i.e., the number of busy connection resources increases more and more with respect to the free/released ones, we continuously monitored the overall network power demand and the network efficiency expressed by the blocking probability. Recall from Section IV-B that QoS-related parameters, modeling the traffic engineering objective, are threshold-based, which means that, given a connection request $r = (s, d)$, all lightpaths connecting s and d that do not satisfy the QoS requirements of r are not eligible to accommodate the connection request. Such a restriction is guaranteed by the dynamic online constrained shortest-path-first selection employed by the proposed RWA scheme. Among the feasible paths satisfying the specified QoS requirements, the RWA scheme evaluates network-related parameters (α_{LB}), modeling the network engineering objective (LB), and the energy-related parameters (α_{EA}), modeling the EA objective, eventually choosing the lightpath minimizing the cost function of (8). Note that, in (8), since α_χ appears as an exponent of x, with x being a parameter normalized in the interval [0, 1], $\alpha_\chi > 1$ values will lower the relative weight of the x parameter, while $0 < \alpha_\chi < 1$ will increase its relative weight, and $\alpha_\chi = 1$ will leave the x value unchanged (useful when we want to compare two different parameters, assigning them the same relative weight).

In order to assess the effectiveness of the proposed RWA scheme, in a first set of simulations, we present the performance of our approach (referred to as "NeatSPF," standing for "Network, Energy Aware and Traffic engineered Shortest Path First") varying α_{LB} (LB) and α_{EA} (EA) parameters. Then, in a second set of simulations, we compare these results with well-known state-of-the-art RWA algorithms. In Figs. 2 and 3, the results of NeatSPF are reported in terms of blocking and power consumption, respectively. Several assignments of the α_{LB} and α_{EA} parameters determine different behaviors of the NeatSPF algorithm. In detail, we set $\alpha_{\text{LB}} = 1$ and $\alpha_{EA} = \text{big}M$, with $\text{big}M \ll 1$ being a large constant, in order to obtain the extreme case in which NeatSPF only considers the network-related parameters to achieve the network engineering objective of maximizing the overall LB, thus minimizing the blocking probability. On the other hand, the opposite assignment of $\alpha_{\text{LB}} = \text{big}M$ and $\alpha_{EA} = 1$ makes NeatSPF pursue only the EA objective, discarding any network engineering constraint, thus minimizing the energy consumption. These two extreme cases are useful to study the lower and upper bounds of the NeatSPF performance. Then, in order to study the tradeoffs between the different objectives, two intermediate cases biasing the LB and the EA goals are considered. The *NeatSPF more LB* is obtained by assigning a higher weight ($\alpha_{\text{LB}} = 0,5$) to the network-related parameter and a lower one to the energy-related parameter ($\alpha_{EA} = 2$), slightly privileging the network engineering objective over the energy consumption. Speculatively, the *NeatSPF more EA* is obtained with the reverse assignment of weights to the network and energy-related parameters ($\alpha_{\text{LB}} = 2$ and $\alpha_{EA} = 0,5$), slightly privileging the EA objective over the LB. Finally, the *NeatSPF LB and EA equally weighted* is obtained by assigning the same weights to both parameters ($\alpha_{\text{LB}} = 1$ and $\alpha_{EA} = 1$), making the two objective directly comparable, in an effort to achieve a balance between the network and energy engineering objectives.

In Fig. 2, *NeatSPF only LB* achieves the lowest blocking probability, by routing the highest number of connection requests, followed by *NeatSPF more LB*. On the other hand, *NeatSPF only EA* rejects the highest number of connections, since it will select longer routes in order to pass through the least energy-consuming network elements (nodes, links, optical amplifiers, etc.). A slightly better performance is observed in *NeatSPF more EA* which considers some network-related

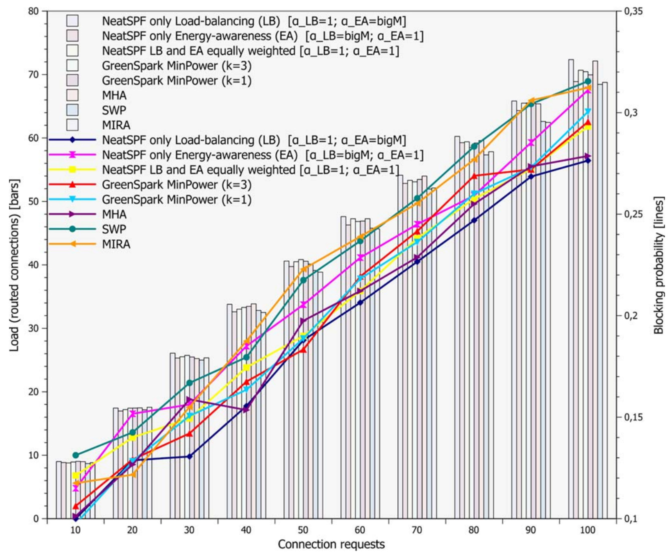

Fig. 4. Blocking probability (lines) and the load (bars) of the well-known state-of-the-art RWA schemes compared with the NeatSPF algorithm versus the connection requests.

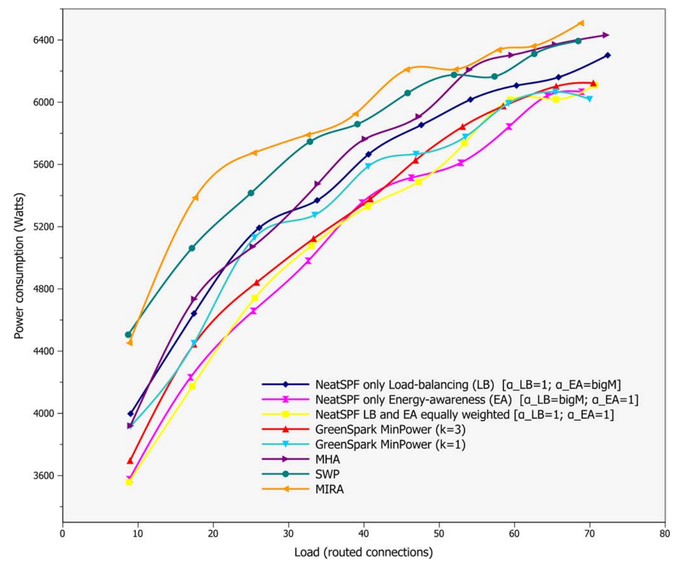

Fig. 5. Power consumption of the well-known state-of-the-art RWA schemes compared with the NeatSPF algorithm versus the load (routed connections).

parameter in its decision process, even if with quite low relative weight. Finally, the *NeatSPF LB and EA equally weighted* exhibits a well-balanced performance, standing just in the middle among the previous cases. The power consumption of the NeatSPF RWA scheme is reported in Fig. 3. As expected, the lowest energy consumption is achieved by the *NeatSPF only EA*, where the energy-related parameter assumes the highest weight and no LB is pursued, while the worst performance in terms of power consumption is exhibited by *NeatSPF only LB*, which is completely energy unaware. However, it is worth to note that the power consumption is easily decreased by assigning even a small weight to the energy-related parameter. The *NeatSPF more LB* sensibly decreases its power consumption with respect to *NeatSPF only LB*, while achieving good performance in terms of connection blocking. The hybrid *NeatSPF LB and EA equally weighted* exhibits a very low power consumption, in some points even lower than *NeatSPF only EA*. This phenomenon is due to the better load distribution achieved by the hybrid NeatSPF, which leaves more free resources to be used by future requests with respect to the pure *NeatSPF only EA* algorithm which, in turn, by occupying all of the lowest emitting routes at the beginning, will possibly not have enough resources and will be forced to select longer routes which will lead to slightly increased power consumption. In other words, the greedy choice of *NeatSPF only EA* made at each connection request may punctually lead to suboptimal routing in the long run (in this case, we can see that "the perfect is the enemy of the good" as explained numerically by the Pareto principle in its 80–20 rule [32]); therefore, a more balanced selection of parameters can lead to better results.

In Figs. 4 and 5, we report the blocking and the power consumption of several well-known state-of-the-art routing algorithms, whose implementation details are publicly available, compared with NeatSPF. Only some instances of NeatSPF are reported for comparison, since the other cases have been already presented. In particular, we show the performance of

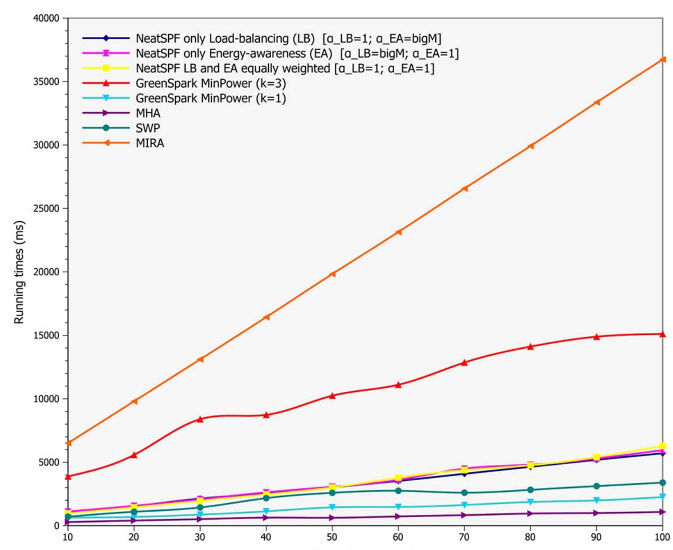

Fig. 6. Running times of the well-known state-of-the-art RWA schemes compared with the NeatSPF algorithm versus the connection requests.

minimum hop algorithm (MHA; [33]), shortest widest path algorithm (SWP; [34]), minimum interference routing algorithm (MIRA; [35]), and green smart parametric adaptive RWA algorithm based on K-shortest path (GreenSpark; [16]). MHA selects the shortest route (in terms of hop count) among source and destination; SWP selects, among the shortest routes, the widest one, i.e., the one with the highest residual capacity. MIRA selects the route that is foreseen to less interfere with future connection requests that are likely to come in the network. GreenSpark is based on a two-stage selection process: in the first step, the k best balanced paths are selected, according to an exclusive network engineering objective of optimizing the LB and thus minimizing the congestion and the consequent

This article has been accepted for inclusion in a future issue of this journal. Content is final as presented, with the exception of pagination.

blocking. In the second step, according to a pure EA objective, the lowest energy-consuming route among the k is finally selected to route the connection. In Fig. 4, we can observe that the lowest blocking is achieved by *NeatSPF only LB*, followed by MHA which achieves good performance owing to the limited number of connections in the network. Similar performances are obtained by *GreenSpark MinPower* $(k = 3)$ and *NeatSPF LB and EA equally weighted*. MIRA performs quite well at the beginning, but its performance degrades as the load increases, since it does not take into account the current traffic load in routing decisions. In Fig. 5, we can observe that the lowest power-consuming algorithms are *NeatSPF only EA* and *NeatSPF LB and EA equally weighted*. *GreenSpark MinPower* is the second less consuming algorithm, with *GreenSpark MinPower* $(k = 3)$ better than *GreenSpark MinPower* $(k = 1)$ as expected, since it has a higher degree of choice to lower the power consumption of connections. Following the increasing power-consuming algorithm, the *NeatSPF only LB* performs better than all of the remaining algorithms, which are, in the order, MHA, SWP, and MIRA. MHA, by selecting the shortest paths, achieves a lower power consumption than SWP and MIRA, which, in turn, select longest routes in an effort to reduce blocking. It is worth to note that MHA, SWP, and MIRA are totally energy unaware, since they do not consider energy-related parameters in their routing decision; GreenSpark, instead, was designed with both LB and EA in mind. However, here it suffers for the lack of grooming capability, for which it was originally conceived. Furthermore, being based on the k-shortest path, the computational complexity of GreenSpark linearly depends on k, while the NeatSPF family has the advantage of being faster, since its complexity does not depend on any parameter of the algorithm but just on the size of the network. Such a consideration leads us to the last chart shown in Fig. 6, in which the running times of the algorithms have been plotted during the simulations. NeatSPF exhibits a very low computational complexity (regardless of the α_χ values), overcome only by a constant factor by MHA, *GreenSpark MinPower* $(k = 1)$, and SWP. MHA, which has an extremely simple shortest path routing algorithm, has, however, well-known drawbacks in terms of blocking and power consumption too. *GreenSpark MinPower* $(k = 1)$ is slightly slower than MHA due to its more complex LB edge cost function, followed by SWP which has to add some calculation before selecting the final route. In general, all of these Dijkstra-based algorithms perform very well in terms of running times, routing 100 connections in less than 5 s (0,05 seconds per connection). Notably higher times are shown by *GreenSpark MinPower* $(k = 3)$, which shows the effects of the k-shortest path calculation. Finally, the slowest algorithm is shown to be MIRA, which suffers for the maximum flow calculation to identify the "critical" links performed each time a new lightpath has to be established. In conclusion, NeatSPF, owing to its parametric cost function, can be easily tuned to achieve either the best LB or the lowest power consumption with respect to the other algorithms with which it has been compared. It also demonstrated that an optimal tradeoff can be achieved by an appropriate selection of the network and energy-related parameters according to the objective of the network operator while maintaining computational complexity very low and therefore providing more than satisfactory network responsiveness.

VI. Conclusion

We have presented a simple but effective RWA framework, based on shortest path routing with an adaptive link weighting function. It is designed to be suitable for real-time network control and management as well as effective in providing good wavelength utilization together with low blocking probabilities, leading to efficient usage of the network's resources. It also integrates EA in its decision process, driven by a flexible and configurable energy model, in order to support sophisticated strategies for containing the network's energy consumption and reducing the associated costs. Apart from being a successful wavelength routing scheme, the most significant added value of the proposal is the inherent flexibility of the multiobjective optimization model, where multiple tunable parameters can be used to drive the solution toward several sections of the Pareto curve. This leads to suboptimal solution to the aggregate problem that may privilege some specific objective (e.g., the containment of energy consumption) over the others, according to the dynamically changing carriers' needs, while maintaining an affordable polynomial time complexity which makes it suitable for online routing employed by modern control planes. Extensive simulation experiments, conducted on several real network topologies, resulted in a good tradeoff between the different involved (and apparently conflicting) users' and carrier's optimization objectives, demonstrating that the proposed approach is computationally inexpensive, easy to implement, quite balanced in its results, and, hence, ready for deployment in real-world optical networks.

References

[1] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: An approach to high bandwidth optical WAN's," *IEEE Trans. Commun.*, vol. 40, no. 7, pp. 1171–1182, Jul. 1992.

[2] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen, "Reducing the energy consumption of Ethernet with adaptive link rate (ALR)," *IEEE Trans. Comput.*, vol. 57, no. 4, pp. 448–461, Apr. 2008.

[3] X. Yuan, R. Melhem, and R. Gupta, "Distributed path reservation algorithms for multiplexed all-optical interconnection networks," *IEEE Trans. Comput.*, vol. 48, no. 12, pp. 1355–1363, Feb. 1999.

[4] W. Ben-Ameur and E. Gourdin, "Internet routing and related topology issues," *SIAM J. Discret. Math.*, vol. 17, no. 1, pp. 18–49, 2004.

[5] A. Bley, "Inapproximability results for the inverse shortest paths problem with integer lengths and unique shortest paths," *Networks*, vol. 50, no. 1, pp. 29–36, Aug. 2007.

[6] R. Ramaswami and A. Segall, "Distributed network control for optical networks," *IEEE/ACM Trans. Netw.*, vol. 5, no. 6, pp. 936–943, Dec. 1997.

[7] H. Zang, L. Sahasrabuddhe, J. Jue, S. Ramamurthy, and B. Mukherjee, "Connection management for wavelength-routed WDM networks," in *Proc. GLOBECOM*, 1999, vol. 2, pp. 1428–1432.

[8] A. Mokhtar and M. Azizoğlu, "Adaptive wavelength routing in all-optical networks," *IEEE/ACM Trans. Netw.*, vol. 6, no. 2, pp. 197–206, Apr. 1998.

[9] I. Chlamtac, A. Farago, and T. Zhang, "Lightpath (wavelength) routing in large WDM networks," *IEEE J. Sel. Areas Commun.*, vol. 14, no. 5, pp. 909–913, Jun. 1996.

[10] W. Liang, G. Havas, and X. Shen, "Improved lightpath (wavelength) routing in large WDM networks," in *Proc. 18th ICDCS*, 1998, pp. 516–523.

[11] M. Gupta and S. Singh, "Greening of the Internet," in *Proc. SIGCOMM*, New York, NY, USA, 2003, pp. 19–26.

[12] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, "An energy saving routing algorithm for a green OSPF protocol," in *Proc. INFOCOM IEEE Conf. Comput. Commun. Workshops*, 2010, pp. 1–5.

[13] A. Cianfrani, V. Eramo, M. Listanti, and M. Polverini, "An OSPF enhancement for energy saving in IP networks," in *Proc. IEEE INFOCOM WKSHPS*, 2011, pp. 325–330.

[14] M. Hasan, F. Farahmand, J. Jue, and J. Rodrigues, "A study of energy-aware traffic grooming in optical networks: Static and dynamic cases," *IEEE Syst. J.*, vol. 7, no. 1, pp. 161–173, Mar. 2013.

[15] L. Chiaraviglio, M. Mellia, and F. Neri, "Reducing power consumption in backbone networks," in *Proc. IEEE ICC*, Piscataway, NJ, USA, 2009, pp. 2298–2303.

[16] S. Ricciardi *et al.*, "An energy-aware dynamic RWA framework for next-generation wavelength-routed networks," *Comput. Netw.*, vol. 56, no. 10, pp. 2420–2442, Jul. 2012.

[17] F. Palmieri, U. Fiore, and S. Ricciardi, "SPARK: A smart parametric online RWA algorithm," *J. Commun. Netw.*, vol. 9, no. 4, pp. 368–376, Dec. 2007.

[18] S. Ricciardi, J. Wang, F. Palmieri, D. Careglio, and L. Dittman, "An energy-aware engineered control plane for wavelength-routed networks," *Trans. Emerg. Telecommun. Technol.*, vol. 26, no. 2, pp. 231–249, Feb. 2015.

[19] M. Saad and Z.-Q. Luo, "Wavelength selective WDM network design: How bad is shortest path routing?" in *Proc. IEEE ICSPC*, Nov. 2007, pp. 532–535.

[20] H. R. Lewis and L. Denenberg, *Data Structures and Their Algorithms*, 1st ed. Boston, MA, USA: Addison-Wesley, 1997.

[21] P. Serafini, "Some considerations about computational complexity for multiobjective combinatorial problems," in *Recent Advances and Historical Development of Vector Optimization*, vol. 294, ser. ser. Lecture Notes in Economics and Mathematical Systems. Berlin, Germany: Springer-Verlag, 1986.

[22] M. Caramia and P. Dell'Olmo, *Multi-Objective Management in Freight Logistics: Increasing Capacity, Service Level and Safety With Optimization Algorithms*. London, U.K.: Springer-Verlag, 2008.

[23] J. Chabarek *et al.*, "Power awareness in network design and routing," in *Proc. 27th IEEE INFOCOM*, Apr. 2008, pp. 457–465.

[24] A. Adelin, P. Owezarski, and T. Gayraud, "On the impact of monitoring router energy consumption for greening the Internet," in *Proc. 11th IEEE/ACM GRID*, Oct. 2010, pp. 298–304.

[25] S. Ricciardi *et al.*, "Towards an energy-aware Internet: Modeling a cross-layer optimization approach," *Telecommun. Syst.*, vol. 52, no. 2, pp. 1247–1268, Feb. 2013.

[26] BONE project, "WP 21 TP Green Optical Networks, D21.2b Report on Y1 and Updated Plan for Activities," 2009. [Online]. Available: http://www.ict-bone.eu/portal/landing_pages/bone_deliverables/bone_216863_-_d21-2b_-_report_y1.pdf

[27] W. Vereecken, W. Van Heddeghem, D. Colle, M. Pickavet, and P. Demeester, "Overall, ICT footprint and green communication technologies," in *Proc. 4th ISCCSP*, 2010, pp. 1–6.

[28] Geant, "The Geant2 Network," Jun. 2011. [Online]. Available: http://www.geant2.net/server/show/nav.00d007009

[29] F. Palmieri, U. Fiore, and S. Ricciardi, "Simulnet: A wavelength-routed optical network simulation framework," in *Proc. IEEE ISCC*, Jul. 2009, pp. 281–286.

[30] S. Uhlig, B. Quoitin, J. Lepropre, and S. Balon, "Providing public intradomain traffic matrices to the research community," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 83–86, Jan. 2006.

[31] F. Idzikowski, "Power Consumption of Network Elements in IP Over WDM Networks," Telecommun. Netw. Group, Techn. Univ. Berlin, Berlin, Germany, TKN Techn. Rep. Ser. TKN-09-006, Jul. 2009. [Online]. Available: http://www.tkn.tu-berlin.de/publications/papers/powerNumbers_final.pdf

[32] V. Pareto and A. N. Page, *Manuale di Economia Politica*. New York, NY, USA: A.M. Kelley, 1971.

[33] D. Awduche *et al.*, "Extensions to RSVP for LSP tunnels," *Internet Draft Draft-Ietf-Mplsrsvp-Lsp-Tunnel-04.Txt*, 1999.

[34] R. Guerin, A. Orda, and D. Williams, "QoS routing mechanisms and OSPF extensions," in *Proc. IEEE GLOBECOM*, Nov. 1997, vol. 3, pp. 1903–1908.

[35] K. Kar, M. Kodialam, and T. V. Lakshman, "Minimum interference routing of bandwidth guaranteed tunnels with MPLS traffic engineering applications," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 12, pp. 2566–2579, Dec. 2000.

Francesco Palmieri received the M.S. degree and the Ph.D. degree in computer science from the University of Salerno, Fisciano, Italy.

He is an Associate Professor with the Department of Computer Science, University of Salerno. He serves as the editor-in-chief of an international journal and participates in the editorial board of others. His research interests concern advanced networking protocols and architectures, and network security.

Sergio Ricciardi received two M.Sc. degrees in computer science from the University of Naples Federico II, Naples, Italy, and the Technical University of Catalonia (UPC), Barcelona, Spain, and the Ph.D. degree in computer architecture from the UPC.

He is a Research Associate with the Advanced Broadband Communications Center, Department of Computer Architecture, UPC. His research interests mainly focus on energy-aware network protocols and energy-oriented optimizations for grids and clouds.

Ugo Fiore received M.S. degree from University Federico II of Naples, Italy, and Ph.D. degree from University of Salerno, Italy.

He is currently with Department of Molecular Medicine and Medical Biotechnologies of Federico II University. He leads the Network Operations Center, Federico II University, Naples, Italy. He began his career with the Italian National Council for Research, and he has also more than ten years of experience in the industry, developing software support systems for telco operators. His research interests focus on optimization techniques and algorithms aimed at improving the performance of high-speed core networks and on security-related algorithms and protocols.

Aniello Castiglione (S'04–M'08) received the B.S. and Ph.D. degrees in computer science from the University of Salerno, Fisciano, Italy.

He joined the Department of Computer Science, University of Salerno, in 2006. He serves as a reviewer for several international journals and has been a member of international conference committees. His research interests include communication networks, information forensics, security, and cryptography.

Dr. Castiglione is a member of various associations, including IEEE and ACM.

Davide Careglio received the M.Sc. and Ph.D. degrees in telecommunications engineering from the Technical University of Catalonia (UPC), Barcelona, Spain, and the Laurea degree in electrical engineering from Politecnico di Torino, Torino, Italy.

He is an Associate Professor with the Department of Computer Architecture, UPC. His research interests include networking protocols, with emphasis on optical switching technologies, and algorithms and protocols for traffic engineering and QoS provisioning.