Energy-Efficient Resource Provisioning Algorithms for Optical Clouds

J. Buysse, K. Georgakilas, A. Tzanakaki, M. De Leenheer, B. Dhoedt and C. Develder

Abstract—**Rising energy costs and climate change have led to an increased concern for energy-efficiency (EE). As Information and Communication Technology (ICT) is responsible for about 4% of total energy consumption worldwide, it is essential to devise policies aimed at reducing it. In this paper, we propose a routing and scheduling algorithm for a cloud architecture, which targets minimal total energy consumption by enabling switching off unused network and/or Information Technology (IT) resources, exploiting the cloud-specific anycast principle. A detailed energy model for the entire cloud infrastructure comprising wide area optical network and IT resources is provided. This model is used to make a single-step decision on which IT end points to use for a given request, including the routing of the network connection towards these end points. Our simulations quantitatively assess the EE algorithm's potential energy savings, but also assess the influence this may have on traditional Quality of Service parameters such as service blocking. Furthermore, we compare the one-step scheduling with traditional scheduling and routing schemes, which calculate the resource provisioning in a two-step approach (selecting first the destination IT end point, and subsequently using unicast routing towards it). We show that depending on the offered infrastructure load, our proposed one-step calculation considerably lowers the total energy consumption (reduction up to 50%) compared to the traditional iterative scheduling and routing, especially in low to medium load scenarios, without any significant increase in the service blocking.**

Index Terms—**Energy Efficiency, WDM networks, Cloud Computing, Resource Provisioning**

I. INTRODUCTION

ICT equipment, facilities and the processes to control this equipment consume up to 4% of the world's total energy budget, implying a considerable environmental impact in terms of greenhouse gas emissions [3], [33]. This paper addresses the energy expenditure for an integrated network and IT infrastructure that can support cloud and grid architectures. The blueprint for the Grid architecture was laid out in [22]: in analogy with a power grid, users could get access to computing power on demand. Grid customers would generally create an application, submit it using the grid middleware, and wait until the job finishes in order to collect the results. A more commercial version, the cloud infrastructure, extends this concept and applies the Infrastructure-as-a-Service (IaaS) concept. The consumer decides on a number of Virtual Machines (VMs), which are to be deployed on real physical devices, to which access is granted during a certain time. Cloud computing is seen as an energy-efficient architecture, as end users are limited to low-power devices, while processing power (and hence also a large part of energy consumption) is moved to the cloud [3]. Moreover, cloud architectures provide aggregation points for workloads that would otherwise be run on separate devices. This means that demands can be consolidated through statistical multiplexing and hosts can be better utilized. Grid and cloud architectures both require the pooling and coordinated allocation of a large set of distributed resources, and we aim to optimize their utilization to reduce the overall energy consumption. As the network prerequisites for the applications we envisage are very demanding (e.g., high bandwidth and low latency), we assume an optical circuit-switched network based on Wavelength Division Multiplexing (WDM) and thus consider an optical grid/cloud context (see [16] for a recent overview on such optical grids/clouds). We jointly optimize energy consumption of network and IT resources using a scalable algorithm by exploiting the anycast principle. Anycast reflects the idea that a user is generally not interested in the location where his workload is processed "in the cloud"), as long as the requirements (which have been set in advance by so-called Service Level Agreements, SLAs [35]) are met. Hence, freedom arises as to where to execute a job or to place a VM. This paper presents a heuristic that for a given request finds (i) an IT end point to process the request (the scheduling problem) and (ii) a route from the requesting source to that IT end point in the optical network (the routing problem). Requests arrive sequentially and we are solving the online routing problem, as opposed to the offline version (e.g., [8]), which has an a priori known request vector, expressing for each source the number of requests which need to be served. Our algorithm minimizes energy consumption by either trying to share as much active resources as possible (avoiding a startup cost for each newly activated resource) or by allowing switching-off idle resources. The remainder of this paper is structured as follows. Section II starts off with an overview of related work, where we indicate the novelty of our contribution. Next, in Section III, we present our power model for the grid/cloud infrastructure (including quantified power consumption figures). In Section IV we detail

the routing/scheduling algorithms, which are subsequently investigated by a detailed simulation case study in Section V. Final conclusions and future work are discussed in Section VI.

II. Related work

A. Optical network energy models

Optical network technology is incontestably energy-efficient. The authors of [30] present a comparison of different IP-over-WDM architectures, demonstrating that a translucent optical architecture (i.e., the optical signal is periodically regenerated by all-optical 3R regenerators) can save up to 60% of energy compared to classical technologies (e.g., where optical signal regeneration is done in the electronic IP layer). Comparable conclusions are drawn in [2], [4], [36], [37]: optical nodes generally consume less power than electronic ones, especially optical circuit-switched architectures based on MEMS switching devices. Furthermore, it has been demonstrated that an energy-efficient network design is coincidentally a cost-efficient design since router ports play a dominant role in both energy and capital cost. In Section III we will further discuss the model for the network energy consumption based on [2].

B. IT energy models

Regarding electricity consumption of servers and data centers, [28] indicates that power usage of all servers in the U.S. accounts for a substantial fraction of total US electricity consumption, which even doubles when auxiliary infrastructure (cooling, water pumps, etc.) is included. This is the reason that our energy model takes this supporting infrastructure into consideration. The authors of [21] investigate the power properties for servers, individual racks and clusters. They also demonstrate that nameplate ratings (manufacturer's prediction of power use) have little or no value as they tend to overestimate actual peak usage which explains why we take the parameters for a server's energy consumption from real life measurements. Secondly, they investigate the influence of Dynamic Voltage Scaling (DVS): this method reduces energy consumption by slowing down the rate of CPU processing since the faster the processing rate, the higher the energy consumption. Our energy model for a server is based on this work, while we changed the model for racks and data centers using up-to-date cooling techniques. Another strategy for IT energy minimization is server consolidation. The authors of [32] have investigated this while also trying to predict which nodes will need to be powered down/on in the future. These previous ideas, i.e., server consolidation and DVS, are combined into a single formalism in [10].

C. Energy-efficient operation in optical networks

Switching off network elements to save energy has been evaluated in [11] for an offline scenario (i.e., traffic is known beforehand - as opposed to our approach). The authors demonstrate that, for the scenarios under consideration, there is an energy saving potential of total network energy. Similar conclusions are drawn in [9], which extends [11] with an empirical study for power consumption of a router. Scaling down the logical IP topology in an IP-over-WDM network is investigated in [34]. The authors assign a higher cost for IP links having a load below a certain threshold, deviating traffic flows from these links to remove the IP links from the IP topology. Results show that a high threshold only favors architectures which make use of equipment with high idle power (e.g., as demonstrated in [9]), as for the more EE equipment longer paths (which lead to more transit traffic in core interfaces) lead to an increase in power consumption, as the power requirements are proportional with interface bandwidth. The effect of putting clusters of network nodes in a sleep state, by routing to an appropriate location (thus using anycast as described in Section I), is examined in [5]. Our work differs in that we allow powering down individual network nodes, as well as network links. Power-awareness combined with resiliency aspects is investigated in [25], but only considers the network resources and a unicast scenario: the authors achieve power reduction by putting network resources into a sleep state when they are used as backup resources and demonstrate the effectiveness by comparing different routing algorithms. Although in our work we do not consider protection, we are using a similar network energy model where different components of network entities can be shut down. In [23] the authors propose to groom sub-wavelength traffic into light paths, while allowing a modular network node to offer energy savings by powering on/off chassis, modules or ports depending on traffic entering the network node. They conclude that at off-peak hours, a traditional (minimizing the number of light path setups per request) and energy-aware approach have about the same energy consumption. In peak conditions however, the energy-aware approach outperforms the traditional strategy (regarding energy consumption) since more traffic requests can be routed through already active components. A comprehensive overview of ongoing research regarding energy efficiency in telecom networks, with a specific emphasis on optical technologies, is presented in [39]. For several network architectures (metro, access and core), energy minimization opportunities are investigated and related ongoing standardization efforts are overviewed. They also indicate that there might be a potential in scheduling jobs in a grid context, allowing servers to be switched-off. We build on this concept, while also considering the energy consumed in the optical core

network in between the IT end points and the data centers.

D. Energy-efficient operation in data centers

The work in [7] reviews methods and technologies currently deployed for energy-efficient operation of computer hardware and network infrastructure, particularly in cloud contexts. They demonstrate that data center scheduling can influence energy consumption and that virtualization of resources can be beneficial from an energy consumption perspective. These policies only focus on one part of the cloud, either the network or the data center, but no work tries to combine both realms. The authors indicate possible improvements, such as reducing energy consumption due to communications, which is the aim of this paper. In [1] the authors investigate how to build a cluster-scale network (within the data center premises) whose power consumption is proportional to the amount of traffic it is transmitting. They demonstrate that a flattened butterfly topology (similar to a fully connected torus) operated at a data rate proportional to the offered traffic intensity of the data center, is the most energy-efficient intra data center network design. The work in [27] presents an intra data center scheduling approach (for a three-tier network) that combines energy efficiency and network awareness: it allows analyzing data received from the switches and links and takes actions based on the network feedback. The scheduling approach avoids hotspots within a data center while minimizing the number of computing servers required for a job execution (job consolidation). In our work however, we do not consider advanced intra data center scheduling of jobs, but enforce a First Come First Serve (FCFS) policy. Note that this work complements ours, where we do not provide detailed modeling of the intra data center network. We believe that incorporating such more advanced intra data center scheduling will not impact our qualitative discussions pertaining to the importance of jointly considering (core) network and data center energy consumption.

E. Energy-efficiency in an integrated infrastructure

Dynamically powering on/down servers to address actual demand in a grid context has been investigated in [17]. The authors propose a power-aware scheduling scheme that reduces IT power consumption. The penalty is an increase in network utilization because longer paths are used. Our work builds on this concept by also considering the optical network, jointly optimizing the utilization of IT and network resources used to serve all demands. Chapter six of [20] proposes two ways to reduce energy consumption: (i) a novel, integrated optical network and IT infrastructure and (ii) an energy-aware service plane architecture. The first optimization consists of distributing a fraction of IT nodes from IT resource sites at the network edge into the network core so that network operators can benefit from the existing space, cooling and power of switching nodes in the core of the network. The second optimization consists of a resource orchestration formulation taking into account energy-aware parameters, such that the selection of network and IT resources is optimized to reduce the overall power consumption. Depending on the scenario, the new integrated infrastructure can improve energy efficiency up to 45% and the EE resource orchestration up to 10%. Another attempt to define a comprehensive energy model where network and IT resources are treated in an integrated way has been examined in our earlier work [8]: it addresses the energy-efficient operation of integrated network and IT infrastructures in the context of cloud computing in an offline scenario. There, we proposed energy-efficient routing and IT allocation algorithms using MILP, by allowing switching-off several IT and networking elements and by exploiting the anycast principle. More specifically, comparing joint minimization of both network and IT energy provides energy savings of the order of 3% to 55% compared to the network energy minimization only approach, depending on the granularity of a data center to switch on/off a set of servers. On the other hand, pure network-energy minimization allows energy savings of the order of 1% to 2% of the total energy budget compared to shortest path routing (i.e., energy-unaware). Although [20] and [8] indicate that treating network and IT resources jointly allows for energy optimization, their approaches are difficult to adopt in real settings since they suffer from scalability issues and cannot produce results in a reasonable time frame. Therefore, we extend this earlier work in two ways: (i) we update the energy model to include energy-efficient cooling units (In-Row Cooling) and (ii) we tackle the problem in an online scenario to obtain results in a faster time frame.

F. Contribution of this paper

Our study extends previous works in several ways. Our first main contribution consists in the integration of the network and the IT realm: by considering optical and IT resources in the same scheduling and routing step, we lower the overall energy consumption considerably. Moreover, we provide a one-step anycast calculation and compare it with a sequential computation (two-step, first IT data center selection, then routing towards it) and show the benefits of our unified approach in terms of power consumption and service blocking. Furthermore, we allow switching off network nodes, links, servers, racks and data centers in contrast to previous works which mainly focused on either the core network or the IT infrastructure. Secondly, our unified energy model considers a cooling system, namely in-row cooling, which proves to be the
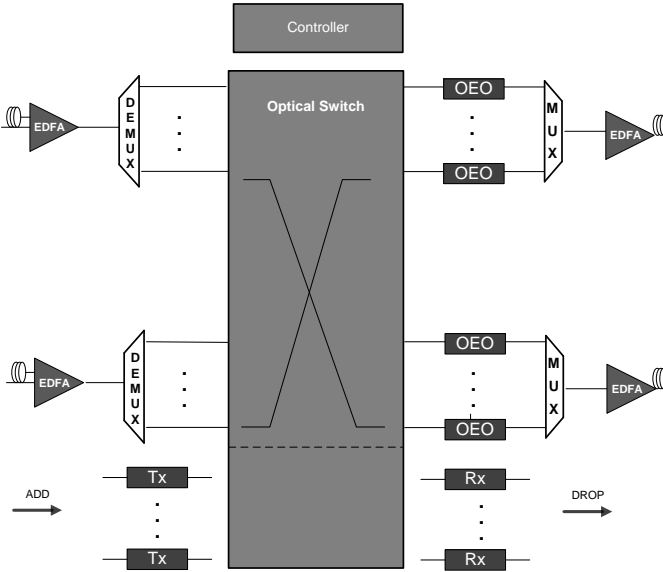
Fig. 2: Layout of an opaque OXC. All elements except for the (de)multiplexer consume energy.

most energy-efficient cooling system for data centers available today [6]. Thirdly, we treat the problem from an online perspective, as opposed to the offline scenario, resulting in an algorithm that is able to dynamically allocate resources in a short time frame. Lastly, we focus not only on energy consumption, but we also investigate the influence of EE scheduling and routing on traditional QoS parameters such as service blocking and average resource load (as opposed to e.g. [20]).

III. MODELING

A. Topology modeling

We model the optical network as a bidirectional graph $G = (S, C, E)$ where S is the set of source nodes, comprising optical cross-connects (OXCs) generating requests. C is the set of core OXCs, which (as opposed to source OXCs) may be switched off completely. E is the set of optical fiber links connecting all OXCs ($S \cup C$). Each fiber is assumed to have W wavelengths. The topologies used in our study are presented in Fig. 1. Furthermore we define $D \subseteq S$ as the set of destination sites, i.e., these OXCs $d \in D$ are connected to a data center. Our graph model employs auxiliary links between the data center objects and $d \in D$, which we will denote as virtual links, as they do not represent actual physical links. All fiber links incident to $d \in D$ have $2W$ wavelengths, as these are the end points of all paths and need more capacity to prevent network blocking. We assume that all data centers have the same characteristics: each data center d has n racks, each containing s servers with idle and peak power characteristics described and measured in [13].

B. Network energy modeling

We assume OXCs based on a photonic switching matrix that is realized by 3D Micro-Electrical-Mechanical-Systems (MEMS) [31]. Each OXC supports a number of input and output fibers ports, each employing a maximum number of wavelengths W. It is assumed that each OXC is equipped with wavelength converters at the output so that a light path (a wavelength path including all used wavelength links from source to destination) can be established between any source-destination pair as long as there is a free port, avoiding situations of wavelength blocking. Apart from the passive elements, being the Multiplexers (MUX) and De-Multiplexers (DEMUX), Fig. 2 illustrates the active elements of the OXC: the switch matrix, one Erbium-Doped Fiber Amplifier (EDFA) per input/output fiber port and one transmitter (Tx) and one receiver (Rx) pair per light path. The OEO transponders support full wavelength conversion. The number of through (express) ports ($ports_{through}$) is calculated as the number of input fibers times the fiber wavelength capacity W. The add/drop ports (e.g., for traffic from/to a local data center) are denoted as $ports_{a/d}$). The active incoming/outgoing fibers are represented as f_{in} and f_{out} respectively. The network power is completely determined by the power consumption of all the OXCs and the optical fiber links. The power expenditure of an OXC (P_{OXC}) depends on the constant power consumption of (i) the switch fabric (P_{sf}), (ii) the receivers and transmitters (P_{transc}), (iii) the wavelength converters (P_{conv}), (iv) the optical amplifiers (P_{ampl}) and (v) the controller power ($P_{control}$) for the OXC. Eq. 1 show how these figures are used in the total power consumption model of the OXC, while Table I shows typical values for their parameters.

$$P_{OXC} = P_{control} + P_{sf} + P_{transc} + P_{conv} + P_{ampl} \quad \text{(1a)}$$
$$P_{transc} = ports_{a/d} \times P_{Tx/Rx} \quad \text{(1b)}$$
$$P_{conv} = ports_{through} \times P_{transponder} \quad \text{(1c)}$$
$$P_{ampl} = (f_{in} + f_{out}) \times P_{edfa} \quad \text{(1d)}$$

Regarding the fiber links of the optical networks, the power consuming elements are the optical amplifiers installed per span. The amplifier span length ($span$) is assumed to be 80km. Hence, the power consumption P_l of a fiber link depends on its length ($|l|$) and can be calculated as shown in Eq. 2 (Note that -1 is used because the first span can be covered by the EDFA at the fiber output port of the OXC).

$$P_{link}(l) = \left(\frac{|l|}{span - 1} \right) \times P_{edfa} \quad \text{(2)}$$

The total network energy consumption is then computed by Eq. 3. Note that we multiply the network energy with a factor called the Power Usage Effectiveness (PUE), to account for energy used for cooling and power delivery for the network resources, and typically

Fig. 1: The topologies considered in this study, containing 28 OXCs. The circled OXCs are the eight core nodes. The dotted lines between the DC's and the network nodes are the virtual links. All topologies were gathered from [15].

(a) Dense topology (57 fiber links, node de-(b) Basic network (40 fiber links, average(c) Sparse network (33 fiber links, average
gree 4.03)   node degree 3.07)   node degree 2.4)
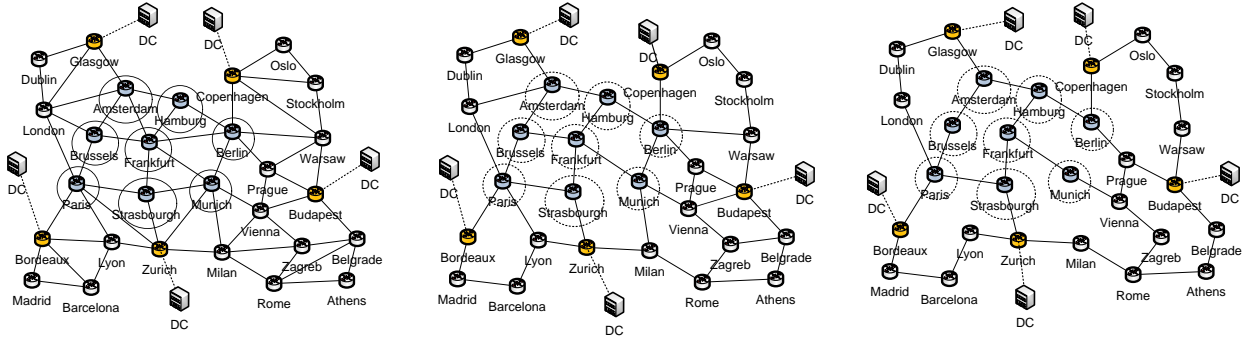
amounts to around 2 [4]. We have chosen not to model the power delivery and cooling chain in more detail for the network. Indeed, the values for cooling and power delivery for a data center and an OXC differ in several orders of magnitudes. Hence, a more accurate power-cooling model for OXCs would not change our results qualitatively (while a simple PUE approach as opposed to our current model for the data center would).

$$P_{network} = PUE \times \left(\sum_{n \in S \cup C} P_{OXC} + \sum_{l \in E} P_{link}(l) \right) \quad (3)$$

C. IT energy modeling

1) Power consumption of a server: We express the capacity of a server using floating-point operations per second (FLOPS). A server's power consumption is accurately estimated by Eq. 4 given its current load ϕ_{server} expressed in FLOPS, its maximum processing capacity z_{server} (also expressed in FLOPS), the power in idle state P_{idle} and the power at maximum load P_{max} [21].

$$P_{server}(\phi_{server}) = P_{idle} + \frac{P_{max} - P_{idle}}{z_{server}} \times \phi_{server} \quad (4)$$

2) Power consumption of a data center: We formalize the energy consumption of a data center based on a typical state-of-the-art deployment (see Fig. 3). In this model, a data center consists of rows of IT equipment which contain servers, storage devices and other supporting hardware such as coolers, water pumps (to move the cooling water) and uninterruptible power supply (UPS) systems. All power issued to these racks first passes a UPS unit which serves as a battery backup to prevent IT equipment failures in case of power interruptions. Power leaving the UPS enters a power distribution unit (PDU) that sends the power directly to the racks and servers. Note that (i) the electricity consumed by the power delivery chain (PDU+
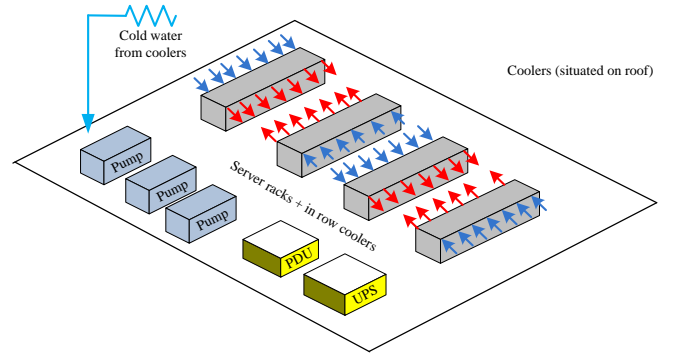

Fig. 3: Energy consuming devices in our data center model.

UPS) accounts for a substantial portion of the overall power consumption of the data center (depending on the technology and load up to half of the total energy consumption) and that (ii) this power delivery chain on top of the pure IT power wastes some energy, which is mainly caused by energy loss at the UPS [26]. Another important factor in a data center regarding power consumption is air flow. The predominant architecture for delivering cooled air is raised floor air delivery from perimeter Computer Room Air Handlers (CRAH). CRAHs are placed around the room and distribute cold air through a raised floor with perforated floor tiles. This kind of architecture suffers from a couple of imperfections: (i) the distance between the cooling units and the heat source makes it difficult to remove the heat without mixing with the supply air and (ii) a considerable amount of energy is needed to drive the fans [6]. To overcome this, we consider an air-circulating solution that addresses these problems, called in-row or rack-based cooling. In this approach, the air cooling systems are integrated into a rack; it makes the air paths shorter, and significantly reduces the power required to operate the fans [18]. We model the power consumption of such an in-row cooler, given

the current capacity of all the rack's servers , the same way as a server; linearly interpolated between a P_{min}^{inrow} and P_{max}^{inrow}, as in Eq. 5.

$$P_{rack} = P_{min}^{inrow} + \frac{P_{max}^{inrow} - P_{min}^{inrow}}{\sum_{server \in rack} z_{server}} \times \sum_{server \in rack} \phi_{server} \tag{5}$$

Apart from air flow, we still need a cooling mechanism. The assumed deployment uses k dry coolers / free coolers which cool the water to about 17-18 C. Finally, the pumps that circulate the cooled water to the racks have to be accounted for. Concluding, the power consumption of our data center prototype is shown in eq. 9 while Table I shows values for these parameters, based on actual readings of the Ghent data center (which serves as our state-of-the-art example, both in technology and in dimensions) or equipment data sheets. Our model allows switching off certain parts of a data center, which gives us freedom in our request scheduling:

- When a server is not in use, we switch it off completely.
- Whenever a rack has no active servers we allow to switch off the in-row coolers
- When no racks are active we allow switching-off the coolers, pumps and UPS system (start up cost for a data center).

$$P_{DC} = P_{base} + \sum_{r \in racks} P_{rack} + \sum_{s \in servers} P_{server}(\phi_s) \tag{6a}$$

$$P_{base} = \begin{cases} 0 \text{ if not in use} \\ P_{UPS} + P_{pumps} + P_{cooler} \text{ otherwise} \end{cases} \tag{6b}$$

IV. PROVISIONING ALGORITHM

We investigate two approaches of scheduling and routing. The first algorithm is based on an integrated scheduling approach, where the destination site and the route towards that destination are found in a single pass, optimizing the network and IT infrastructure utilization simultaneously. We will refer to this approach as Full-Anycast (FA). In a second approach, we first decide where to handle the request and find the route towards that destination subsequently. This means that scheduling a request consists of two separate calculations: in a first step it optimizes the IT infrastructure, followed by the best possible routing given the IT destination. This latter approach (denoted Assisted Anycast or AA) constitutes the state-of-the-art technique in commercial cloud infrastructures. As a last remark, both FA and AA only consider data centers still having enough capacity to fulfill the request. For both FA and AA, when a request has been scheduled to a data center, the data center enforces a First Come First Serve policy (FCFS): it first tries to schedule the requests to the first active server (in an

Symbol	Description	Value
S	Set of source nodes generating requests	20
C	Set of core nodes. These do not generate requests and can be switched off completely.	8
E	Set of bidirectional links.	56
D	Set of OXCs which are connected to a data center.	5
W	Amount of wavelengths per fiber link.	16
n	Number of servers per rack	20
π	Number of racks per data center	45
κ	Number of free coolers	3
P_{max}	Power consumption of a server when at 100% load. [13]	268 W
P_{min}	Power consumption of a server when unused [13]	144 W
P_{idle}^{inrow}	Power consumption of in-row cooler when unused. (AR)	300 W
P_{max}^{inrow}	Power consumption of in-row cooler when all its servers are at 100% load. (AR)	500 W
P_{pumps}	Average power consumption of the pumps for the cooling water (AR)	28500 W
P_{cooler}	Average power consumption of the coolers. (AR)	13000 W
P_{ups}	Average power consumption of UPS. (AR)	12500 W
$P_{transponder}$	O/E/O: Power consumption of a line-side WDM Transponder (10Gbit/s) [24]	35 W
$P_{control}$	Power consumption of a controller [13]	150 W
P_{sf}	Power consumed by the switching fabric. [29]	30 W
$P_{tx/rx}$	E/O,O/E: Power consumed by either a transmitter or a receiver [2]	5.9 W
P_{edfa}	Power consumption for an EDFA. [24]	15 W
$span$	Amplifier span length	80 km

TABLE I: Parameters and power consumption figures for the network and IT resources. References are provided where possible and "AR" (actual reading) indicates that average power was measured on site at the Ghent University data center(Jan. 2012).

active rack) it finds. Only after deciding there are no active servers that can process the request, a new rack is activated with the necessary servers.

A. Full Anycast (FA)

The FA routing algorithm uses a function $P^{FA} : (E \times \mathbb{N}) \to \mathbb{R}$ found in Eq. 7 (expanded in Eq. 7a, Eq. 7b, Eq. 7c and Eq. 7d) for assigning link weights for link l when a request r needs to be scheduled, after which it computes the shortest path based on these weights using Dijkstra's algorithm. We assume $\phi \in \mathbb{N}$ to be the amount of requested IT capacity. Note that in Eq. 7b we add 1 in the sum, to also account for the EDFAs situated in the source and destination OXC of link l. $P_{node}(l)$ returns either the base power for an OXC, or the extra power needed to switch a path. $P_{DC}(l, \phi)$ only works for virtual links, i.e., the graph edges which connect an OXC with a data center. The function $P_{DC}(\phi)$ returns the additional power needed if request r were to be scheduled to data center DC. Assume we have a function $P(DC)$ which returns the current power of data center DC and $P'(DC)$ the power

of the same data center after scheduling request r, then $P_{DC}(l, \phi)$ is given by $P'(DC) - P(DC)$.

$$P^{FA}(l, \phi) = \alpha.P_{link}(l) + \beta.P_{node}(l) + \gamma.P_{DC}(l, \phi) \quad \text{(7a)}$$

$$P_{link}(l) = \begin{cases} 0 \text{ if link } l \text{ is in use} \\ \left(\left\lceil \frac{|l|}{span} \right\rceil + 1 \right).P_{edfa} \text{ otherwise} \end{cases} \quad \text{(7b)}$$

$$P_{node}(l) = \begin{cases} P_{sf} + P_{control} + P_{transponder} \\ \text{ if end of } l \text{ is inactive} \\ P_{transponder} \text{ otherwise} \end{cases} \quad \text{(7c)}$$

$$P_{DC}(l, \phi) = \begin{cases} P_{base}(\phi) + P_{DC}(\phi) \\ \text{ if adjacent DC of } l \text{ is inactive} \\ P_{DC}(\phi) \text{ otherwise} \end{cases} \quad \text{(7d)}$$

We mention that when $\alpha = \beta = \gamma = 1$, the function $P^{FA}(l, \phi)$ attributes each link the extra power it requires if that link (virtual or actual) were to be used to handle request r. By changing the values of α, β and γ, we change the relative importance of power contributions of links, OXCs or data centers, which has been shown to impact the QoS (e.g., blocking [38]). Moreover, by choosing a value different from one for α, β and γ the algorithm is no longer a greedy one and inactivates network resources, although minimizing the infrastructure's energy would activate them. This could potentially be beneficial as the newly activated IT resources (momentarily consuming more power than necessary) can later be reused to service future requests, reducing the temporary energy penalty in the future (see Section V-A2). In our performance evaluation, we will demonstrate a relation between energy consumption and QoS by changing the values for α, β and γ. In this work we will denote a parameter set as $\{\alpha, \beta, \gamma\}$.

B. Assisted Anycast (AA)

As mentioned above, the assisted anycast algorithm consists of two steps. First we select the data center to handle the request after which we find a route to that data center. We investigate four heuristics to select the data center:

- *Closest*: chooses the data center physically closest to the requesting source;
- *L-max*: chooses the data center with the highest current load (concentrating IT requests as much as possible);
- *L-Min*: chooses the data center with the lowest current load (performing IT load balancing);
- *Random*: randomly chooses a data center (as a benchmark strategy).

When assigning link weights to the graph edges, we only use the network-related terms from FA. More specifically we assign weight to the links using $P^{AA} : E \rightarrow \mathbb{R}$ found in Eq. 8.

$$P^{AA}(l) = \alpha.P_{link}(l) + \beta.P_{OXC}(l) \quad \text{(8)}$$

V. PERFORMANCE EVALUATION

We will show results for the simulations performed for the dense EU topology, portrayed in Fig. 1a, with 28 nodes, of which 8 are core nodes and the remaining 20 source nodes. Section V-A presents results assuming communication-intensive requests, while Section V-B will confirm that our conclusions hold for an IT-intensive request scenario. In Section V-C we will present results for the other topologies found in Fig. 1. All source sites $s \in S$ adopt a Poisson process to generate requests, with mean arrival rate λ and mean holding time μ, which accurately fits real world Grid job traces [12]. Consequently the load per source site is expressed in Erlang (λ/μ). Each request corresponds with a single bandwidth unit (i.e., one wavelength) and a fixed amount of IT capacity (which correspond to a number of servers) which needs to be provisioned at a single data center. The dense topology contains 57 bidirectional fiber links, with each link supporting $W = 16$ wavelengths, apart from the links between an OXC and a data center node. Such OXC-DC links are assumed to have $W = 32$ wavelengths. The link lengths correspond to the actual distance between adjacent vertices (cities). Each data center is equipped with 20 racks, each containing 45 servers. We have performed 20 simulations (with a certain warm up period) with different seeds for every load and averaged the results; where possible the graphs show error bars, indicating the 95% confidence interval. We stopped the simulation after having served 200.000 requests. We have used a custom-built simulator [14], developed in the context of the GEYSERS [19] project.

Simulations are initially performed for a scenario where network connectivity is important (we require 3.3 servers per request) and named this the *network-intensive scenario* and later we perform the same set of simulations with identical seeds where we increase the requested IT capacity per bandwidth unit to 8.3 servers per request, which we denote as the *computing-intensive scenario*. We start with a thorough analysis of the FA algorithm, which we compare to AA in Section V-A4. The parameters α, β and γ have been ranged between 0.001 and 1 of which we show results for the most important parameter sets.

A. Network-intensive scenario (FA/Dense topology)

1) Pure IT vs. pure network optimization: In order to compare savings made by parameter sets which either emphasize network or IT power minimization, we illustrate in Fig. 4 the total power consumption for (i) the parameter choice with a high focus on network optimization $\{1, 1, 0.001\}$ denoted as *Net. Opt.* and for (ii) the parameter set with a large focus on IT optimization $\{0.001, 0.001, 1\}$ (*IT Opt.*) We also mention the percentage that network and IT resources contribute to the total energy budget (depicted as the numbers on the corresponding bars), to demonstrate the balance
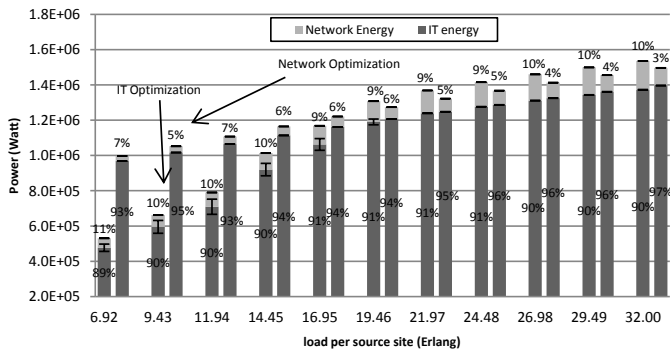
Fig. 4: Power consumption (divided into consumption for network and IT resources) for two parameter sets: first bar *IT Opt.* and second bar *Net. Opt.* The numbers on the bars indicate the contribution of either network or IT resources to the total energy budget. Up until 16.95 Erlang, *IT Opt.* is best, after which *Net. Opt.* has lower values.
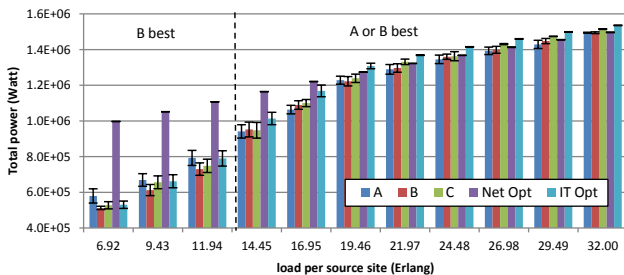

Fig. 5: Total power consumption for parameter sets A, B,C, *IT. Opt.* and *IT Opt.* Up until 11.94 Erlang B minimizes total power consumption, after which A and B attain about the same values.

between network and IT. Fig. 4 shows (i) that *IT Opt.* leads to minimal energy consumption in low load conditions while *Net. Opt.* achieves this in high load conditions and (ii) that minimizing network energy leads to an increase in IT energy and vice versa.

The large variations in total energy in low load situations (a difference up to 48%) mainly stem from switching on all data centers to optimize network power consumption for the *Net. Opt.* scenario, while fewer active data centers could serve all requests. However, starting from 19 Erlang this situation changes and *Net. Opt.* achieves a total power reduction compared to *IT Opt.* of about 3%. In these cases all data centers have to be switched on and the reduction of IT power for *IT Opt.* (on average 15.1 Watt lower IT power consumption than *Net. Opt.*) is too small for the network power energy savings achieved by *Net. Opt.* (on average difference of 58.2 Watt more savings in network energy than *IT. Opt.*). In what follows we will investigate how the values for α, β and γ can be chosen to lower overall power consumption even further.

2) Parameter set minimizing total energy consumption: Our goal is to find the parameters leading to minimal energy consumption, while keeping an acceptable level of service blocking. The first parameter set we investigate is $\{1,1,1\}$, which we will denote as C. In practice, C is a greedy algorithm which chooses the best routing and scheduling achievable at the moment of calculation. In our simulations we have performed a parameter sweep for the values for α, β and γ where we have chosen all possible combinations out of 1, 1/10, 1/100, 1/1000. Results of those simulations point out two extra important parameter sets: two parameter sets with a less explicit focus on network resources than *Net. Opt.* $\{0.1, 0.1, 0.001\}$ (denoted as A) and $\{0.1, 0.01, 0.001\}$ (denoted as B). Figure 5 shows the total power consumption achieved by those parameter sets while in Table II we show the difference in power consumption for these parameter sets compared to the absolute minimum from the parameter sweep. There are three general conclusions which can be derived from Fig. 5 and Table II: (i) in the low load range $[6.92 - 11.94]$ parameter set B achieves minimal power consumption, while in the other end either A or B is best, (ii) neither C, IT Opt. or Net. Opt. reaches this minimal power consumption and (iii) making efficient use of network resources pays off in high load conditions. In order to explain the difference in power values for each parameter set, we need to look into the ability of switching off resources, for which we refer to Fig. 6, Fig. 7 and Fig. 8 where we have plotted the number of inactive data centers, OXCs and fibers per parameter selection. Fig. 9 shows the average path length each algorithm requires.

B has minimal power consumption in the $[6.96 - 11.94]$ end, as it is more effective in switching-off data centers than A (about half a data center). The reason for this is that A sometimes reaches situations where the contribution of IT power ($\gamma.P_{DC}(l,\phi)$) is minimized to such an extent that the contributions of needed network power ($\alpha.P_{link}(l) + \beta.P_{oxc}(l)$) to reach
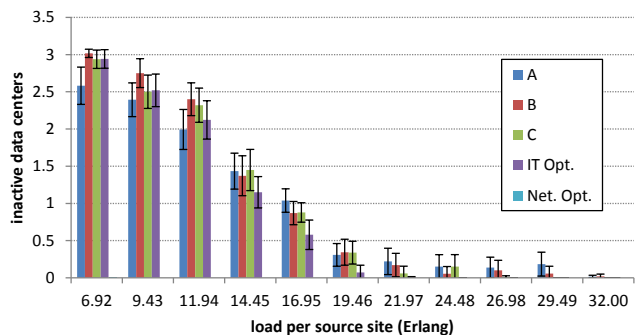

Fig. 6: Number of data centers that are turned off. Parameter sets with a high focus for IT power minimization are clearly best in switching off complete data centers (*Net. Opt.* is unable to switch off data centers).

TABLE II: Difference in total power consumption for the different parameter sets compared to absolute minimum from all simulations over all parameter sets. Gray cell indicate that the corresponding parameter set achieved minimal energy consumption (over all runs).

Offered Load	6.92	9.43	11.94	14.45	16.95	19.46	21.97	24.48	26.98	29.49	32
A	11.60%	8.40%	7.90%	0.00%	0.00%	0.60%	0.00%	1.20%	0.00%	0.00%	0.20%
B	0.00%	0.00%	0.00%	1.10%	2.40%	0.10%	0.60%	2.20%	0.40%	1.30%	0.20%
C	2.80%	6.60%	2.50%	0.70%	3.30%	1.50%	3.20%	2.40%	2.60%	3.00%	1.60%
Net. Opt.	48.60%	41.80%	34.10%	19.10%	12.90%	4.20%	2.50%	2.80%	1.40%	1.80%	0.40%
IT Opt.	3.40%	7.50%	7.60%	7.10%	9.00%	6.70%	5.80%	6.00%	4.60%	4.70%	2.90%
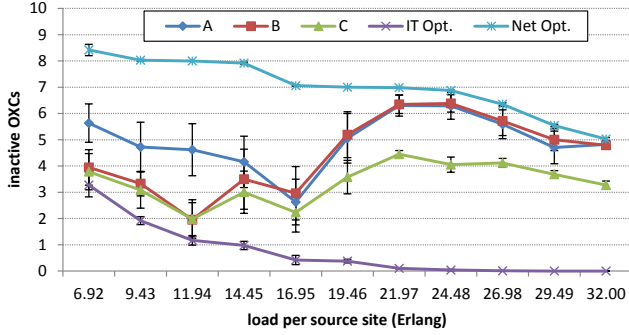
Fig. 7: Number of OXCs which are inactive. Network focused parameters sets are switching off more OXCs. The increase around 19.46 erlang stems from switching on all data centers (see Fig. 6), thus reduces the need for longer paths. Note the ability of A, B and C to turn off OXCs in higher load scenarios: as more and more data centers are turned on, the need to go through the core of the network diminishes and more OXCs can be turned off.

any of the active data centers is too large compared to the adjusted value for the start up cost of an inactive data center. So instead of taking a relatively long path, where additional network resources need activation, the algorithm chooses another path (using already active network resources) and boots up an extra data center. (Note that, since our algorithm does not perform a rescheduling or rerouting step at certain time intervals, this penalty stays during the complete simulation.)
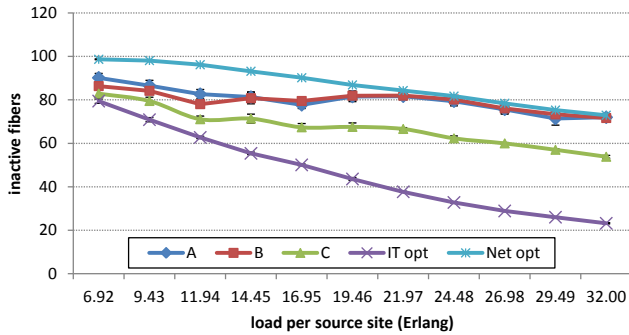
Conversely, in terms of switching off network resources, A is more successful: it is able to switch off on average 2 more OXCs than B in the $[6.96 - 11.94]$ region, as it sometimes has one active data center more than B and hence shorter paths can be used (see Fig. 9). These network savings however do not counter the actual startup cost for the extra data center.

In the right region of the graphs ($[14.45 - 32]$) we see that A and B are able to switch off the same amount of fibers, OXCs and IT resources thus reaching about the same level of energy consumption (given that almost all data centers are active, see Fig. 6). As the heavy startup cost for a data center is not included anymore (only rack/server cost) in the term for IT energy, $P_{DC}(l,\phi)$, the factor $\gamma = 0.001$ minimizes the IT energy contribution to a number which is eight times smaller than the contribution of OXC power ($\beta.P_{oxc}(l)$). As paths constitute multiple hops, making β 10 times smaller (B has a $\beta = 0.01$ compared to A which has $\beta = 0.1$) does not affect the routing much and A and B reach the same routing and scheduling.

When we focus on the greedy algorithm C, Table II indicates that it never reaches the minimal total power consumption, which is also reflected in its ability for switching off resources. The intuitive reason is that C attributes the real incremental power to service a new request, and does not account for the possible reuse of newly activated resources by later requests.

Fig. 8: Number of fiber links which can be switched off. A, B and Net. Opt. are able to switch off significantly more fiber links.
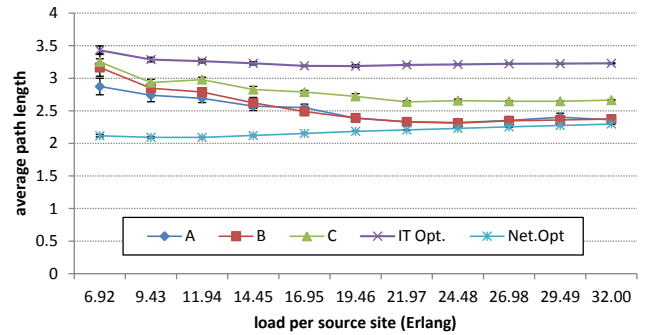

Fig. 9: Average path length per parameter set. IT Opt. produces longer paths, to reach the IT energy minimizing datacenter.

Looking at Fig. 6, in the $[6.96 - 11.94]$ region, B is able to switch off a higher number of data centers. As the contribution of IT power that C accounts for is higher than that for B (or A, for that matter), longer paths are required to avoid activation of a new rack (see Fig. 9). As C thus requires more network resources to reach the data centers, situations occur where for a certain source node there is no (sufficient) free network capacity towards particular data centers, making it necessary to start up another data center to process the request. In the $[14.45 - 32]$ area however, almost all data centers need to be switched on in any case. Yet, for C the accounted contribution of IT power for the algorithm is still large enough (even without data center start-up costs) compared to the network resources ($P_{DC}(l, \phi)$ is about 10 times larger than $\beta.P_{oxc}(l)$ or $\alpha.P_{link}(l)$ for C). Thus, following longer paths is still cheaper with the cost metrics at hand (i.e., IT power minimization is still preferred over network power minimization). Consequently, C is unable to switch off network resources as much as A or B (see Fig. 7), which explains the difference in total power consumption between C and A/B.

Lastly we note that the contribution of link power (i.e., EDFAs) in the algorithm is minimal because (i) whenever a link has already been activated, its contribution (as part of the algorithm) is neutralized ($P_{link}(l) = 0$) as it can be freely used and (ii) the average number of EDFAs per link is five, resulting in an average contribution of only $PUE \times (5 \times 15)$ Watt, which is small compared to the contributions of the OXCs (about 3 times when only one wavelength is routed over the OXC) and the IT resources (about 4 times for 1 rack with one server). We see that *Net. Opt.* is able to switch off significantly more fiber links than the other strategies (up to 48% compared to *IT Opt.*), as EE routing is equivalent to switching-off network resources. In low load conditions, A is able to switch off 4% more fiber links than B. As stated above, B requires this to reach better destinations to keep as much IT resources inactive as possible. Lastly, we find that *IT Opt.* is unable to switch off links as efficiently as the other strategies, as longer paths are needed to reach the best IT site.

3) Influence on QoS: In this section we investigate the influence of parameter options on request blocking in Fig. 10 (due to unavailable network or IT resources), show the average network load in Fig. 12 and mention the data centers load. In the considered *network-intensive scenario*, there is sufficient data center capacity to meet all requests in the considered load scenarios. The only reason for requests not to be provisioned is lack of network resources, i.e., we fail to find a light path to a given server. We see that when optimizing for *IT Opt.*, we have slightly higher blocking, since paths are somewhat longer (see Fig. 9), thus saturating links more (see Fig. 11). Differences among the other strategies are minimal.
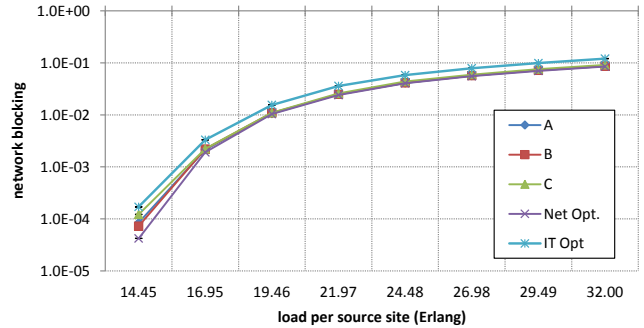

Fig. 10: Network blocking per parameter set. Apart from IT Opt., the A, B, C and Net. Opt. have no significant differences.
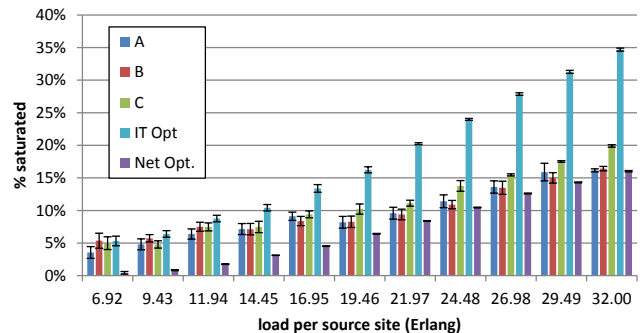

Fig. 11: Percentage of link that has an average load higher than 85%. Parameter sets with a large focus on IT power, have a high saturation value.

We thus find that the strategies leading to the lowest energy consumption (A or B, see higher) are also those with lower blocking. This may sound contradictory to earlier work described in [38], showing a trade-off between energy efficiency and blocking due to network resource fragmentation resulting from long EE paths. However, this work is different in several ways. We consider a network with wavelength conversion, whereas they assume the wavelength continuity constraint. Hence, the effect of resource fragmentation on
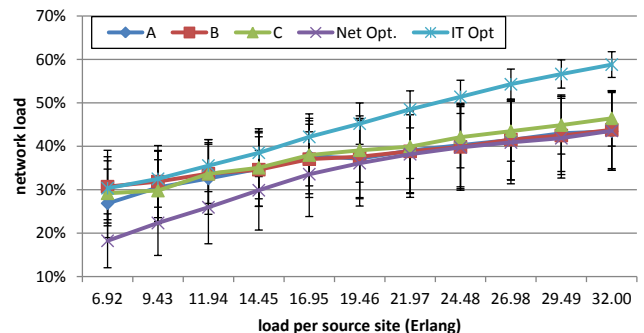

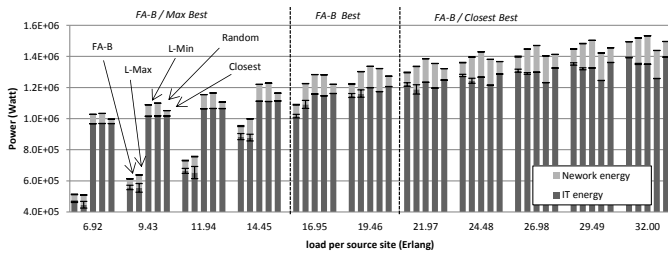Fig. 12: Average network load. A, B C, *Net. Opt.* and *IT Opt.*

Fig. 13: Distribution of power (network and IT energy) comparing FA (parameter set B) with AA
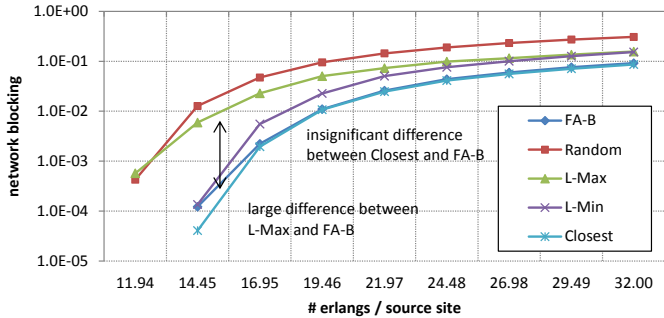

Fig. 14: Network blocking figures comparing FA (parameter set B) with AA for different scheduling algorithms. FA-B and Closest attain about the same blocking value. *L-max* and *Random* reach unacceptable blocking figures.

blocking in our use case is not present, as blocking only occurs when there is no capacity left anymore. Secondly, they assume a random traffic pattern, where each node of the network is a possible destination and lastly they are not switching off nodes (i.e., transponders, switching fabric, etc.) but only the optical links (EDFAs).

4) Difference between FA and AA (Dense topology): In this section we study whether we can achieve the same results as the FA algorithm with a simple AA approach. In Fig. 13 we show the total power consumption for the AA scheduling algorithms, with parameter settings ($\alpha = \beta = 1$). Based on results not detailed here (because of space constraints), we have concluded that the total power consumption for AA is hardly influenced by either α or β. The reason for this is that independently choosing the IT site, forces the algorithm to use OXCs (most dominant network resources), although another destination choice could leave the considered network resource inactive. In Fig. 13 we plot the total power consumption for the AA greedy approaches ($\alpha = \beta = 1$) together with the FA values for parameter set B (FA-B). We compare the corresponding blocking probability in Fig. 14.

Looking at the power consumption in Fig. 13, we find as expected that FA-B performs best (with the notable exception of the highest loads; see our comment at the end of this subsection). Nevertheless, some AA approaches do come very close, but the exact one

depends on the load region. For low loads (until 14.45 Erlang in our case study at hand), the *L-max* strategy seems the best AA approach (and only 3% above FA-B), while for higher loads *Closest* is to be preferred. The fact that *L-max* seems best at low loads is intuitively clear: it is possible to aggregate requests in a limited number of data centers (which is what *L-max* aims for) and turning off the rest. Yet, at these low loads, *L-max* leads to significantly higher blocking ratios (see Fig. 14) than any other AA strategy or FA-B. For higher loads (21.97 erlang and above), intuition also expects *Closest* to be best, since there all data centers need to be powered on, and selecting the nearest data center minimizes network resource usage. Network blocking for *Closest* is also similar to that of FA-B, thus making it a valid (and less complex from an implementation point of view) alternative. Only at mid loads (16.95-19.46 erlang), none of the AA approaches consumes as few power as FA-B. In conclusion, to have a single approach that attains lowest power consumption under all load conditions, none of the AA alternatives does the job, and we should resort to the FA approach.

As a final note, we mentioned that Fig. 13 suggests that *Random* attains the lowest total power consumption for the highest considered loads (starting from 26.98 Erlang). Yet, Fig. 14 shows that *Random* has a very high blocking ratio and consequently the apparent power decrease does not stem from intelligent scheduling/routing, but merely because requests are blocked and we get lower data center/network utilization.

B. Computing-intensive scenario (Dense topology)

When we increase the desired number of servers per request, we change our scenario from one where requests resemble network intensive applications (e.g., video streaming services) to applications where computation is more important. We have also simulated such a use case (for the same FA strategies with parameter settings A, B, C, *IT Opt.* and *Net Opt.*) and have reached the same qualitative conclusions as for the *network-intensive scenario*:

- There is no "universally best" option amongst IT-only or network-only optimization.
- By considering network and IT resources together, we achieve minimal energy consumption of the complete infrastructure.
- This energy reduction does not come with a service blocking penalty.
- True optimum can be reached using FA; it is not possible to reach the same optimum with a simple AA heuristic.

The difference between maximum and minimum power consumption amounts to 38%, which is 10% less than the previous use case (more power intensive IT resources need to be activated). The preference
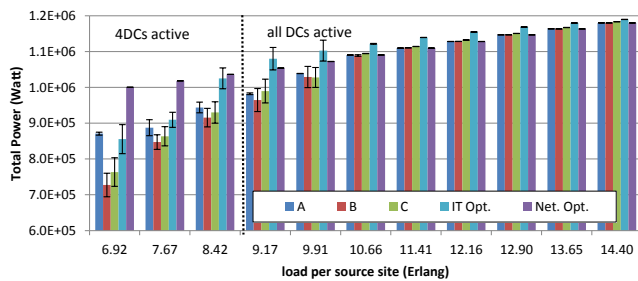
Fig. 15: Power values for the basic network.

Fig. 16: Network blocking for the basic network.

Fig. 17: Power values for the sparse network.

for using either parameter set A or B, depending on the load remains: in low load conditions B is still preferred, but reaches in higher load conditions the same optimal value as parameter set A. The ability of all parameter sets to switch off network resources does not disappear, but is merely shifted to lower load conditions: from a certain point all network resources need to be activated in order to reach certain IT end points.

Although the relation for service blocking between parameters sets stays unchanged, they differ in values. In high load conditions there is not enough IT capacity left to process a request and IT blocking occurs. Although IT blocking for the *IT Opt.* parameter set is lower than for the other strategies (there is only an insignificant difference in IT blocking among A, B and C), network blocking for *IT Opt.* is prevailing, rendering the IT blocking penalty for A, B or C still small enough to outperform *IT Opt.* where total service blocking is concerned. This is also reflected in the network load, which slightly differs between *IT Opt.* and the other FA cases, leading to the difference in blocking.

C. Influence of topology

We demonstrate that our conclusions for the dense topology also hold for the sparse and basic EU topology (with one major difference for the sparse topology when comparing FA and AA). Using the A, B, C, *Net. Opt.*, *IT Opt.* and the different AA algorithms, we have again performed 20 simulations and averaged the results for the basic and sparse EU topologies. The number of requested servers per request is 3.3. (We also performed these simulations for a requested 8.3 servers per request and found that the same qualitative conclusions hold.)

1) Basic topology: The difference between the basic and the dense topology is the number of fiber links (40 vs. 57). There is one major consequence with respect to energy minimization: the number of possible paths between source and destination pairs is smaller for the basic topology compared to the dense topology. This means fewer opportunities for choosing a route between one of the source $s \in S$ and one of the destination nodes $d \in D$. This results in (i) fewer
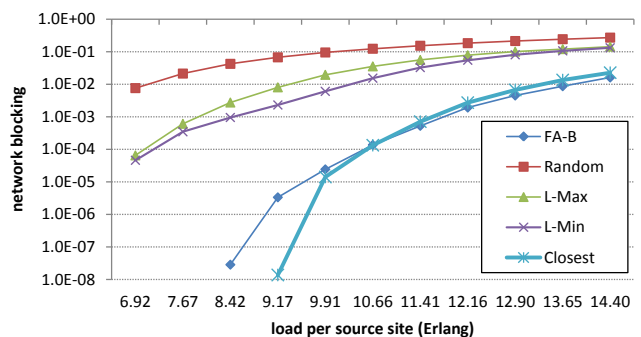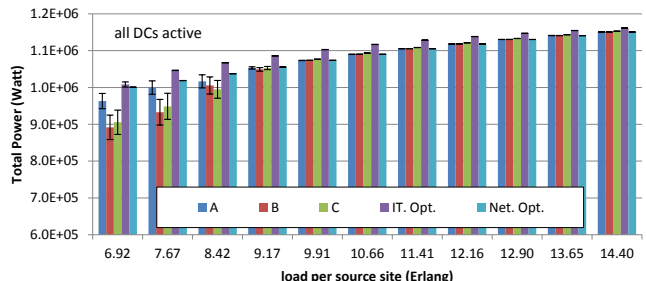
opportunities for switching off network resources and (ii) fewer opportunities for switching off data centers as there is less network capacity. This is also reflected in Fig. 15 where we plot the total power consumption for the different strategies (with an adjusted load per source (λ/μ) site as we keep the number of wavelengths per link the same as for the dense topology). We conclude that all qualitative results for the dense topology also apply for the basic topology. The difference between *IT Opt.* and *Net. Opt.* is considerably lower (up to 14% compared to 48% for the sparse topology) and we see that *Net. Opt.* very quickly reaches minimal energy consumption (starting from 9.17 Erlang): all data centers need to be switched on to overcome network blocking as important links get saturated. The relative difference between A, B and C stays unchanged compared to the dense topology: in the $[6.92 - 10.66]$ region B is the best parameter choice while in the other end A, B, C, *Net. Opt.* reach almost the same optimal power consumption figures. We note that there is no significant difference for the network blocking figures for A, B, C and *Net. Opt.* and that *IT Opt.* has blocking figures which differ from the other parameter set in the orders of several magnitudes. The conclusion regarding the comparison between FA-B and AA also still applies: (i) in the low load scenario AA L-Max approximates the FA algorithm in terms of power consumption, but with a network blocking penalty and (ii) in high load conditions FA-B has similar energy values and service blocking figures as *Closest* scheduling.
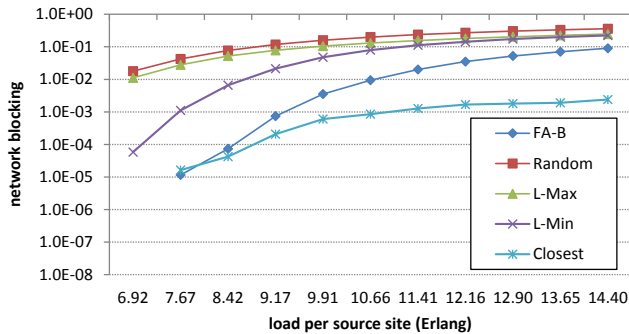
Fig. 18: Network blocking for the sparse network.

2) Sparse topology: The number of fiber links for the sparse topology is even less that the basic topology (33 vs. 40), thus opportunities for EE routing and scheduling are even more limited. Focusing on total power consumption (Fig. 17) we see that even in low load scenarios, *IT Opt.* is outperformed by the other strategies as all data centers need to be switched on to overcome network blocking. The relative difference between A, B, C and *Net. Opt.* is similar as for the basic and dense topology, with a preference for B in the low load scenarios. The relation between AA *L-max* and FA-B is also unchanged: AA *L-max* approximates FA power consumption in a low load scenario, with a service blocking penalty. In high load scenarios however, the service blocking figures for FA and *Closest* are different, although reaching the same optimal energy values. Trying to route with a power minimization objective leads to longer paths in a sparse topology. These longer paths consume precious network capacity, leading to a larger service blocking, while the power optimization seems to have no effect (compared to choosing the closest data center). The reason for the latter, is that EE routing of a single newly arriving request temporarily allows to provision it without activating new resources, but the advantage is lost quite soon when subsequent requests still require to activate new (scarce) network resources. The latter effect seems not to play in less network constrained conditions (i.e., the basic and dense topologies).

VI. Conclusions and future directions

Energy reduction in optical networks received a considerable amount of attention in the research community. In this work, we have ported a number of ideas presented in previous works to an optical cloud context. More specifically, we have presented a unified, online and weighted routing and scheduling algorithm for a typical optical cloud infrastructure for which we have developed an energy consumption model jointly considering network and IT resources.

We can summarize our findings as follows:

- There is no "universally best" option amongst IT-only or network-only energy optimization. Only

by considering network and IT resources jointly, we are able to reach the infrastructure's minimal energy consumption.

- This energy reduction does not lead to a larger service blocking (apart for highly loaded sparse networks).
- Minimal energy consumption can only be reached using a unified "Full Anycast" approach; it is not possible to reach the same optimum with a simple two step heuristic ("Assisted Anycast") which first considers IT resources after which routing is performed, in particular for low to medium load conditions.

Possible extensions and investigations can be devised. Our scheduling algorithm only considers data center selection after which a first server selection strategy is performed over all servers and racks. Consequently, adapting the algorithm with different in-data center scheduling algorithms could lower total energy consumption even further. Another direction for future work is enforcing the wavelength continuity constraint, relieving the need for OEO conversion at OXCs (consequently lowering network energy as transponders are not necessary) and investigating different wavelength selection algorithms. Lastly, resiliency could be explored: how can we protect the integrated network and IT infrastructure, providing resiliency for both network and IT resources, by allowing sharing inactive protection resources (links, OXCs and servers/racks).

References

[1] Dennis Abts, Mike Marty, Philip Wells, Peter Klausler, and Hong Liu. Energy proportional datacenter networks. In *Proc. Int. Symp. on Computer Architecture (ISCA 2010)*, pages 338–347, Saint-Malo, France, 19–23 Jun. 2010.

[2] Slaviša Aleksić. Analysis of power consumption in future high-capacity network nodes. *IEEE/OSA J. Optical Comm. and Netw.*, 1(3):245–258, Aug. 2009.

[3] Paul Anderson, Gaynor Backhouse, Daniel Curtis, Simon Redding, and David Wallom. Low carbon computing a view to 2050 and beyond. Technical report, Technol. & Standards Watch, Bristol, Nov. 2009.

[4] Jayant Baliga, Robert Ayre, Kerry Hinton, Wayne V. Sorin, and Rodny S. Tucker. Energy consumption in optical IP networks. *IEEE J. Lightwave Technol.*, 27(3):2391–2403, Jul. 2009.

[5] B.G. Bathula and J.M.H. Elmirghani. Green networks: Energy efficient design for optical networks. In *Proc. Conf. on Wireless and Optical Commun. Networks. (WOCN)*, pages 1–5, Cairo, Egypt, Apr. 2009.

[6] John Bean and Kevin Dunlap. Energy efficient cooling for data centers: A close-coupled row solution. *ASHRAE J.*, 1, Oct. 2008.

[7] Andreas Berl, Erol Gelenbe, Marco Di Girolamo, Giovanni Giuliani, Hermann De Meer, Minh Quan Dang, and Kostas Pentikousis. Energy-efficient cloud computing. *The Computer J.*, 57(7):1045–1051, August 2010.

[8] Jens Buysse, Kostas Georgakilas, Anna Tzanakaki, Marc De Leenheer, Bart Dhoedt, Piet Demeester, and Chris Develder. Calculating the minimum bounds of energy consumption for cloud networks. In *Proc. IEEE Int. Conf. on Computer Commun. Networks (ICCCN)*, pages 1–7, Maui, Hawaii, USA, 31 Jul. - 4 Aug. 2011.

[9] J. Chabarek, J. Sommers, P. Barford, C Estan, D. Tsiang, and S. Wright. Power awareness in network design and routing. In *Proc. 27th Conf. on Computer Commun. (INFOCOM 2008)*, pages 457–465, Phoenix, AZ, U.S., 15–17April 2008.

[10] Yiyu Chen, Amitayu Das, Wubi Qin, Anand Sivasubramaniam, Qian Wang, and Natarajan Gautam. Managing server energy and operational costs in hosting centers. In *Proc. ACM SIGMETRICS Int. Conf. on Measurement and Modeling of Comp. Sys.*, pages 303–314, Banff, Canada, 6–10 Jun. 2005.

[11] Luca Chiaraviglio, Marco Mellia, and Fabio Neri. Energy-aware backbone networks: a case study. In *Proc. IEEE Int. Conf. on Commun. Workshops*, pages 1–5, Dresden, Germany, Jun. 2009.

[12] Konstantinos Christodoulopoulos, Emmanouel Varvarigos, Chris Develder, Marc De Leenheer, and Bart Dhoedt. Job demand models for optical grid research. In *Proc. 11th Int. IFIP TC Conf. on Optical Network design and modeling (ONDM 2007)*, pages 127–136, Athens, Greece, 29–31 May 2007.

[13] Standard Performance Evaluation Corporation. Specpower. Technical report, SPEC, http://www.spec.org/power_ssj2008/results/, 2008.

[14] Marc De Leenheer, Jens Buysse, Kevin Mets, Bart Dhoedt, and Chris Develder. Design and implementation of a simulation environment for network virtualization. In *Proc. 16th IEEE Int. Workshop Comp. Aided Modeling, Analysis and Design of Commun. Links and Netw. (CAMAD 2011)*, pages 87–91, Kyoto, Japan, Jun. 2011.

[15] Sophie De Maesschalck, Didier Colle, Ilse Lievens, Mario Pickavet, Piet Demeester, Christian Mauz, Monika Jaeger, Robert Inkret, Branko Mikac, and Jan Derkacz. Pan-european optical transport networks: An availability-based comparison. *Photonic Network Commununications*, 5(3):203–225, May 2003.

[16] Chris Develder, Marc De Leenheer, Bart Dhoedt, Mario Pickavet, Didier Colle, Filip De Turck, and Piet Demeester. Optical networks for grid and cloud computing applications. *Proc. IEEE*, 100(5):1149–1167, May 2012.

[17] Chris Develder, Mario Pickavet, Bart Dhoedt, and Piet Demeester. A power-saving strategy for grids. In *Proc. 2nd Int. Conf. on Networks for Grid Applications (GRIDNETS)*, Bejing, China, Oct. 2008.

[18] Keving Dunlap and Neil Ramussen. The advantages of row and rack-oriented cooling architectures for data centers. In *American Power Conversion White Paper*, 2006.

[19] E. Escalona, Shuping Peng, R. Nejabati, D. Simeonidou, J.A. Garcia-Espin, J. Ferrer, S. Figuerola, G. Landi, N. Ciulli, J. Jimenez, B. Belter, Y. Demechenko, C. de Laat, Xiaomin Chen, A. Yukan, S. Soudan, P. Vicat-Blanc, J. Buysse, M. De Leenheer, C. Develder, A. Tzanakaki, P. Robinson, M. Brogle, and T.M. Bohnert. Geysers: Generalised architecture for dynamic infrastructure services. In *GEYSERS: A novel architecture for virtualization and co-provisioning of dynamic optical networks and IT services*, pages 1–8, Warsaw, Poland, Jun. 2011.

[20] Chinwe Esther Abosi. *Towards a service oriented framework for the future optical internet*. PhD thesis, School of Computer Science and Electronic Engineering University of Essex, Apr. 2011.

[21] Xiaobo Fan, Wolf-Dietrich Weber, and Luiz Andre Barroso. Power provisioning for a warehouse-sized computer. In *Proc. 34th Int. Symp. of Computer architecture, SCA*, pages 13–23, San Diego, CA, USA, 7–9 Jun. 2007.

[22] Ian Foster and Carl Kellelman. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, 1999.

[23] M.M. Hasan, F. Farahmand, and J.P. Jue. Energy-awareness in dynamic traffic grooming. In *Proc. Conf. on Optical Fiber Commun. (OFC)*, pages 1–3, San Diego, CA, USA, Mar. 2010.

[24] Filip Idzikowski. Power consumption of network elements in IP over WDM networks. Technical report, Telecommunication Networks Group (TKN), TU Berlin, 2009.

[25] Amornrat Jirattigalachote, Cicek Cavdar, Paolo Monti, Lena Wosinska, and Anna Tzanakaki. Dynamic provisioning strategies for energy efficient WDM networks with dedicated path protection. *Optical Switching and Netw.*, 8(3):201–213, Mar. 2011.

[26] Alan Katz. Maximizing energy cost savings using high efficiency UPS. In *The electricity forum*, 2006.

[27] Dzmitry Kliazovich, Pascal bouvry, and Samee Ullah Khan. DENS: Data center energy-efficient network-aware scheduling. In *Proc. IEEE/ACM Int. Conf. on Green Computing and Commun. & Int. Conf. on Cyber, Physical and Social Computing*, pages 69–75, Washington, DC, USA, Dec. 2010.

[28] G. Jonathan Koomy. Estimating total power consumption by servers in the U.S. and the world. In *Lawrence Berkeley National Labs*, 2007.

[29] Makato Murakami and Oda Kazuhiro. Power consumption analysis of optical cross-connect equipment for future large capacity optical networks. In *Proc. 11th Int. Conf. on Transparent Optical Networks ICTON 2009.*, pages 1–4, Sao Miguel, Azores, Portugal, 28 Jun. – 2 Jul.. 2009.

[30] Francesco Musumeci, Massimo Tornatore, and Achille Pattavina. A power consumption analysis for IP-over-WDM core network architectures. *IEEE/OSA J. Optical Comm. and Netw.*, 4(2):108–117, Feb. 2012.

[31] A. Olkhovets, P. Phanaphat, C. Nuzman, D.J. Shin, C. Lichtenwalner, M. Kozhevnikov, and J. Kim. Performance of an optical switch based on 3-D MEMS crossconnect. *IEEE Photonics Technol. Letters*, 16(3):780–782, Mar. 2004.

[32] A.-C. Orgerie, L. Lefevre, and J.-P. Gelas. Save watts in your grid: Green strategies for energy-aware framework in large scale distributed sys. In *Proc. 14th IEEE Int. Conf. Parallel and Distributed Sys.(ICPADS 2008)*, pages 171–178, Washington, DC, USA, 10 Dec. 2008.

[33] Mario Pickavet, Willem Vereecken, Sophie. Demeyer, Pieter. Audenaert, Brech Vermeulen, Chris Develder, Didier Colle, Bart Dhoedt, and Piet Demeester. Worldwide energy needs for ICT- the rise of power-aware networking. In *Proc. 2nd Int. Symp. Advanced Netw. and Telecommun. Syst.. ANTS*, pages 1–3, Mumbai, India, 15–17 Dec. 2008.

[34] B. Puype, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester. Power reduction techniques in multilayer traffic engineering. In *Proc. 11th Int. Conf. on Transparent Optical Networks*, pages 1–4, Sao Miguel, Azores, Portugal, 28 Jun. – 2 Jul. 2009.

[35] Philip Robinson, Alexandru-Florian Antonescu, Luis M. Contreras-Murillo, Jos Aznar, Sbastien Soudan, Fabienne Anhalt, and Joan A. Garca-Espn. Towards cross stratum SLA management with the GEYSERS architecture. In *Proc. IEEE Int. Symp. Parallel and Distributed Processing with Applications (ISPA)*, Legans, Madrid, Spain, 10–13 Jul. 2012.

[36] Gangxiang Shen and Rodney S. Tucker. Energy-minimized design for IP over WDM networks. *IEEE/OSA J. Optical Commun. and Netw.*, 1(1):176–186, Jun. 2009.

[37] R.S. Tucker, R. Parthiban, J. Baliga, K. Hinton, R.W.A. Ayre, and W.V. Sorin. Evolution of WDM optical IP networks: A cost and energy perspective. *IEEE J. Lightwave Technol.*, 27(3):243–252, Feb 2009.

[38] P. Wiatr, P. Monti, and L. Wosinska. Green lightpath provisioning in transparent WDM networks: Pros and cons. In *Proc. 4th Int. Symp. on Advanced Networks and Telecommunications Sys. (ANTS 2010)*, pages 10–12, Mumbai, India, Dec. 2010.

[39] Yi Zhang, P. Chowdhury, M. Tornatore, and B. Mukherjee. Energy efficiency in telecom optical networks. *IEEE Commun. Surveys & Tutorials*, 12(4):441–458, Jul. 2010.