A Comparative Performance Study of Load Adaptive Energy Saving Schemes for IP-Over-WDM Networks

Marcel Caria, Mohit Chamania, and Admela Jukan

***Abstract*—Load adaptive energy saving schemes for backbone IP networks use dynamic transport circuit services to adapt the active network resources to the current traffic demand in order to reduce the network's energy consumption. Recently, several approaches, categorized as *Switch-Off* schemes, have been proposed which attempt to reduce the energy consumption of already existing networks by *switching off* IP ports and links during periods of low traffic. Although it has been shown that these schemes can notably decrease the network's energy consumption, they are prone to instabilities in the IP routing service and decreased resilience due to reduced connectivity, and they may induce monitoring reconfigurations. To address these challenges, we propose the *Switch-On* scheme in an IP-over-WDM network, where the network is designed so that the essential IP connectivity is maintained during low traffic periods while dynamic circuits are *switched on* in the optical layer to boost network capacity during periods of high traffic demand. Switching on the optical links during peak network loads can address some of the challenges associated with switching off IP ports and links during the low traffic periods. In this paper, we provide a comparative analysis of load adaptive energy saving schemes and present a discussion of the trade-off between energy efficiency and routing stability. The performance results and analytical study show that the multilayer approaches in IP-over-WDM networks carry significant potential for improvement in energy efficiency.**

***Index Terms*—Energy efficiency; IP-over-WDM; Load adaptive; Network optimization; Routing stability.**

I. Introduction

Global IP traffic is predicted to increase 4.3 times from 2009 to 2014, reaching 63.9 exabytes per month in 2014 [1]. This growth rate has not only driven up demand for bandwidth from the Internet backbone infrastructure but has also presented providers with new challenges, especially the energy consumed by the network hardware. It is assumed that the Internet is responsible for about 1% of the electricity consumption in broadband enabled countries, with a tendency to rise [2], with 2% to 3% of telecom companies' operational expenditures (OPEX) already spent on energy [3]. The environmental aspect is critical as the information and communication technologies (ICT) sector is considered a major contributor to global warming, producing about 2% of global carbon dioxide emissions [4]. Although the semiconductor industry is still continuing to follow Moore's law so that the ongoing technological progress of network equipment is assumed to reduce the consumed energy per transmitted bit by around 13% per year [5], this progress is far from keeping up with Internet traffic growth. Without new and efficient energy conservation approaches, energy consumption is likely to become the main bottleneck for the growth of the Internet. Significant efforts are therefore underway to address the energy conservation in IP backbone networks.

Recently, several load adaptive energy saving approaches for backbone IP networks have been proposed which *switch off* IP ports and links during periods of low traffic demand in order to reduce the network's energy consumption. This method, which we hereafter refer to as the *Switch-Off* scheme, is illustrated in Fig. 1(a). As can be seen, the Switch-Off scheme takes advantage of the fact that traffic load in backbone networks follows a stable daily pattern with a significantly lower demand during the night. In this way, the network's capacity can be reduced during the night by shutting down idle resources, which in turn saves power on the corresponding interfaces. However, shutting down parts of the IP topology presents several challenges to network operators. First, any switching off of the IP links can cause rerouting of a large number of IP flows, which not only leads to service disruptions but also requires reconfiguration of the provider's tools for operation, administration, and maintenance. Second, and as illustrated in the shaded part of Fig. 1(a), with a decreasing number of links the average hop count increases, which may increase the end-to-end packet delay [6] and lower the network connectivity, and hence its resilience.

To address these challenges, dynamic optical circuits have been proposed as an innovative solution to offload IP traffic in case of traffic surges. This method is energy efficient not only due to the optical transmission and switching systems being the "greenest" of all networking solutions but also for its ability to deploy advanced multilayer networking methods for joint packet and circuit operation. This is illustrated in Fig. 1(b), where we depict the scheme we refer to as *Switch-On*. Unlike the Switch-Off scheme, in which the IP network is designed for peak traffic, here we guarantee the same connectivity for off-peak traffic demands using low capacity links (shaded side of Fig. 1(b)). During periods of high traffic demand, dynamic transport circuits are comissioned in order to increase the network's capacity. As will be seen later,

Manuscript received April 13, 2011; revised November 14, 2011; accepted January 16, 2012; published February 3, 2012 (Doc. ID 145855).

The authors are with the Institute of Computer and Network Engineering, Technische Universität Carolo-Wilhelmina zu Braunschweig, Germany (e-mail: caria@ida.ing.tu-bs.de).

Digital Object Identifier 10.1364/JOCN.4.000001

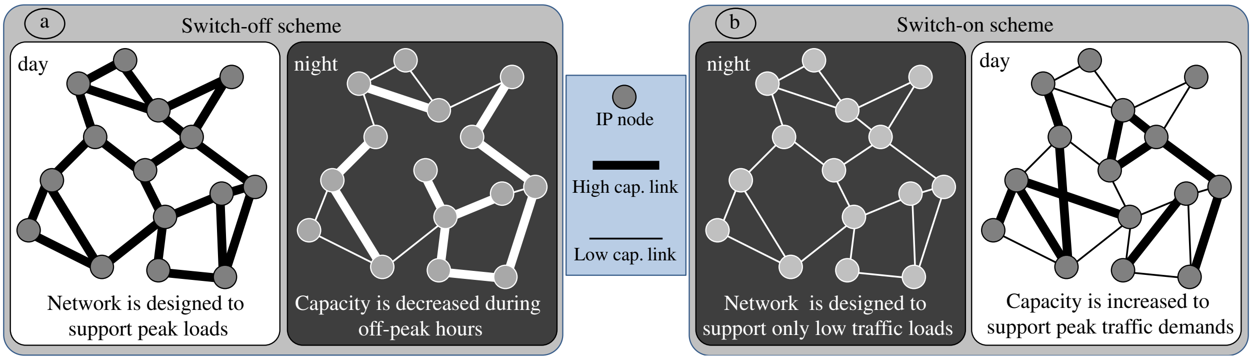

Fig. 1. (Color online) Illustration of the Switch-Off and Switch-On schemes. (a) *Switch-Off*: IP links are designed to support peak loads in the basic topology during the day and some links are switched off or reduced in capacity during off-peak hours in the night. (b) *Switch-On*: IP links are designed for optimum operation only during off-peak hours in the night, while, for high daytime loads, dynamic circuits are added and/or the capacity of existing IP links is increased using dynamic circuits.

however, the Switch-On scheme is not without challenges: the newly established circuits can also destabilize the IP routing, are often inadequate in size due to the established standards on circuit capacity, and require advanced planning and reconfiguration methods.

In this paper, in order to provide fundamental insights into the energy efficiency of the proposed schemes, we provide a comparative analysis of the multilayer Switch-On scheme and the conventional Switch-Off scheme. For the Switch-On scheme, we analyze two classes of circuit setup operations: (a) when the dynamic circuit is advertised in the IP layer routing service as a new link and (b) when the circuit is hidden from the IP layer routing. The hidden circuit is especially beneficial for maintaining the stability of the IP routing service. However, this practice carries challenges since an IP router is required not to advertise a hidden circuit via the routing protocol; at the same time, the router must be configured to reroute traffic onto the same circuit only on a per flow basis. We formulate the Switch-Off scheme and both Switch-On schemes as Integer Linear Programming (ILP) models, with all three models reflecting the same architectural characteristics of the IP-over-WDM network, such as shortest path first (SPF) routing in IP, power consumption of network equipment, and maximum utilization of link capacities. However, they differ in the way the network is designed and operated, which is reflected in the fundamentally opposite traffic scenarios that they are initially optimized for.

We also introduce novel routing constraints for bypass usage such as those which ensure that IP traffic can only be offloaded on an optical bypass which connects a pair of IP routers on the original routing path of an IP flow. Considering the fact that the complexity class of our models is NP-hard, attention must be paid to the size of the network to be optimized with our ILP models. In the performance results, we have compared the different methods based on the following performance metrics: energy saving, routing changes, average hop count per transferred bit, average nodal degree, and interface size. Our results are based on a real 14-node network topology using a reference (but linearly scaled-up) traffic matrix. In order to respect the special case in which a few (but very large) traffic flows dominate in the network, we also studied the effect of elephant flows regarding their number, path length, and time of occurrence. The results have shown that the Switch-On schemes outperform the Switch-Off scheme in terms of energy saving capability, whereas for Switch-On using conventional network engineered links, there exists an adjustable trade-off between the routing stability in the IP layer and the energy savings achieved. The Switch-On schemes have been shown to carry significant potential to save energy, while maintaining routing stability.

The rest of the paper is organized as follows. Section II shows the related work and our contribution to the topic. Section III presents the architecture and explains the use of an optical bypass hidden from the IP routing service. Section IV presents the ILP formulation which is used to determine the optimal Switch-On/Switch-Off operations for the compared schemes. Section V presents a performance study of the three schemes, and Section VI concludes the paper and provides directions for future research.

II. Related Work and Our Contribution

The recently proposed load adaptive energy efficiency schemes are mechanisms to dynamically adapt the network's capacity to the actual traffic demand in order to shut down idle equipment so that electrical power can be saved.

As shown in [7], link loads in backbone networks carry peak loads from eight in the morning till midnight, whereas most network resources are largely underutilized during the night. In [8], a general upper bound for the energy saving potential of load adaptive networks is calculated. It is shown there that those schemes can significantly improve the network's energy efficiency if their capacity is dynamically adjusted to the daily backbone traffic patterns. In [9], the corresponding mathematical problem is referred to as the Minimum Edges Routing Problem, and the authors analyzed the complexity of the same, showing that it is NP-hard. The authors of [10] analyze the relationship between reliability, performance, and power consumption, with the result that using power consumption as the foremost parameter in network planning creates networks with concentrated connections, leading to reliability problems. The performance of IP link Switch-Off schemes was analyzed in [11] and [12], where not only idle links but also certain nodes can be completely switched off, since the authors assume that zero traffic is generated or terminated at some nodes. For the Switch-Off scheme in [13], it is assumed that the network links consist of multiple cables whose capacities are aggregated and that individual cables can be powered down to save energy.

A layer 2 Switch-Off approach is proposed in [14], where an ILP model is presented for power aware routing and wavelength assignment (PA-RWA) with protection. The authors of [15] extend the PA-RWA model in a way that low bandwidth connections can be groomed on the same lightpath, while the authors of [16] extend the PA-RWA model for the *Drop-and-Continue optical switch* architecture, which supports flexible connection grooming. Here, connections can be groomed if they originate in the same source node and terminate at the destination node or any intermediate node of the same lightpath.

Multilayer approaches are given in [17–19], where the total power in the IP and the optical layer is minimized, and [20], where minimization of the number of lightpaths and minimization of electronically processed traffic is combined into a grooming problem, both leading to models of very high complexity which are solvable only for small networks of no practical importance. The authors of [21] extended the multilayer Switch-Off approach with reconfiguration constraints for IP traffic stability and applied it to a daily traffic pattern partitioned into segments, each of two hours' length, so that adaptations follow the actual traffic demand more precisely.

In our preliminary work [22], we were the first to consider the use of hidden bypasses for energy saving, which we refer to here as the Switch-On approach. We compared the latter to the traditional Switch-Off approach, where we showed that bypass-based offloading reduces the total energy consumption and the total installed capacity, as well as the number of reconfigurations required in the network.

In [23], we presented another Switch-On scheme which uses dynamic circuits to establish conventional IP links that are advertised in the IP routing service. In order to reduce the number of IP routing reconfigurations during energy saving operations, we formulated a "punishment" for routing changes in the objective of the optimization process. Again, we compared our Switch-On scheme to the traditional Switch-Off approach, showing that significant energy conservation is possible while taking routing stability into consideration. In this paper, which completes the studies started in [22] and [23], we compare all three models, i.e., the traditional Switch-Off scheme, the Switch-On scheme using hidden bypasses, and the Switch-On scheme using network engineered links, while extending the range of our observations to network performance in terms of average hop count per transferred bit and network resilience measured as the average nodal degree, as well as the effects of traffic pattern anomalies caused by elephant flows.

III. IP-Over-WDM Network Architecture

In this section, we describe the network architectural model and provide definitions for the optical bypass.

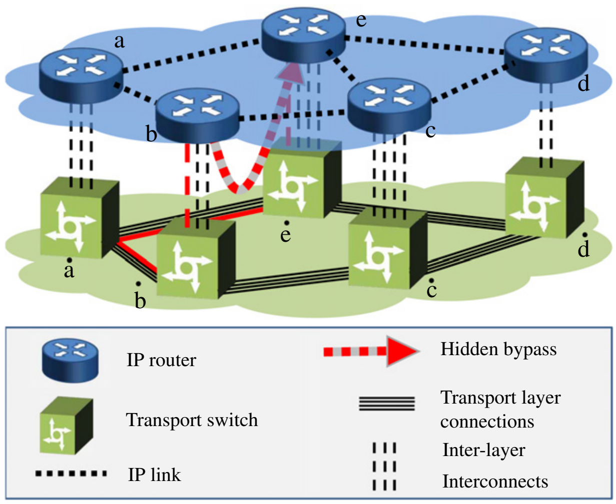

Fig. 2. (Color online) Example for a hidden bypass (highlighted red): The WDM connection via ($\dot{b}$, $\dot{a}$) and ($\dot{a}$, $\dot{e}$) is commissioned in order to offload traffic from b to e.

A. Network Topology

In order to facilitate the use of dynamic circuit services provided by the optical transport layer, we assume that each IP router is co-located with an optical network switch as shown in Fig. 2, which is typically the case in current carrier networks. It is assumed that the optical transport network supports both leased-line as well as dynamic circuit services, where static IP links are established using leased lines. An optical transport network deployed can be opaque or transparent. *Opaque* means that an incoming wavelength is terminated at each optical node, and multi-hop transmission is facilitated via optical–electrical–optical (OEO) conversions, for the purposes of signal regeneration. *Transparent*, on the other hand, means that an optical transport network employs an optical switching technology such as an optical cross-connect (OXC) or a reconfigurable optical add–drop multiplexer (ROADM) which permits WDM connections between physically non-adjacent nodes. In the remainder of the paper, we will simply use the term dynamic optical circuit to refer to either an opaque or transparent optical path, unless otherwise noted.

While, in the traditional Switch-Off scheme, the base topology is designed for high traffic loads, and we therefore know the total installed capacity on IP routers, the same is not true for the Switch-On scheme, where the initial link capacities are designed for low loads, and capacity is added for facilitating high traffic demands. While this can affect the size of the IP router and thereby affect the power consumption, we assume that both models use the same IP router chassis, and our results will later show that the total capacity required in the Switch-On scheme is actually lower than that required for the traditional link Switch-Off scheme. We also assume here that the total capacity available to interface the IP router with its co-located transport switch is very large and is therefore not a bottleneck.

B. Network Engineering Paradigms

Optical bypass is a term associated with the phenomenon in which a dynamic optical circuit (opaque or transparent) is used to interconnect two non-neighboring IP routers. However, we also differentiate optical bypasses based on the way they are advertised in the IP network. In conventional Network Engineering (NE) parlance, an optical bypass created between a pair of IP routers is usually advertised in the IP routing protocol as a new IP link. On the other hand, as proposed in our previous work in [24], a multi-hop circuit setup between a pair of IP routers which is not advertised in the IP routing protocol is referred to as a *hidden bypass*. In contrast to the traditional bypass with NE, in which dynamic circuits established in the transport network are used as regular new IP links, a hidden bypass offloads IP traffic onto the optical layer under the following conditions:

- The bypass must not be advertised as a link in the IP routing service.
- Only the router at the ingress of the bypass is configured to offload specific predefined IP flows onto the bypass, while other routers remain unaffected by the existence of the bypass.
- An IP flow can only be offloaded onto a bypass if the ingress and egress of that bypass lie on the original IP routing path of that flow.

Under these conditions, we minimize the impact of the bypass setup on the IP routing stability compared to the introduction of a conventional network engineered IP link. An example of such a hidden bypass is shown in Fig. 2 between routers b and e. Assuming shortest path routing and that traffic from b to e would originally use the path $(b \to a \to e)$, this flow could be offloaded onto the bypass, since the bypass ingress router b and the bypass egress router e lie in the original routing path. A counter-example is the flow from b to d, which could not be offloaded onto the bypass since the bypass egress router e does not lie in the original routing path $(b \to c \to d)$.

In this paper, we use both types of transport connection for the Switch-On scheme, i.e., the traditional network engineered IP link, as well as the hidden bypass. The possibility to establish hidden bypasses actually depends on the hardware deployed in the operator's network, since an IP router must be configurable in a way that it a) does not advertise certain new links via the routing protocol and b) can reroute traffic on a per flow basis. In order to distinguish the two Switch-On schemes in this paper, we use the abbreviations *Switch-On BP* when we use the hidden bypasses, whereas we say *Switch-On NE* when we use conventional Network Engineering.

C. Port Sizes

In our model, we assume that a *logical* IP link (or hidden bypass) is established by using a single or multiple (aggregated) *physical* IP links. There are no restrictions (in terms of numbers or size) of IP links that can be aggregated, e.g., a 10 Gbit/s and a 40 Gbit/s circuit can be aggregated in order to establish a 50 Gbit/s logical IP link. However, and without the loss of generality, we do not allow the virtualization of an IP port: for example, a single 100 Gbit/s IP port cannot be used to provision two different links. In the link Switch-Off scheme, either all physical links constituting a logical IP link are switched off to shut down the connection completely or some of the physical links are shut down to reduce the capacity of a logical link. On the other hand, in the Switch-On scheme, additional physical links can be activated to boost the capacity of an existing link or to establish a completely new logical IP link or hidden bypass. As mentioned before, we do not constrain the total number of ports installed on an IP router. Whereas our results will show that the *total capacity* installed in the Switch-On scheme is lower, making this a fair assumption, we see in our performance evaluation (Section V) that the optimal solution does require a larger number of small (low capacity) IP ports. Also note that in our formulation we assume a perfect matching between the capacity of the port and the corresponding circuit, which means that a 10 Gbit/s circuit cannot be associated with a 40 Gbit/s IP port.

D. Routing in the IP and the WDM Network

The IP network under consideration uses an SPF protocol variant in our model. Whereas we do not restrict the model to any specific routing protocol, the proposed model does not support multipath routing between a pair of routers (e.g., Equal Cost Multipath Routing in Open Shortest Path First (OSPF)) and assumes that a (single) unique route is used to forward all packets from a source to a destination. In our model, we do not explicitly model capacity constraints in the transport network, and we assume that the transport network has sufficient capacity to route all circuit requests for leased lines and dynamic circuits. In order to perform a fair energy comparison, we also assume that all circuits between a given pair of routers use the *same (precomputed) path in the transport network*. Therefore, the power used by two circuits of the same granularity between the same pair of routers is equal.

E. Traffic Profile

The model that we use in this paper assumes complete knowledge of the IP traffic matrix, including comprehensive knowledge of the daily traffic patterns. To model the variations in traffic patterns over a day, we create distinct time intervals, and for each interval we assume knowledge about the traffic loads between all pairs of routers in the network. To reduce the complexity, we only use two time intervals, namely, the *peak* traffic interval from 8 am till 12 midnight and the *off-peak* traffic interval during the night, assuming that the traffic load during the night is simply a rather small fraction of the peak traffic. However, the model can be extended to more than two intervals per day in order to allow for adaptation of the network's capacity to the actual traffic demand in shorter time intervals. For instance, in [21], the authors applied their Switch-Off model to a daily traffic pattern partitioned into 12 distinct two-hour time intervals and used a sequential top-down order (i.e., from the traffic peak in subsequent steps to the time intervals with lower traffic demand) for Switch-Off computation. Their results show that by splitting the 24 hours of a day into such short intervals, the power consumption can smoothly follow the actual traffic demand. While we have not analyzed the influence of the number of time intervals on the results presented in this paper, we assume that more

TABLE I
SUMMARY OF NOTATION

Parameter	Type	Meaning
$\dot{V} = \{\dot{x}, \dot{y}, \ldots\}$	Set	Set of circuit switches
$\dot{E} = \{(\dot{x}, \dot{y}), (\dot{x}, \dot{z}), \ldots\}$	Set	Set of physical links (fibers) between switches
$V = \{x, y, \ldots\}$	Set	Set of IP routers in layer 3, where router x, co-located with switch $\dot{x}$, etc.
$E = \{(x, y), (x, z), \ldots\}$	Set	Set of IP links in layer 3
M	Integer	Maximum number of ports at any router in the network
$\hat{L}_{xy}$	Boolean	Parameter for existence of link (x, y) during the preceding time interval
$N = \{\text{peak}, \text{off-peak}\}$	Set	Set of time intervals
τ^n	Real	Duration of a certain time interval, $\sum_n \tau^n = 24$ h
T	Set	Set of available bandwidth rates for circuits
C^t	Integer	Capacity of a circuit of a certain bandwidth rate $t \in T$
λ^n_{xy}	Integer	Traffic from x to y in time interval n
$\dot{P}^t_{xy}$	Integer	Power consumption of a circuit from x to y with bandwidth rate $t \in T$
P^t	Integer	Power consumption of a pair of IP ports with bandwidth rate $t \in T$
α	Real	Maximum IP link utilization, $0 \leq \alpha \leq 1$
w_{xy}	Real	IP routing weight
$\sigma \in \mathbb{R}^+$	Real	Weighting factor of the IP rerouting penalty
Variable	**Type**	**Meaning**
$X^{n,t}_{xy}$	Integer	Number of circuits of rate t used in time interval n for IP Link (x, y)
C^n_{xy}	Integer	Capacity of IP Link (x, y) in time interval n
L^n_{xy}	Boolean	Link Existence Variable: *true*, if the IP link (x, y) exists in time interval n
$R^n_{sd}(i, j)$	Boolean	IP Routing Variable: *true* if Path from s to d uses link (i, j) in time interval n
$\Re^n_{sd}(x, y)$	Boolean	Extended routing parameter: *true* if path from s to d uses loose path $s \to x \to y \to d$
P^n_{xy}	Integer	Power consumption of IP link from x to y in time interval n
RC^n_{sd}	Real	IP routing cost from s to d in time interval n
FT^n_{ijd}	Boolean	IP forwarding variable for time interval n: *true* if router i forwards packets for d to j
$Punish^n$	Integer	Penalty for routing changes in time interval n

(i.e., shorter) time intervals would enable any load adaptive energy conservation scheme to save more energy and at the same time would likely lead to a higher number of routing changes per day. We plan to examine this issue in our future work.

F. Power Consumption

For a general and accurate power model, there are too many parameters to consider, e.g., technology, vendor, performance, generation and utilization of the network equipment, applications run in the network, etc., see [25], with many of them not easy to obtain. Therefore, in this paper, we use a simplified power model, which depends on two parameters only: 1) the number and type of active IP ports during a time interval and 2) the power consumption of the circuits provisioned in the transport layer to establish IP links or hidden bypasses. The power consumption of an IP port is a (known) function of the port granularity. Note here that we assume a bandwidth discount, i.e., the power per bit decreases with increase in port sizes, so that the power consumed by a 40 Gbit/s IP port would be lower than the power consumed by four 10 Gbit/s IP ports. As mentioned before, we assume that the WDM network is used to provision IP links, and that, for a given pair of IP routers, all circuits use the same physical path. The power consumed by a circuit is defined as a function of the optical technology used and the number of physical hops in the optical layer. The technology used can be an opaque or a transparent optical network, depending on the regeneration requirement of optical signals. In general, however, our power model is technology agnostic, since the power parameters used can be adapted to reflect conditions in any optical transport network technology, with or without optical signal regeneration. Finally, it is assumed that the IP router chassis used is the same in all cases, so that the chassis power consumption is not included in our model. This simplification is justified by our results, which show that the *total capacity* installed is lower in the Switch-On scheme, ensuring that no unfair advantage is given to it.

IV. Analytical Model

We will now present the analytical model based on ILP used to determine the operation of the Switch-Off and Switch-On schemes. The notation used in this formulation is listed in Table I.

In our formulation, we present a model which can optimize the energy consumption of the network under the constraint that the link capacity can be composed flexibly using the available circuit granularities. In all cases, the model follows some basic constraints such as IP link capacity utilization thresholds and SPF-based routing which were originally presented in [26]. However, the modeling of the link capacity is kept flexible, wherein 1) a new link or an optical bypass hidden from the IP routing service may be added, 2) the capacity of an existing IP link may be modified (increased/decreased), or 3) an existing link may be completely switched off. As stated before, the model also assumes complete knowledge of the peak and the off-peak traffic load, and it uses this information to compute the IP topology and link capacities and consequently the power consumption for different schemes.

Power Consumption: We use the information about the IP links and the circuits used to provision these links to compute the power used by the IP network. The power consumed by an IP link (or a hidden bypass) from x to y during time interval n is given by

$$P_{xy}^{n} = \sum_{t} X_{xy}^{n,t}\left(\dot{P}_{xy}^{t} + P^{t}\right), \tag{1}$$

which, applied to IP link (b,e) in Fig. 2, incorporates the total power used by all the IP ports in b and e as well as the power used by all the circuits (in this case only one, via $(\dot{b}, \dot{a})$ and $(\dot{a}, \dot{e})$) in the WDM network used to provision the IP link. Please note that, due to our simplified power model (see Subsection III.F), we use normalized power values without any unit.

Routing Penalty: To model the effect of routing changes, we define the routing penalty parameter $Punish$ for each traffic profile transition from profile $n-1$ to profile n, which gives an extent of the reconfiguration effort required in the event correlation database, which is proportional to the total number of different links between the old and the new routing paths:

$$Punish^{n} = \sum_{sdij}\left|R_{sd}^{n}(i,j) - R_{sd}^{n-1}(i,j)\right|. \tag{2}$$

If, for instance, in our example network in Fig. 2 the route from b to e changes from $b \to a \to e$ to $b \to e$, we have a rerouting punishment value of 3, since the failure event for links (b,a), (a,e), and (b,e) must be reconfigured in the event correlation database.

Objective Function: The objective function, which minimizes the total energy consumed while accounting for the routing penalty, is defined as

$$\text{Minimize} \sum_{n}\sum_{xy}\left(\sigma \cdot Punish^{n} + P_{xy}^{n}\right), \tag{3}$$

where σ is the routing stability parameter ($\sigma \in \mathbb{R}$, $\sigma \geq 0$) that balances the trade-off between energy consumption and routing penalty. This parameter is set to zero in the Switch-Off scheme, where routing changes are generally not considered. Also, when we use hidden bypasses for the Switch-On scheme, σ is set to zero, since large scale reroutings are avoided here by design. Please note that incorporating routing stability into the objective of the Switch-Off scheme would also be possible, but this was not the focus in our analysis, since we decided to compare both our Switch-On schemes only with the *conventional* Switch-Off scheme (i.e., without considering routing stability). However, we assume that the extension of the Switch-Off model with the routing penalty would lead to a comparable reduction of IP reroutings that we saw in our results for Switch-On, albeit at the cost of performance loss in terms of energy efficiency.

Existence, Capacity, and Utilization of IP Links: For all routers $x,y \in V$, and for both time intervals $n \in \{\text{peak}, \text{off-peak}\}$, the IP link existence variable L_{xy}^{n} is bounded by

$$\left\lceil \frac{\sum_t X_{xy}^{n,t}}{M} \right\rceil \leq L_{xy}^{n} \leq \sum_{t} X_{xy}^{n,t}. \tag{4}$$

Applied to the topology in Fig. 2, the right-hand side of the constraint ensures that there is no link (b,e) if no circuit is commissioned between $\dot{b}$ and $\dot{e}$, whereas the left-hand side of the constraint ensures that the link (b,e) exists, if at least one circuit is commissioned from $\dot{b}$ to $\dot{e}$. The capacity of the IP link (i,j) in a time interval $n \in \{\text{peak}, \text{off-peak}\}$ is given as the sum of all circuits that are commissioned between $\dot{i}$ and $\dot{j}$ during that time interval:

$$C_{ij}^{n} = \sum_{t} X_{i,j}^{n,t} \cdot C^{t}. \tag{5}$$

We constrain the total capacity of an IP link by

$$\sum_{s}\sum_{d} \lambda_{sd}^{n} \cdot R_{sd}^{n}(i,j) \leq \alpha \cdot C_{ij}^{n} \tag{6}$$

to ensure that it supports the actual IP traffic load while limiting the link utilization to a predefined threshold α for $n \in \{\text{peak}, \text{off-peak}\}$ and all $i,j \in V$.

General Routing Constraints: To model the routing constraints in the IP network, we first model the route continuity constraints and then incorporate additional constraints to ensure that the routing conforms to SPF. For all routers $s,d,i,j,k \in V$ (with $k \neq s$ and $k \neq d$) and $n \in \{\text{peak}, \text{off-peak}\}$, the following routing constraints must hold:

$$R_{sd}^{n}(i,j) \leq L_{ij}^{n}, \tag{7}$$

$$\sum_{p} R_{sd}^{n}(s,p) = \sum_{q} R_{sd}^{n}(q,d) = 1, \tag{8}$$

$$\sum_{p} R_{sd}^{n}(p,k) = \sum_{q} R_{sd}^{n}(k,q). \tag{9}$$

Applied to Fig. 2, Eq. (7) ensures that the IP link from a to e can only be used if it exists. Equation (8) ensures that, for a routing path from a to c, exactly one outgoing link from a and exactly one incoming link at c is used, whereas Eq. (9) ensures routing continuity inside the network.

Shortest Path First: IP networks typically employ a variant of the SPF routing algorithm, and constraints are introduced here to ensure that only least cost paths can be used to route IP traffic. For every possible link between any routers x and y, we use a predefined routing weight metric w_{xy}, which is used to evaluate the cost of the routing path. Using this metric and the routing variable, the routing cost RC_{sd}^{n} from s to d ($s \neq d$) during the time interval n is given by

$$RC_{sd}^{n} = \sum_{x}\sum_{y} R_{sd}^{n}(x,y) \cdot w_{xy}. \tag{10}$$

We ensure minimum routing cost from s to d for $n \in \{\text{peak}, \text{off-peak}\}$ and $s,d,x \in V$ with $s \neq d$ by

$$RC_{sd}^{n} \leq RC_{sx}^{n} + L_{xd}^{n} \cdot w_{xd} + \left(1 - L_{xd}^{n}\right) \cdot w_{\infty}. \tag{11}$$

The constraint ensures that the routing cost is less than or equal to the routing cost from s to an immediate neighbor x of d plus the link cost from x to d. For example, the links (a,e), (e,d), and (d,c) in Fig. 2 cannot be used to route traffic from a to c since there is the intermediate node e which is already connected to the destination c. Note that, if the regarded intermediate node is not a neighbor of d, a very large weight w_{∞} is introduced to render the constraint useless.

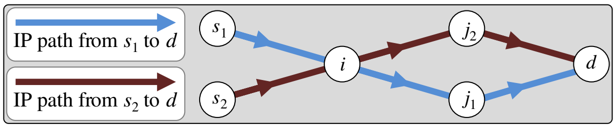

Fig. 3. (Color online) Example of the resulting invalid IP forwarding model, when only the shortest path constraint is considered. Router i forwards packets to the same destination on different outgoing links.

Packet Forwarding: While the routing constraints ensure that only shortest cost paths are used, they do not mimic the operation of an SPF routing algorithm in the IP network, especially when multiple shortest paths can exist between a source and a destination. For instance, both IP paths $s_1 \to i \to j_1 \to d$ and $s_2 \to i \to j_2 \to d$ in Fig. 3 satisfy the shortest path constraint (Eq. (11)). However, they could not exist simultaneously in an IP network. The reason is simple: all packets regardless from which source must follow the same IP routing path to a common destination as soon as they reach a common intermediate node, since that intermediate node has only one entry for that destination in its forwarding table. In the example shown, router i does not satisfy this constraint, since it forwards packets to the same destination d on different outgoing links. To model this *destination-based SPF packet forwarding*, we introduce a Boolean variable FT^n_{ijd} which mimics a forwarding table, with

$$\sum_j FT^n_{ijd} = 1 \tag{12}$$

for all $i,d \in V$ and $n \in \{\text{peak}, \text{off-peak}\}$, ensuring that a router always has exactly one forwarding table entry to a certain destination. By constraining

$$R^n_{sd}(ij) \le FT^n_{ijd} \le L^n_{ij} \tag{13}$$

for all $s,d,i,j \in V$ and $n \in \{\text{peak}, \text{off-peak}\}$, we ensure that the IP routing follows the forwarding table (left-hand side of Eq. (13)) and also that the link to the next hop exists (right-hand side of Eq. (13)). As an example, router e in Fig. 2 must have exactly one entry in its forwarding table for router a so that, regardless from which source router the traffic is routed via e to a, it is always using the same outgoing link from e (Eq. (12)). Furthermore, all routing paths must follow the routing tables (left-hand side of Eq. (13)) and forwarding on a link requires that the link exists (right-hand side of Eq. (13)).

Using Hidden Bypasses for Switch-On: To model the bypassing of flows, a new Boolean variable f^{sd}_{xy} is introduced, which indicates if the traffic from s to d is offloaded to a *bypass* from x to y. It is constrained by

$$\forall s,d,x,y \in V, \quad \hat{L}_{xy} = 0: \; f^{sd}_{xy} \le \sum_t X^t_{xy}, \tag{14}$$

$$\forall s,d,x,y \in V, \quad \hat{L}_{xy} = 0: \; f^{sd}_{xy} \le \Re^n_{sd}(x,y), \tag{15}$$

where Eq. (14) ensures that traffic can only be offloaded on a bypass if at least one circuit is established from x to y, whereas Eq. (15) ensures that the traffic from s to d can only be offloaded on a bypass whose end points lie on the original routing path from s to d, i.e., no IP link will experience any traffic increase by rerouted flows, since a flow can exit its designated IP routing path only to enter a bypass, and it can reenter the IP network from that bypass only at routers that are on its designated IP routing path. In this model, the number of bypasses used along the routing path between a pair of routers s and d is not constrained. However, no two bypasses used along the same route may overlap, which is modeled by

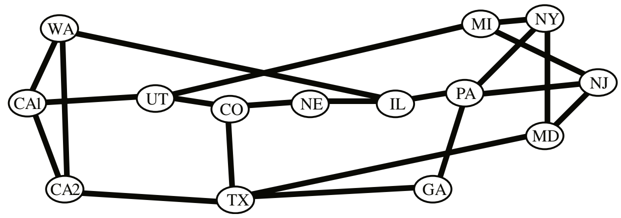

Fig. 4. The expanded T1 topology of the NSFnet from 1991.

$$\forall s,d,i,j \in V: \; \sum_{xy} R^n_{xy}(i,j) \cdot f^{sd}_{xy} \le 1. \tag{16}$$

Finally, we include a constraint for the capacity utilization of the established bypasses and extend the IP link capacity constraint to take the IP offloading into account. For all existing IP links (i,j) $(\hat{L}_{ij} = 1)$,

$$\sum_{sd} \lambda^{\text{peak}}_{sd} \cdot R^n_{sd}(i,j) \left(1 - \sum_{xy} R^n_{xy}(i,j) \cdot f^{sd}_{xy} \right) \le \alpha C_{ij} \tag{17}$$

must hold, so that the total capacity used on the original IP link is not exceeded: the original traffic on the link (i,j) is given by $\lambda^{\text{peak}}_{sd} \cdot R^{sd}_{ij}$, and the term $(1 - \sum_{xy} R^{xy}_{ij} \cdot f^{sd}_{xy})$ excludes traffic that is bypassed over the link (i,j), whereas the term R^{xy}_{ij} ensures that the original route from x to y (now bypassed) includes the link (i,j). Finally, the capacity of the established bypasses is constrained by

$$\forall x,y \in V, \quad \hat{L}_{xy} = 0: \; \sum_{sd} \lambda^{\text{peak}}_{sd} \cdot f^{sd}_{xy} \le \alpha \sum_t X^t_{xy} \cdot C^t. \tag{18}$$

Modeling Switch-On and Switch-Off: Note that the formulation currently does not restrict the circuit setup or tear-down between any pair of nodes in the network. In both the Switch-Off and Switch-On model, we first compute the base topology (n = peak for Switch-Off and n = off-peak for Switch-On), where the IP connectivity is fixed but the link capacity required needs to be determined. After the initial step, we allow Switch-On or Switch-Off on this topology to save energy in the next time interval.

For the initial time interval, we introduce constraints on L^n_{xy} to guarantee that only a specific set of links exist in the network. For example, to model the initial IP topology as shown in Fig. 2, we introduce constraints such as $L^n_{ce} = 1$ and $L^n_{ac} = 0$ to ensure that link (c,e) exists and link (a,c) does not exist in the initial topology.

In the Switch-Off scheme, the initial time interval is given by n = peak, and in the interval n = off-peak we introduce a

TABLE II
TRAFFIC MATRIX (IN MEGABITS PER SECOND) OF THE NSFNET BACKBONE

	WA	CA1	CA2	UT	CO	TX	NE	IL	PA	GA	MI	NY	NJ	MD
WA	4.72	23.84	10.41	2.42	17.47	0.78	4.78	22.14	3.04	1.65	27.71	8.6	3.93	17.01
CA1	63.92	3.48	54.23	26.78	52.12	23.27	35.45	137.75	10.17	19.03	71.05	91.68	49.1	68.97
CA2	9.7	42.28	0.03	41.43	7.56	32.33	7.7	76.15	8.92	4.11	45.9	5.52	12.37	19.18
UT	6.24	5.52	12.13	0	1.7	0.54	0.62	2.56	1.78	2.9	11.66	10.81	2.13	6.2
CO	109.13	142.22	16.91	3.05	0.32	3.59	9.58	55.31	21.35	15.93	64.1	105.39	11.72	19.34
TX	1.64	14.7	3.05	4.91	3.02	0	2.32	2.39	0.78	3.44	5.38	4.29	1.37	6.18
NE	32.89	55.12	90.94	3.98	19.59	7.02	0	101.49	17.62	19.52	136.91	82.96	21.04	145.67
IL	13.29	208.49	186.98	7.58	25.08	2.37	86.3	0.28	39.07	29.34	80.05	63.25	17.96	79.02
PA	75.49	17.73	33.2	5.34	22.21	6.06	22.28	54.24	0	35.22	98.39	131.21	40.6	56.12
GA	1.66	37.27	9.12	3.32	19.86	8.43	4.43	50.74	6.08	0.12	32.29	23.26	11.28	12.77
MI	9.93	33.43	51.82	4.5	8.4	11.55	16.7	33.68	18.2	22.33	40.44	53.04	28.69	33.06
NY	27.76	117.19	17.66	13	38.22	6.36	15.4	50.95	35.2	26.16	188.12	66	24.89	58.64
NJ	35	49.19	16.53	6.7	7.48	0.76	3.99	21.69	104.61	31.73	61.49	70.41	6.28	46.4
MD	72.8	201.79	48.26	20.41	79.36	28.29	29.07	81.65	27.21	1.48	115.29	122.31	55.78	108.11

Notes.
Source: Data taken from [27].

constraint to ensure that circuits can only be switched off:

$$X_{xy}^{\text{off-peak},t} \le X_{xy}^{\text{peak},t}. \tag{19}$$

Similarly, in the Switch-On scheme, the initial time period is given by n = off-peak, and during n = peak we allow only Switch-On by making sure that the number of circuits provisioned during the peak interval is always greater than or equal to the number of circuits provisioned during the off-peak interval. This constraint is mathematically formulated exactly the same as Eq. (19).

V. PERFORMANCE EVALUATION

In our performance evaluation, we used the expanded T1 topology from the NSFnet from 1991 (see Fig. 4), and we linearly scaled up the traffic matrix, since the original traffic values λ_{sd} would be too low for today's standards in core network traffic. The scaled-up values are used as the peak daytime traffic, and the corresponding traffic matrix is shown in Table II. (In that matrix, we set all λ_{sd} with $s = d$ to zero, since self-traffic has no effect on the backbone's link loads.) The ratio of peak/off-peak traffic may vary in different networks, so we surveyed five different traffic scenarios, differing in the nighttime traffic load defined as a fraction of daytime traffic, namely, 30%, 40%, 50%, 60%, and 70%. Since the optimum routing and the corresponding dimensioning of link capacities is sensitive to traffic variations, we introduced random noise to the scale up so that the actual size of each source–destination flow varies in a certain range, shown in Table III. For each traffic scenario, we generated multiple cases, in order of hundreds, each with a peak and an off-peak traffic matrix, and solved the ILPs relative to the different energy saving schemes and used the averages. As the ILP solver, we used the Gurobi Optimizer [28] on a standard desktop computer with a 2.5 GHz quad-core CPU.

For the purpose of this evaluation, the NSFnet topology was used as both the basic IP topology and the transport network topology. The IP network port granularities and their normalized power consumption, as well as the power consumption of optical transport circuits, are shown in

TABLE III
SCALING OF TRAFFIC VALUES FROM TABLE II

Traffic scenario	Value range
Daytime traffic	[60,000, 70,000] times original traffic
Nighttime 30%	[25%, 35%] of daytime traffic
Nighttime 40%	[35%, 45%] of daytime traffic
Nighttime 50%	[45%, 55%] of daytime traffic
Nighttime 60%	[55%, 65%] of daytime traffic
Nighttime 70%	[65%, 75%] of daytime traffic

TABLE IV
POWER CONSUMPTION

Cap (Gb/s)	$\dot{P}_{xy}^t$	P^t
10	1	5
40	3	16
100	7	34

Table IV. We assume an opaque optical network, and, as can be seen from that table, the power consumption of a transport connection depends on its physical hop count (i.e., $\dot{P}_{xy}^t \neq 0$) and its capacity ($\dot{P}_{xy}^{10} < \dot{P}_{xy}^{40} < \dot{P}_{xy}^{100}$).

We used the hop count as the metric for our SPF variant ($w_{xy} = 1$), and the maximum link utilization threshold was assumed to be $\alpha = 0.7$. The durations of the peak and off-peak time intervals were fixed as $\tau_{\text{peak}} = 16$ h and $\tau_{\text{off-peak}} = 8$ h.

The first set of results is shown in Fig. 5, where we measure the energy saved per day for all three schemes compared to a static network that is optimized for peak traffic and uses no saving scheme at all (i.e., it always consumes the same power). The results show that considerable energy savings are always possible with load adaptive schemes and that they depend linearly on the off-peak/peak traffic ratio. In other words, networks with a low difference between daytime and nighttime traffic do not benefit from load adaptive energy saving schemes as much as networks in which that difference is large. It can also be observed that both Switch-On schemes clearly outperform the Switch-Off scheme. This is primarily due to the fact that the use of additional connections means that the maximum hop count decreases as compared to the base topology used in the Switch-Off scheme, which leads to

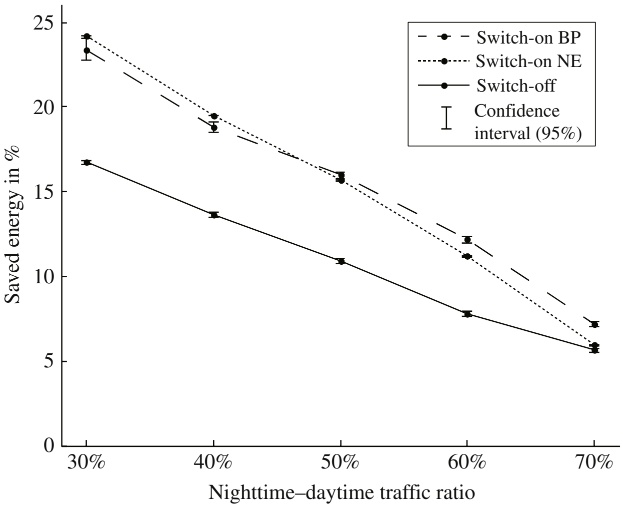

Fig. 5. The energy saved (compared to a static network with no saving scheme deployed) for each scheme depending on the nighttime–daytime traffic ratio.

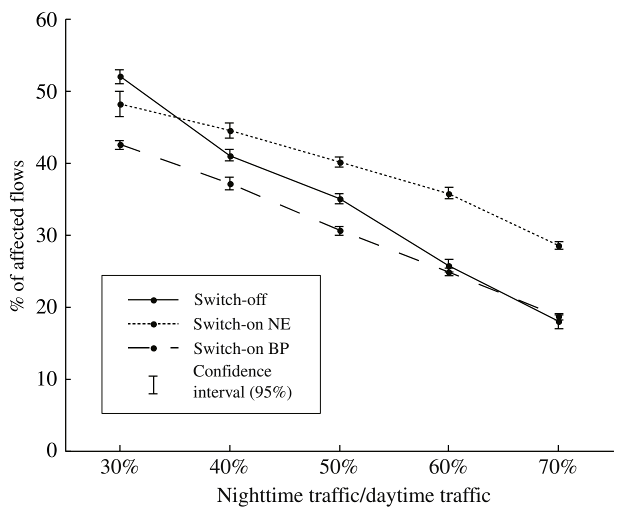

Fig. 6. Fraction of flows that undergo rerouting during the transition from the peak to the off-peak time interval depending on the nighttime–daytime traffic ratio.

less packet processing and lower installed router capacity, and therefore lower power consumption.

As we mentioned in Subsection III.B, the bypass-based Switch-On scheme minimizes the impacts of the establishment of additional connections on the IP routing stability (i.e., it does not require any routing reconvergence), and therefore it is clear that the network operation is much more stable during the transition. However, any changes in the traffic path mean that reconfigurations in the monitoring and the event correlation database in the Network Management System are required for proper operation. Figure 6 shows the fraction of flows that undergo rerouting during the transition from the peak to the off-peak time interval, whereas Fig. 7 shows the total number of the mentioned reconfigurations in the management system, which was measured as follows. For each source–destination

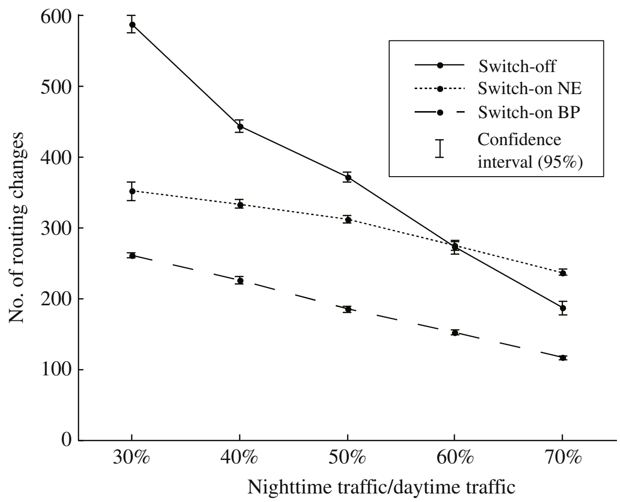

Fig. 7. Number of routing changes for each scheme depending on the nighttime–daytime traffic ratio.

TABLE V
INSTALLED INTERFACES IN THE CASE OF 50% OFF-PEAK/PEAK TRAFFIC RATIO

Mechanism	10 Gb/s	40 Gb/s	100 Gb/s
Switch-Off	42.4	22.8	1.6
Switch-On BP	92.0	7.4	0
Switch-On NE	94.2	6.1	0

flow that undergoes rerouting during the transition from the peak to the off-peak time interval, the old (i.e., original) path and the new (i.e., rerouted) path are compared, whereas every link belonging to the original path but not to the rerouted path, as well as every link belonging to the new path but not to the old path, counts as one unit in Fig. 7. The results show that, even though the three schemes do not differ significantly in the number of rerouted flows (see Fig. 6), it can be seen in Fig. 7 that the Switch-On schemes can notably decrease the number of necessary routing reconfigurations (especially in the case of small daytime–nighttime traffic ratios), thereby reducing the effort required to reconfigure the monitoring system and the event correlation database.

For all results presented so far, the IP routing stability parameter σ was kept constant at 0.05, meaning that energy saving has a high priority, while routing stability has a lower priority. As we explained in Section IV, this parameter is used only in the Switch-On scheme with conventional Network Engineering, although it is worth mentioning that this method is not restricted to this scheme and can also be applied in the Switch-Off scheme, in order to reduce its impact on routing stability. However, we did not consider this option in our performance evaluation, mainly because the Switch-Off scheme is already performing badly compared to the other schemes, and the introduction of the rerouting penalty would lower its performance even more. In our comparison, we used the Switch-Off model as it was originally proposed.

In order to study the effect of σ-variation on the Switch-On scheme where it was introduced, we made an additional analysis, with the result shown in Fig. 8. While Fig. 8(a)

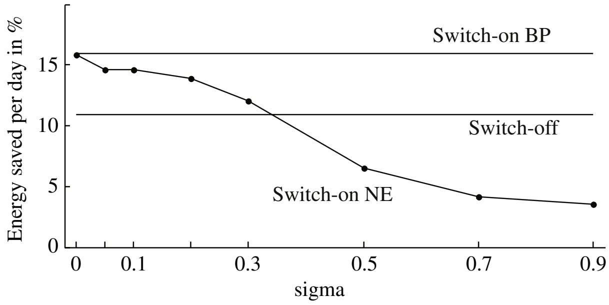

(a) Comparison of energy savings with variation of σ.

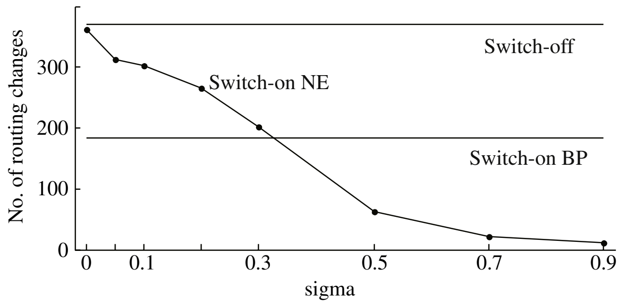

(b) Comparison of number of routing changes with variation of σ.

Fig. 8. For the Switch-On scheme using conventional Network Engineering: trade-off between the number of routing changes and energy saved per day depending on σ (weighting factor for IP routing stability). The constant values of Switch-On using hidden bypasses and Switch-Off are given as straight lines for orientation.

shows the energy saved per day, Fig. 8(b) shows the number of routing changes. We can see from that figure that this rerouting punishment method very much allows one to adapt the trade-off between routing stability and energy conservation to the network operator's preferences. On the one hand, by setting a very high σ (e.g., $\sigma > 0.7$), the number of routing changes can be reduced below 20, but at the cost of poor energy saving capabilities. On the other hand, by setting a very low σ (e.g., $\sigma < 0.05$), the energy saving capabilities almost reach those of the Switch-On scheme using hidden bypasses, but at the cost of rather poor routing stability.

In order to reduce energy consumption while keeping the network stable, the Switch-On schemes employ a large number of interfaces of small granularity. As seen in Table V, they employ more than twice the number of 10 Gbit/s interfaces as the Switch-Off scheme, while they use a significantly lower number of high capacity interfaces. Even though the total capacity installed per router is lower, the Switch-On scheme may trigger more frequent network planning cycles as the increase in traffic may require that new interfaces be installed. Also, frequent upgrades of small capacity interfaces may prove to be more costly from the CAPEX perspective.

The network performance in terms of transmission delay greatly depends on the average IP hop count per transferred bit, since each IP hop adds some delay to the transmission time due to buffering and processing of packets. All three schemes examined have the same tendency to use fewer links in the night than in the day due to the lower traffic demand, resulting in an increased hop count during the night. But when we compared the three schemes with each other regarding the average hop count per transferred bit over a whole day, we observed that the Switch-On schemes show values that are about 17% lower, suggesting that the Switch-On schemes are able to reduce the average network delay and thus increase the user's quality of experience.

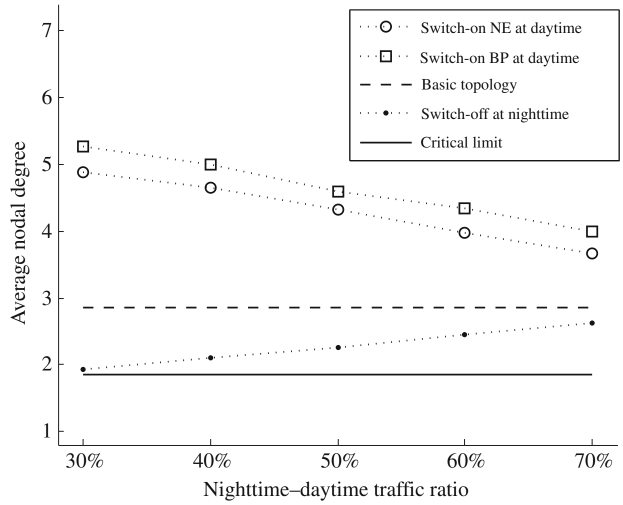

Fig. 9. Network resilience measured in the average nodal degree. The static value of the basic topology represents the values for both Switch-On schemes in the night as well as the value for the Switch-Off scheme in the day. The critical limit shows the lower bound below which the 14-node network becomes disconnected.

Our results for the average nodal degree depending on the nighttime–daytime traffic ratio are depicted in Fig. 9. This parameter is an important factor regarding network resilience, since a network with a low average nodal degree is more likely to become disconnected by link failures. In this context, disconnected means that there are node pairs without an interconnecting path. In our 14-node network, there is a critical limit for the average nodal degree of about 1.86, which represents the case where the network is a tree. To reduce the average nodal degree further implies that nodes definitely become disconnected. Since the topology used during the night in both Switch-On schemes as well as the topology used during the day in the Switch-Off scheme are exactly the same (i.e., the basic topology, which they are initially designed for), the three corresponding lines overlap in this graph. The graph shows that the Switch-On schemes can significantly increase the resilience during high traffic conditions, whereas the Switch-Off scheme decreases the average nodal degree during the night to very close to the critical limit, suggesting that Switch-Off is not suitable in scenarios in which network resilience is an important issue.

We were also interested in how much our results depend on the actual traffic matrix, i.e., by how much can the performance of the examined load adaptive energy saving schemes be influenced by traffic anomalies. To this end, we simulated elephant flows (i.e., single but very high bandwidth traffic flows) that may result from collaborative science projects (e.g., high data volume distributed grid computing for the LHC experiments), online backup services, or other bulk file transfers. We also varied the most important parameters of traffic anomalies, which are listed below.

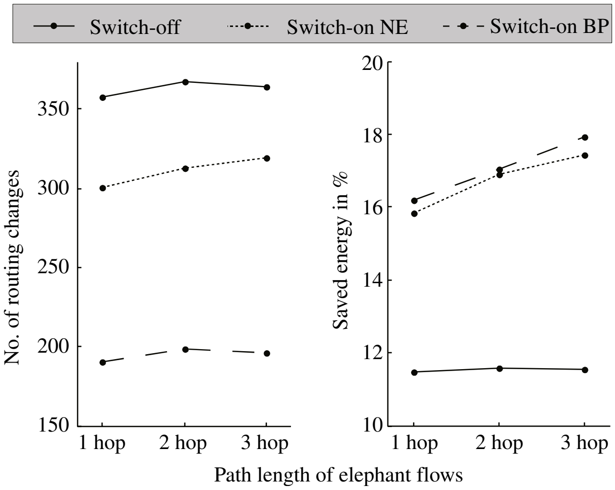

Fig. 10. Impact of elephant flows on the main performance parameters depending on the hop count of the elephant flows.

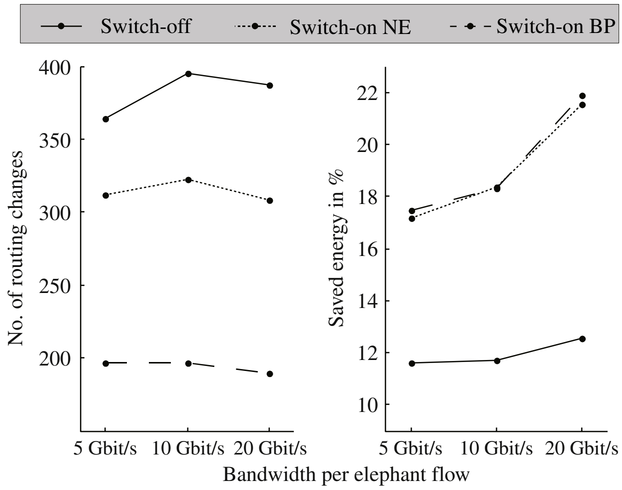

Fig. 11. Impact of five elephant flows on the main performance parameters of the energy saving schemes depending on the bandwidth of the elephant flows.

1) The elephant flows' IP path length: We varied the hop count of the traffic anomalies to examine its impact on the saving schemes.
2) The bandwidth per elephant flow: What is the effect of increasing the bandwidth demand of traffic anomalies on the saving scheme's performance?
3) The number of elephant flows: Are the saving schemes able to handle an increasing number of high demand traffic anomalies?
4) The elephant flows' time of occurrence: Can the schemes handle the scenario in which elephant flows appear only at night (e.g., large scale online backup services)?

Figure 10 shows the impact of daytime elephant flows depending on their IP path distance (i.e., hop count). For this examination, we randomly added five source–destination flows with the same hop count and each with 5 Gbit/s of bandwidth demand. Please note that the limitation to three hop paths in this examination is due to the NSFnet topology used. Again as in Fig. 5, the energy saved in percentage terms (right-hand side of Fig. 10) is related to the static network (i.e., without any load adaptive energy saving scheme deployed). The result is almost identical to the previous result in Fig. 5, with the exception that both Switch-On schemes slightly profit by long distance elephant flows. This is due to the fact that the elephant flows are likely to cause direct connections that bypass IP hops (i.e., avoid unnecessary electronic processing). On the other hand, the number of routing changes does not seem to depend on the IP path length: the differences observed are not significant.

In our next evaluation, we examine the influence of the elephant flows' bandwidth. To this end, we randomly chose five source–destination pairs in order to add a same size elephant flow. Figure 11 shows that the bandwidth per elephant flow has stronger impact than their hop count on the percentage energy saved related to the static network (right-hand side of Fig. 11). Both Switch-On schemes can clearly benefit from bigger flows, whereas the number of routing changes also does not seem to depend on the elephant flows' bandwidth.

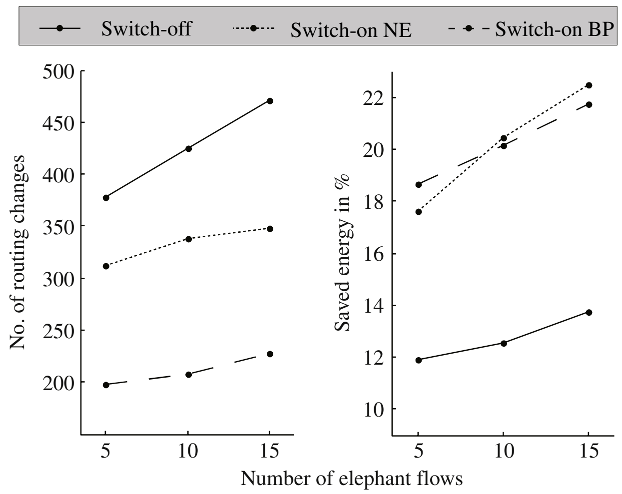

Fig. 12. Impact of the number of elephant flows on the main performance parameters of the energy saving schemes.

We next study how the number of elephant flows influences the performance of the schemes. Therefore we randomly chose 5, 10, or 15 source–destination router pairs and added a 10 Gbit/s flow to each of them. The result is depicted in Fig. 12, where it can be seen that, on the one hand, all three schemes can improve their energy saving performance (again compared to the static network) with increasing number of high capacity flows, while, on the other hand, it can also be seen that the number of routing changes increases accordingly. However, this simulation also does not produce abnormal results, since the increase of energy saving performance as well as the increasing number of routing changes can be explained with the increased daytime–nighttime traffic ratio, i.e., the elephant flows occur during the day and increase the overall daytime traffic of the network by more than 35%. Exactly the same behavior of the three schemes was already observed in our main results (see Figs. 5 and 7).

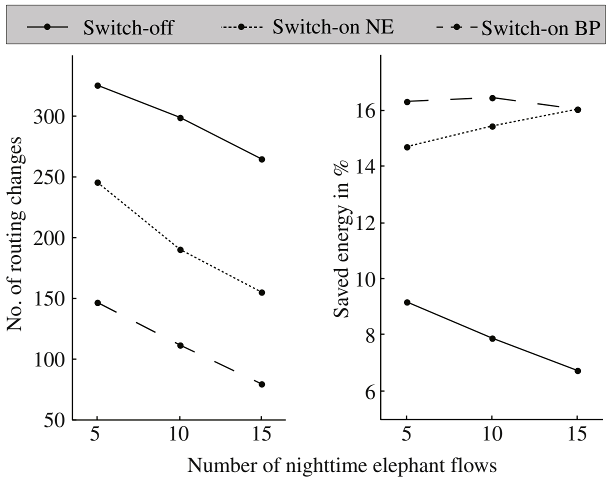

Fig. 13. Impact of elephant flows during the night, e.g., in the case of online backup services.

Finally, the last case that we examine is the occurrence of elephant flows during the night, which can be expected in services such as online backup or scientific data transfers scheduled for periods of low traffic demand. The results are shown in Fig. 13, where it can be seen that, in terms of energy saving performance, all three schemes suffer a little from the reduced nighttime–daytime traffic ratio that arises with nightly elephant flows (compare the values for the 50% traffic ratio in Fig. 5). The Switch-On scheme that uses hidden bypasses shows a decreased but stable performance, whereas the other Switch-On scheme using conventional Network Engineering can even improve its performance slightly when the number of elephant flows increases. This is due to its capability to reuse high capacity links (dimensioned for nighttime elephant flows) during the day by other flows, whereas the bypass scheme, unable to apply regular traffic engineering, can only avoid some of the bypasses for daytime traffic. On the other hand, the Switch-Off scheme rapidly loses performance with the increasing number of nighttime elephant flows. This is caused by the fact that we consider regular network planning for the Switch-Off scheme, i.e., we ensure that the network can handle all traffic peaks. Hence, in order to avoid the case in which a nightly elephant flow overloads a link, its demand must also be considered when the network is dimensioned for daytime traffic. In terms of necessary routing reconfigurations, we can see that all three schemes can decrease the number of routing changes with increasing number of elephant flows, which is again due to the decreasing nighttime–daytime traffic ratio.

VI. Conclusion

Although load adaptive energy efficiency schemes can significantly reduce energy consumption, their impact on routing stability, degradation of quality of service, and network resilience have not been sufficiently studied to date. In this paper, we have provided a comparative performance analysis of the proposed load adaptive schemes and showed that a new energy saving scheme, which we refer to as *Switch-On*, can take into consideration the above-mentioned parameters while saving energy. Our analysis and performance results showed that both Switch-On schemes (i.e., using traditional Network Engineering or a hidden bypass) can outperform the Switch-Off scheme, especially for scenarios in which the difference between off-peak and peak traffic is large. We showed that, unlike the Switch-Off scheme, both variants of the Switch-On scheme do not affect the path redundancy and therefore do not weaken the network resilience. Furthermore, the Switch-On schemes do not increase the average hop count, and thus do not impact the end-to-end packet delay.

An important result of our analysis is that energy efficiency does not need to excessively destabilize routing stability, and to this end we propose that network providers either use hidden bypasses or introduce routing stability into the objective of the energy minimization. Hidden bypasses are especially attractive, as they can greatly reduce the impact of temporary topology changes on the stability of the network. For scenarios in which these hidden bypasses are not available, we showed that, by adjusting the weighting factor for the rerouting penalty, the trade-off between routing stability and energy saving can be balanced precisely to the network operator's preference. From this perspective, we believe that our study is an important contribution to the existing body of research on energy conservation, which typically does not consider the trade-offs relative to IP network routing stability, and network stability in general.

Significant future research is ahead of us. Although we showed that both variants of the Switch-On schemes result in a reduced total network capacity, our approach requires a comparatively larger number of small capacity interfaces. This might be a challenge from a network planning perspective, due to the frequent capacity upgrades, which requires further study. We plan to extend this work to scenarios which do not assume full knowledge of the traffic matrix and allow partitioning of large IP ports to ensure stable network planning cycles in the network. We are also interested in fundamental understanding of the effects of increased number of time intervals on IP routing stability and power consumption. Our future work will also include fast heuristics, which will allow us to study a wider range of network topologies and scenarios.

References

[1] Annual Cisco Visual Networking Index Forecast Projects Global IP Traffic to Increase More Than Fourfold by 2014, *Cisco Press Release* [Online]. Availble: http://newsroom.cisco.com/dlls/2010/prod_060210.html.

[2] J. Baliga, K. Hinton, and R. S. Tucker, "Energy consumption of the Internet," in *COIN-ACOFT 2007*, Melbourne, June 2007.

[3] M. Walker, *Surfing the Green Wave in Telecom*, Ovum, May 9, 2008.

[4] S. Mingay, "Green IT: The new industry shock wave," *Gartner Report*, Dec. 2007.

[5] O. Tamm, C. Hermsmeyer, and A. M. Rush, "Eco-sustainable system and network architectures for future transport networks," *Bell Labs Tech. J.*, vol. 14, pp. 311–327, 2010.

[6] B. Wu, K. L. Yeung, and P.-H. Ho, "Virtual topology design for minimizing network diameter and average hop count in WDM networks," *J. Opt. Commun. Netw.*, vol. 2, no. 12, pp. 1077–1086, Dec. 2010.

[7] K. Thompson, G. J. Miller, and R. Wilder, "Wide-area Internet traffic patterns and characteristics," *IEEE Network*, vol. 11, no. 6, pp. 10–23, 1997.

[8] C. Lange and A. Gladisch, "Energy efficiency limits of load adaptive networks," in *OFC/NFOEC'10*, San Diego, CA, Mar. 2010.

[9] G. Frederic, M. Dorian, M. Joanna, and O. Brice, "Minimizing routing energy consumption: From theoretical to practical results," in *GreenCom 2010*, Hangzhou, China, Dec. 2010.

[10] B. Sanso and H. Mellah, "On reliability, performance and Internet power consumption," in *IEEE DRCN'09*, Washington, DC, Oct. 2009.

[11] L. Chiaraviglio, M. Mellia, and F. Neri, "Reducing power consumption in backbone networks," in *IEEE ICC'09*, Dresden, Germany, June 2009.

[12] L. Chiaraviglio, M. Mellia, and F. Neri, "Energy-aware backbone networks: A case study," in *IEEE Int. Workshop on Green Communications (GreenComm'09)*, Dresden, Germany, June 2009.

[13] W. Fisher, M. Suchara, and J. Rexford, "Greening backbone networks: reducing energy consumption by shutting off cables in bundled links," in *ACM SIGCOMM Workshop on Green Networking 2010*, New York, Aug. 2010.

[14] A. Coiro, M. Listanti, A. Valenti, and F. Matera, "Reducing power consumption in wavelength routed networks by selective switch off of optical links," *IEEE J. Sel. Top. Quantum Electron.*, vol. 17, no. 2, pp. 428–436, Apr. 2011.

[15] S. Zhang, D. Shen, and C. Chan, "Energy efficient time-aware traffic grooming in wavelength routing networks," in *GLOBECOM 2010*, Miami, Dec. 2010.

[16] F. Farahmand, I. Cerutti, M. M. Hasan, and J. P. Jue, "Energy-efficiency of drop-and-continue traffic grooming," in *Optical Fiber Communication Conf.*, 2011, OTuR6.

[17] G. Shen and R. S. Tucker, "Energy-minimized design for IP over WDM networks," *J. Opt. Commun. Netw.*, vol. 1, no. 1, pp. 176–186, June 2009.

[18] F. Idzikowski, S. Orlowski, C. Raack, H. Woesner, and A. Wolisz, "Saving energy in IP-over-WDM networks by switching off line cards in low-demand scenarios," in *ONDM'10*, Kyoto, Japan, Feb. 2010.

[19] B. Puype, W. Vereecken, D. Colle, M. Pickavet, and P. Demeester, "Power reduction techniques in multilayer traffic engineering," in *ICTON 09*, Island of Sao Miguel, Azores, Portugal, July 2009.

[20] E. Yetginer and G. N. Rouskas, "Power efficient traffic grooming in optical WDM networks," in *GLOBECOM 2009*, Honolulu, USA, Dec. 2009.

[21] Y. Zhang, M. Tornatore, P. Chowdhury, and B. Mukherjee, "Energy optimization in IP-over-WDM networks," *Opt. Switching Networking*, vol. 8, no. 3, pp. 171–180, July 2011.

[22] M. Caria, M. Chamania, and A. Jukan, "Trading IP routing stability for energy efficiency: A case for traffic offloading with optical bypass," in *ONDM 2011*, Bologna, Italy, Feb. 2011.

[23] M. Caria, M. Chamania, and A. Jukan, "To switch on or off: A simple case study on energy efficiency in IP-over-WDM networks," in *HPSR 2011*, Cartagena, Spain, July 2011.

[24] M. Chamania, M. Caria, and A. Jukan, "Achieving IP routing stability with optical bypass," in *IEEE ANTS 2009*, Delhi, India, Dec. 2009.

[25] "Energy efficiency for network equipment: two steps beyond greenwashing," *Juniper Networks White Paper*, 2010.

[26] M. Chamania, M. Caria, and A. Jukan, "A comparative performance analysis of IP traffic offloading schemes over dynamic circuits," in *INFOCOM 2011—Mini Conf.*, Shanghai, China, Apr. 2011.

[27] B. Mukherjee, *Optical WDM Networks*. Springer, 2006, p. 422.

[28] Gurobi Optimizer 3.0 [Online]. Available: http://www.gurobi.com/.