To Switch *On* or *Off*: A Simple Case Study on Energy Efficiency in IP-over-WDM Networks

Marcel Caria, Mohit Chamania and Admela Jukan
Technische Universität Carolo-Wilhelmina zu Braunschweig
Email: {caria, chamania, jukan}@ida.ing.tu-bs.de

***Abstract*—In this paper, we present a simple case study on energy efficiency in IP-over-WDM networks with dynamic circuit capability and compare two different load adaptive schemes, referred to as *switch-on* and *switch-off*. While the switch-off technique was already proposed for energy conservation, the switch-on approach is a new paradigm which is based purely on one-time implementations of dynamic circuit capability. Our results show that both approaches can significantly reduce the power consumption and decrease the necessary totally installed capacity, but unlike the switch-off scheme, the switch-on scheme does not affect the path redundancy in the network. Furthermore, switch-on can reduce the number of routing reconfigurations required in the network. While this seems ideal, our results also show that the switch-on scheme uses a large number of small capacity interfaces which may not be suitable from a network planning perspective as it might require frequent capacity upgrades, which alone is an interesting avenue for future research.**

I. Introduction

Internet based services play a substantial role in today's world economy and have been exhibiting significant growth rates with predictions that the global IP traffic will increase 4.3 times from 2009 to 2014, reaching 63.9 exabytes per month in 2014 [1]. This trend is mainly driven by the ongoing increase of all forms of consumer video (IPTV, VoD, P2P, YouTube, etc.) coupled with increasing total number of Internet users worldwide. This unprecedented increase in capacity has not only driven up demand for bandwidth from the Internet infrastructure but has also presented providers with new challenges, especially in the energy consumed by the Internet infrastructure. Although the ongoing technological progress of network equipment reduces the consumed energy per transmitted bit by around 13% per year [2], this will likely not be sufficient in itself since there remains at least a 20% gap compared to the traffic growth. Without new and efficient energy conservation approaches, the energy consumption is likely to become the Internet's main bottleneck. This is especially critical in today's Internet backbone, which is typically over-provisioned and thereby lightly loaded at most times [3].

The most obvious way to reduce energy consumption is to selectively switch-off underutilized resources during periods of low loads [4], [5], as Internet backbone traffic typically follows a fairly stable daily pattern. As shown in [6], link loads in backbone networks carry peak loads between 08:00 and 24:00 and resources can be switched-off in off-peak hours to reduce energy consumption. Recently, the so-called multi-layer approaches have been proposed [7], [8], with the main idea of switching off active components in the IP network by re-routing the IP traffic into the WDM network, which has a significantly lower power consumption per transferred bit.

In our work in [9], we highlighted the deficiencies of the current *switch-off* model which refers to the current proposals for network engineering and planning, where the IP network is over-provisioned and optimized for peak network loads, and the network equipment is switched-off following the daily pattern of low traffic load. The two major drawbacks of the switch-off scheme are: reduced path redundancy and large-scale routing re-configurations. Path redundancy is necessary to deal with network failures, and switching-off IP links in order to optimize for energy consumption can leave the IP topology sparsely connected which can be critical in case of a failure. Change in the IP topology also impacts routing stability as it triggers routing re-convergence which can disrupt services in the order of 10s of seconds [10]. In addition, re-routing of traffic requires re-configuration of the monitoring systems and the event correlation databases, which are responsible for correlating alarms from different components to identify the root cause of network faults. Problems with redundancy and reconfigurations triggered due to re-routing are the primary reason why providers today do not switch-off links to reduce energy consumption in core IP networks.

In this paper, we present a new scheme, referred to as link *switch-on*, which is aimed at addressing the aforementioned problems with current energy-saving schemes. In the switch-on scheme, we allocate the IP link capacities for low traffic loads and make use of network engineering paradigms [11] using dynamically switched circuits to boost IP network capacity during periods of high traffic loads in the network. The nature of the switch-on paradigm ensures that network redundancy is improved due to introduction of new links, and our results show that this paradigm can still improve upon energy savings as seen in the traditional switch-off scheme. In order to compare the two proposed schemes, we developed an analytical framework using Integer Linear Programming (ILP) and model both the switch-on and switch-off schemes. We also provide a mechanism to control the extent of routing re-configurations in the network for both the switch-on and switch-off model in our analytical framework, and an interesting study shows that a trade-off between energy conservation and routing stability can be best gained, when the IP link capacities are dimensioned for low traffic demand and adapted to peak traffic loads, instead of designing it for

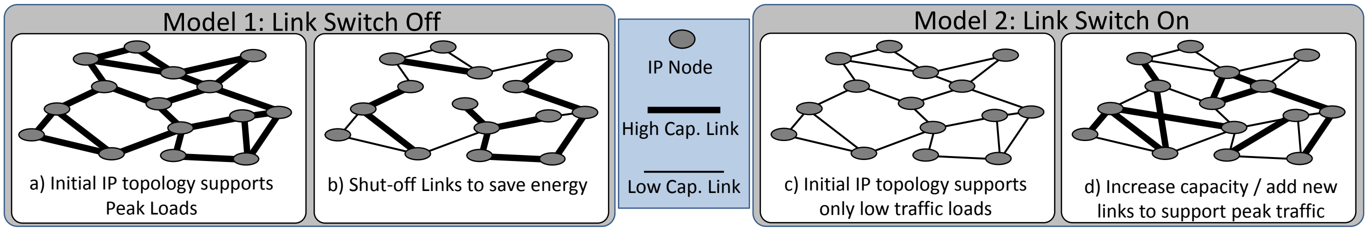

Fig. 1. Explanation of the both energy efficiency schemes: In the Switch-Off scheme, IP links in the initial topology are designed to support peak loads, and some links are switched off during off-peak hours. In the Switch-On scheme, IP links are designed for optimum operation during off-peak hours and new IP links are switched on for high loads.

the maximum traffic demand and switching off underutilized elements in case of low traffic.

The rest of the paper is organized as follows. Section II presents the case study, including the switch-on and switch-off schemes. Section III presents the corresponding analytical model for optimizations. Section IV analyzes the results obtained. Section V concludes the paper and provides directions for further research.

II. A Simple IP-over-WDM Case Study

In this section, we first introduce the basic principles behind the switch-on and switch-off schemes for energy conservation in IP networks. We then present the IP-over-WDM network architecture assumed in analytical formulation in Section III.

A. Switch-on and Switch-off Schemes for Energy Efficiency

The basic principle behind the two schemes for optimizing energy usage in IP networks is presented in Fig. 1. In both schemes, we assume the same initial IP connectivity as seen in Fig. 1(a) and (c). In the traditional link switch-off scheme, the capacity of IP links (represented as bold lines in Fig. 1(a)) is dimensioned to support IP traffic at peak loads, and during off-peak hours, these links are significantly underutilized. In order to save energy, as shown in Fig. 1(b), some of these links can be switched off, or if they consist of multiple aggregated IP ports, they can lower the number of active IP ports to save energy. Previous studies show that this model can significantly reduce the energy consumption. However, as links are switched off, the network topology becomes sparse as can be seen in Fig. 1(b). This phenomenon has a two-fold impact: a sparse network topology means that the average hop-count inside the IP core increases, which can lead to degradation of service quality, and a sparsely connected IP topology is also more vulnerable to network outages due to failures. Note that once the traffic cycles back to peak loads, the network switches on the IP links that were powered down to re-obtain the original IP network state like indicated in Fig. 1(a).

In the switch-on scheme, the capacity of the IP links (represented as thin lines in Fig. 1(c)) are dimensioned to accommodate traffic in off-peak hours, and mechanisms of network engineering using dynamic circuit services are used to upgrade the topology to support peak traffic loads. As seen in Fig. 1(c), the initial link capacities are low compared to its switch-off counterpart in Fig. 1(a). To support traffic at peak loads, dynamic circuits are commissioned between pairs of neighboring IP routers to boost the capacity of existing IP links, or are added between non-neighboring IP routers as new IP links, as illustrated in Fig. 1(d). In contrast to the switch-off scheme, the use of additional IP links means that the average hop count is reduced and the path redundancy is higher during peak load conditions due to the increase in the average nodal degree in the IP network. The reduced hop-count in both the case of peak and off-peak traffic means that the total installed IP capacity required in the network will be lower as compared to switch-off scheme. It also indicates that the switch-on scheme is likely to be more energy efficient than the link switch-off scheme as illustrated by the results in Section IV. However, a drawback of the switch-on scheme is the need for a larger number of IP ports with small capacity, which can affect the traditional network planning cycles in the network.

B. IP-over-WDM Architecture

We now present the basic assumptions on the network architecture (shown in Fig. 2) used for the analytical model presented in section III.

Network Topology: We assume that IP routers are co-located with transport network switches, as shown in Fig. 2. The same architectural structure is used in both, the *switch-on* and the *switch-off* scheme. We assume a transport network that can support dynamic circuit service, and as shown in Fig. 2, each IP router is co-located with a circuit switch (WDM based). IP links are created by establishing circuits between the corresponding (i.e., co-located) switches in the transport network. In our model, we assume that the transport network can provide a dynamic circuit between any pair of IP routers in order to limit the complexity of the ILP model. We also assume here that for both schemes, the IP router chassis size is the same and the total capacity available to interface the IP router with its co-located transport switch is very large, so it does never present a bottleneck. This assumption is primarily made for the switch-on scheme as initial topology and capacities are designed for off-peak loads. Our results will later go on to show that the total IP switching capacity and the router-switch inter-connect for the switch-on scheme are actually lower than what is typically available in over-provisioned networks.

Port Sizes: In this case study, it is assumed that a single port in the IP router cannot be virtualized: e.g., a 10 Gb/s IP

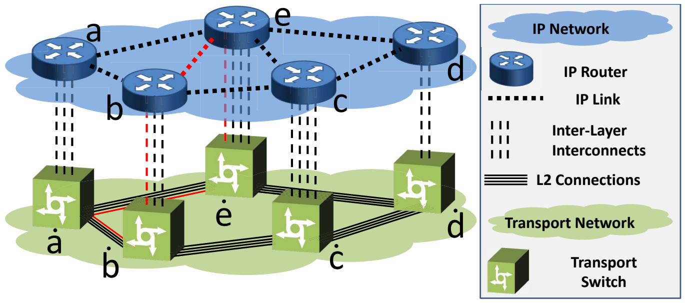

Fig. 2. Example for a new IP Link (highlighted red): The WDM circuit via $(\dot{b}, \dot{a})$ and $(\dot{a}, \dot{e})$ is commissioned in order to establish IP link (b, e) with the inter-layer inter-connections $(b, \dot{b})$ and $(e, \dot{e})$.

port *cannot* be used to provision two IP links to two different routers with capacity 5 Gb/s each. However, multiple ports can be aggregated to increase the capacity of an existing IP link between a pair of routers. In other words, two 10 Gb/s ports can be used to make a 20 Gb/s IP link to the same other router. In the switch-off scheme, IP links can either reduce link capacity by switching off one of the multiple ports used for the link, or all ports constituting one link can be switched off completely. In the switch-on scheme, there are more degrees of freedom as we allow for flexible link capacity alteration. For instance, if a certain IP link capacity is becoming insufficient in the following traffic interval, a new circuit is commissioned. Upon commissioning the circuit, either the port with insufficient IP link capacity can be aggregated with the new circuit capacity and thus increase the capacity of the existing IP link, or a new IP link can be established between a pair of routers which previously were not connected. For the same capacity, the switch-on scheme is likely to employ a larger number of ports of lower bit rate as compared to the switch-off scheme which typically uses a smaller number of ports with high bit rate. However, the total IP routing capacity required per router in the switch-on scheme is likely to be lower due to the possibility to fine tune the planned capacity and the reduction in the average hop-count, which justifies the original assumption.

Routing in the IP and the WDM Network: We assume in both schemes that a variant of the Shortest Path First (SPF) routing protocol is deployed in the IP network to route traffic. We also assume that the routing from source to destination router uses a single unique routing path in the IP layer and mechanisms such as equal-cost multi-path in OSPF are not employed in our network. In order to reduce the complexity of the case study, we further assume that the transport network has sufficient capacity to route all circuit requests for the IP network, and that all circuits used to provision an IP link between a unique pair of IP routers use the same (pre-computed) path in the transport network and are never blocked. In order to reduce the effect of routing changes on the network, we introduced a *routing penalty variable* which measures the total number of reconfigurations required in the monitoring and the event correlation databases which is proportional to the total number of links that are different between the original and re-routed path of each traffic route.

Traffic Profile: We assume that the traffic follows a strong recurring daily pattern which is split into two time intervals, called *peak* and *off-peak*. We assume complete knowledge of the traffic matrix for both intervals and the same traffic scenario is used for both energy conservation schemes. We use a notion of the *basic time interval*, which is the traffic matrix for which the IP topology was originally designed. In the switch-off scheme, the basic time interval is the one with peak network load and its topology is as shown in Fig. 1(a), whereas the basic time interval corresponds to the off-peak load in the switch-on scheme shown in Fig. 1(c).

Power Consumption: As specified before, we assume that the IP routers and the switches used in the transport network are the same for both schemes. Under this assumption, the power consumed by the router and the switch chassis are the same in both cases and are therefore not considered in our formulation. We model the total power consumed by the number and rate of IP ports active during a time interval as well as the power of circuits provisioned in the transport layer to establish these IP links. As we assume that the WDM network is used to provision an IP link, and the circuit between a pair of routers always use the same WDM routed path, we define a fixed power consumption value that depends on the optical technology used, the number of physical hops in the WDM layer and the distances between these hops. The power consumption for an IP port is defined as a function of its granularity. Note here that as port sizes increase, the power per bit decrease, e.g. the power consumed by a 40 Gb/s IP port would be lower than the power consumed by four 10 Gb/s IP ports.

III. Analytical Model

We will now present the analytical model based on Integer Linear Programming used to determine the operation of the switch-on and switch-off techniques. 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 is flexible. 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 [12]. However, the modeling of the link capacity is kept flexible wherein 1) a new link may be switched on, 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 information, and uses this information to compute the IP topology and link capacities and consequently the power consumption for different schemes.

IP Topology and Link Capacity: For all routers $x, y \in V$ and for both time intervals $n \in \{\text{peak}, \text{off-peak}\}$, the IP link

Parameter	Meaning
$\dot{x}$	Switch $\dot{x}$
$(\dot{x}, \dot{y})$	Physical Link between Switches $\dot{x}$ and $\dot{y}$
V	Set of IP Routers in Layer 3
$x \in V$	Router x, co-located with Switch $\dot{x}$
M	Maximum number of ports at any router in the network
E	Set of IP Links in Layer 3
$(x, y) \in E$	IP Link between Routers x and y
$n \in \{\text{peak, off-peak}\}$	A certain time interval
τ^n	Duration of time interval n
T	Set of available bandwidth rates for Circuits
$t \in T$	A certain bandwidth rate for a Circuit
C^t	Capacity of a Circuit of rate t
λ^n_{xy}	Traffic from x to y in time interval n
$\dot{P}^t_{xy}$	Power Consumption of a Circuit from x to y with bandwidth rate t
P^t	Power Consumption of an IP Port of rate t
α	Maximum IP link utilization
w_{xy}	IP routing weight
σ	IP routing stability parameter
β	Constant scaling parameter for $Punish^n$
Variable	**Meaning**
$X^{n,t}_{xy}$	Number of Circuits of rate t used in time interval n for IP Link (x, y)
C^n_{xy}	Capacity of IP Link (x, y) in time interval n
L^n_{xy}	Link Existence Variable: *true*, if the IP link (x, y) exists in time interval n
$R^n_{sd}(i, j)$	IP Routing Variable: *true* if Path from s to d uses link (i, j) in time interval n
P^n_{xy}	Power Consumption of IP link from x to y in time interval n
RC^n_{sd}	IP routing cost from s to d in time interval n
FT^n_{ijd}	IP forwarding variable for time interval n: *true* if router i forwards packets for d over (i, j)
$Punish^n$	Penalty for IP routing changes in time interval n

TABLE I
SUMMARY OF NOTATION

existence variable L^n_{xy} is bounded by

$$\left\lceil \frac{\sum_t X^{n,t}_{xy}}{M} \right\rceil \leq L^n_{xy} \leq \sum_t X^{n,t}_{xy} \tag{1}$$

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 $\dot{e}$ while 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 $n \in \{\text{peak, off-peak}\}$ is given as the sum of the capacity of all IP ports, and consequently all circuits that are commissioned between $\dot{i}$ and $\dot{j}$ during that time interval:

$$C^n_{ij} = \sum_t X^{n,t}_{i,j} \cdot C^t \tag{2}$$

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 from x to y during time interval n is given by:

$$P^n_{xy} = \sum_t X^{n,t}_{xy} \left(\dot{P}^t_{xy} + P^t \right) \tag{3}$$

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.

Routing Penalty: To model the effect of routing changes, we define a *routing penalty* parameter for each profile transition from $n - 1$ to 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 path, and β is the proportionality constant.

$$Punish^n = \beta \cdot \sum_{sdij} \left| R^{n-1}_{sd}(i, j) - R^n_{sd}(i, j) \right| \tag{4}$$

For a simple event like a single link failure in our example network in Fig. 2 for instance, if the route from b to e changes from $b{-}a{-}e$ to $b{-}e$, the failure event for links $b{-}a$, $a{-}e$ and $b - e$ must be reconfigured in the event correlation database.

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

$$\text{Minimize} \sum_n \sum_{xy} \left(\sigma \cdot Punish^n + (1 - \sigma) \cdot P^n_{xy} \cdot \tau_n \right) \tag{5}$$

where σ is the routing stability parameter that balances the trade-off between energy consumption and routing penalty. This parameter is set to zero in the *switch-off* model, where routing changes are not considered.

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

$$R^n_{sd}(i, j) \leq L^n_{ij} \tag{6}$$

$$\sum_p R^n_{sd}(s, p) = \sum_q R^n_{sd}(q, d) = 1 \tag{7}$$

$$\sum_p R^n_{sd}(p, k) = \sum_q R^n_{sd}(k, q) \tag{8}$$

Applied to Fig 2, (6) ensures that the IP link from a to e can only be used if it exists. (7) 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, while (8) ensures routing continuity inside the network.

IP networks typically employ a SPF routing algorithm, and constraints are introduced to model the same. For every possible link between any routers x and y, we use a pre-defined 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^n_{sd} from s to d ($s \neq d$) during the time interval n is given by

$$RC^n_{sd} = \sum_x \sum_y R^n_{sd}(x,y) \cdot w_{xy} \tag{9}$$

We assure 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

$$RC^n_{sd} \leq RC^n_{sx} + L^n_{xd} \cdot w_{xd} + (1 - L^n_{xd}) \cdot w_\infty \tag{10}$$

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 can not 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.

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. To model the *destination-based SPF packet forwarding*, we introduce a boolean variable FT^n_{ijd} which mimics a forwarding table, with

$$\sum_j FT^n_{ijd} = 1 \tag{11}$$

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

$$R^n_{sd}(ij) \leq FT^n_{ijd} \leq L^n_{ij} \tag{12}$$

for all $s, d, i, j \in V$ and $n \in \{\text{peak}, \text{off-peak}\}$, we assure that the IP routing follows the forwarding table (left hand side of (12)) and also that the link to the next hop exists (right hand side of (12)). 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. (11)). Furthermore, all routing paths must follow the touting tables (left hand side of (12)) and forwarding on a link requires that the link exists (right hand side of (12)).

Link Utilization: Using the routing information, we constrain the total capacity of an IP link to be sufficient to support the IP traffic load while limiting the link utilization to a pre-defined threshold α for $n \in \{\text{peak}, \text{off-peak}\}$ and all $i, j \in V$.

$$\sum_s \sum_d \lambda^n_{sd} \cdot R^n_{sd}(i,j) \leq \alpha \cdot C^n_{ij} \tag{13}$$

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 required link capacity 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 like $L^n_{ce} = 1$ 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 $n = \text{peak}$, and in the interval $n = \text{off-peak}$ we introduce a constraint to ensure that circuits can only be switched-off:

$$X^{\text{off-peak},t}_{xy} \leq X^{\text{peak},t}_{xy} \tag{14}$$

Similarly, in the switch-on scheme, the initial time period is given by $n = \text{off-peak}$, and during $n = \text{peak}$, we allow only switch-on by making sure that the number of circuits provisioned during the peak interval is always greater or equal to the number of circuits provisioned during the off-peak interval. This constraint is mathematically formulated exactly the same as (14).

IV. Performance Evaluation

In our performance study, we used the NSFNet topology (14 nodes, 40 unidirectional links) to compare the performance of the switch-on and switch-off energy saving schemes using the Gurobi Optimizer [13] to solve our ILP models on a 2.5 GHz Quad-Core desktop CPU. In this study, the initial IP and the transport network topology is assumed to be the same, and the transport network always uses the shortest path to provision circuits for the IP network. We take a representative traffic matrix [14] and scale this matrix to obtain the peak and off-peak traffic profiles. Constant parameters used in the analytical model are presented in Table II(b). Note that β is only necessary to dimension the value of the routing penalty parameter with respect to the scale of the power consumption in the network, and it must be customized to adapt to different data sets.

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

(a) Power Consumption

Parameter	Value
w_{xy}	$1, \forall x, y \in V$
α	0.7
β	0.1
σ	0.5
$(\tau^{\text{peak}}, \tau^{\text{off-peak}})$	(16, 8)h

(b) Constant Parameters

Case	π^{lo}_{sd}	π^{hi}_{sd}
1	$[0.4, 0.8]$	$[1.0, 1.5]$
2	$[0.4, 0.8]$	$[1.5, 2.0]$
3	$[0.4, 0.8]$	$[2.0, 2.5]$
4	$[0.4, 0.8]$	$[2.5, 3.0]$
5	$[0.4, 0.8]$	$[3.0, 3.5]$

(c) Scaling Factors

TABLE II
Parameters used in the Performance Evaluation

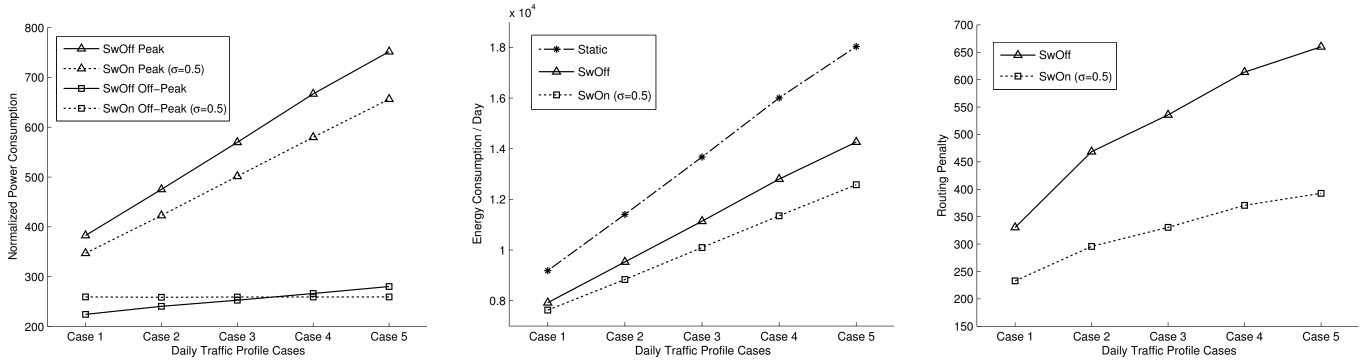

(a) Power Requirements at Peak and Off-peak loads for different traffic profiles with increasing difference between peak and off-peak loads.

(b) Total Energy Requirements for static networks and networks employing switch-on and switch-off schemes for different traffic profiles.

(c) Routing Penalty for networks employing switch-on and switch-off schemes for different traffic profiles with increasing difference between peak and off-peak loads.

Fig. 3. Results of our first performance study

We assume three different bandwidth rates for IP ports, 10 Gb/s, 40 Gb/s and 100 Gb/s. For a general energy model, the power consumption depends on many parameters, e.g., technology, performance, utilization, applications a.s.o. [15]. However, the power consumption in the electrical components usually exceeds that of optical components many times over. Therefore, we used a simplified energy model in order to reduce the computational complexity, where the power consumption of a link simply depends on the IP port size and the number of optical hops (see Table II(a)).

In order to compare the performance of the two schemes, we made two primary performance studies: In the first study, we compared the performance of the switch-on scheme with a 50% penalty on routing changes ($\sigma = 0.5$) for different traffic load scenarios, while in the second study, we compared the performance of the switch-off scheme against the switch-on scheme with varying penalty. (Please note again that $\sigma = 0$ for all *switch-off* optimizations, since that scheme does not consider routing changes.) The scaling function is defined as

$$\lambda_{sd}^{\text{peak}} = \pi_{sd}^{\text{hi}} \cdot \lambda_{sd} \quad \text{and} \quad \lambda_{sd}^{\text{off-peak}} = \pi_{sd}^{\text{lo}} \cdot \lambda_{sd}$$

with the traffic scaling factors π_{sd}^{hi} and π_{sd}^{lo} including random noise which let them vary in a ranges shown in Table II(c). Note that we keep the scaling factor range for the off-peak loads constant and increase the scaling factors during the peak loads to vary the difference between the peak and the off-peak loads in the network.

The average power requirements as computed by the ILP for both the switch-on and switch-off scheme during peak and off-peak intervals for different traffic scenarios (Case 1 to Case 5) is shown in Fig. 3(a). In this study, the switch-on scheme has a routing penalty of $\sigma = 0.5$. It can be seen from the graph that during off-peak hours, the switch-off and switch-on schemes have similar power consumption, with the switch-off scheme requiring marginally lower power in traffic scenarios with smaller peak load scaling factors and requiring higher power in traffic scenarios with higher peak-load scaling factors. This is due to the fact that in the switch-off scheme, the topology is initially optimized for the peak traffic interval, and in scenarios with high peak traffic (e.g., Case 5), the topology uses larger interfaces. The use of large interfaces while optimal for peak loads reduces flexibility in bandwidth reduction leading to higher power consumption during off-peak loads.

It may surprise that the switch-on scheme outperforms the switch-off scheme even under peak load conditions, since the basic time interval for switch-off is defined as the one with peak traffic. However, in contrast to the switch-on scheme, the topology itself (i.e., the adjacecy of the nodes) is fixed to the *initial topology* and only the link capacities are optimized (see Fig. 1(a)). In the switch-on scheme on the other hand, only the off-peak topology is fixed to the initial topology (like shown in Fig. 1(c)), which results in a significantly higher connectivity during the peak traffic interval (see Fig. 1(d)). This increase of the nodal degree in the topology in turn leads to a lowered average IP hop-count, i.e., a lower average number of optical-electrical-optical conversions and thus to a lower total power consumption.

The energy savings are highlighted in Fig. 3(b), which show the total energy requirements per day for the two schemes as well as for a fully over-provisioned (static) network that uses no load adaptive scheme at all. It is clear from this graph that both schemes can lead to significant power savings in carrier networks, and the energy savings are higher as the difference between the day and night loads becomes higher. It is also evident that the switch-on scheme clearly outperforms the switch-off even while hampering it's energy conservation performance for the benefit of fewer routing changes.

Finally, Fig. 3(c) shows how the parameter σ can significantly reduce the routing penalty (i.e., the total number of necessary routing reconfigurations) in the network. It is clear from the graph that by considering the routing penalty, we significantly curb the number of routing changes, which is evident more prominently in traffic scenarios with high

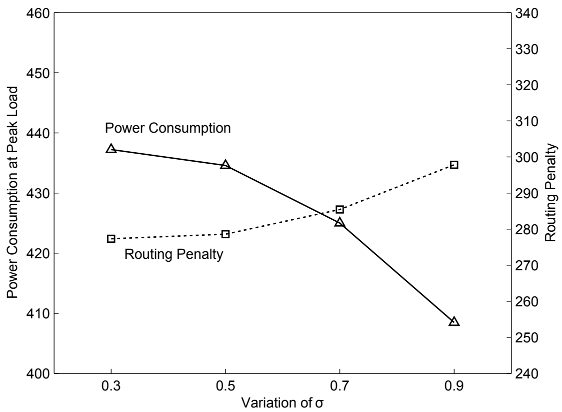

Fig. 4. Trade-off between routing penalty and power consumption at peak loads for switch-on technique.

differences between day and night loads. Here it can be seen that the switch-on scheme has lower than half of the number of routing changes as seen in the switch-off scheme.

The primary drawback of the switch-on scheme lies in the fact that it employs a large number of interfaces of small granularity. As seen in Table III, the switch-on scheme employs almost twice the number of 10 Gb/s interfaces as the switch-off mechanism and a significantly lower number of 40 Gb/s and 100 Gb/s interfaces. While efficient in reducing power consumption, this design may trigger frequent network planning cycles with increase in traffic. Also, frequent upgrade of small capacity interfaces may prove more costly from the CapEx perspective. On the other hand, as the switch-off scheme employs a larger number of high-capacity interfaces, it will not require very frequent planning cycles.

In the second study, we fix the traffic scenario (Case 2 in Table II(c)) and vary the IP routing stability parameter σ to observe the trade-off between the routing penalty and the power requirements for the switch-on scheme. We can see from Fig. 4 that varying the value of σ between 0.3 to 0.5 does not offer any significant change in the power requirements of the network. However, upon further increasing σ, we see that the routing penalty starts increasing and power consumption also comes down as the contribution of the re-routing punishment in the objective function decreases. This trade-off allows carriers to flexibly determine the relative need for network stability as well as power reduction needs and set σ accordingly. It should be noted here that while the parameter σ helps us to define a trade-off between the routing penalty and the power consumption, it does not actually constrain the actual number of routing changes that are made in the network. Another important fact to observe here is that the difference

Mechanism	10 Gb/s	40 Gb/s	100 Gb/s
Switch-Off	42.8	22.5	1.6
Switch-On	88.1	3.25	0

TABLE III
INSTALLED INTERFACES FOR TRAFFIC PROFILE CASE 5

in the number of routing reconfigurations between $\sigma = 0.3$ and $\sigma = 0.9$ is much smaller than the difference between the switch-on and the switch-off schemes, indicating that the routing changes here are significantly lower even when we provide significant flexibility to the switch-on scheme.

V. CONCLUSION

In this paper, we presented a simple case study on energy efficiency in IP-over-WDM networks with dynamic circuit capability, with two basic schemes referred to as *switch-on* and *switch-off*. The switch-on technique for energy conservation was proposed as a fundamentally new paradigm based on dynamic circuit service capability for the IP layer. Unlike the switch-off scheme, the switch-on scheme does not affect the path redundancy in the network and our results show that the switch-on scheme can reduce the power consumption, the total installed capacity as well as the number of routing reconfigurations required in the network. While this seems ideal, our results also show that the switch-on scheme uses a large number of small capacity interfaces which may not be suitable from a network planning perspective as it might require frequent capacity upgrades. Also, while the number of routing reconfigurations can almost be halved, they are still significant in number. To address this issue, new approaches to IP-over-WDM offloading are needed to find a trade-off between routing stability and energy efficiency. Our study is the first step towards this goal.

REFERENCES

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

[2] O.Tamm, C.Hermsmeyer, A.M.Rush: *"Eco-sustainable system and network architectures for future transport networks,"* Bell Labs Technical Journal, Vol.14, pages 311-327, 2010

[3] Internet2 Headroom Practice: https://wiki.Internet2.edu/confluence/download/attachments/17383/Internet2+Headroom+Practice+8-14-08.pdf

[4] L.Chiaraviglio, et al.*"Reducing Power Consumption in Backbone Networks,"* ICC 2009, Dresden, June 2009

[5] L.Chiaraviglio, et al. *"Energy-Aware Networks: Reducing Power Consumption By Switching Off Network Elements"*, National Telecommunications and Information Theory Group Meeting, Italy, June 2008

[6] K.Thompson, G.J.Miller, R.Wilder: *"Wide-Area Internet Traffic Patterns and Characteristics"*, IEEE Network, vol.11, no.6, pp. 10-23, Dec 1997

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

[8] B.Puype, et al. *"Power reduction techniques in multilayer traffic engineering"*, ICTON 09, July 2009

[9] M.Caria, M.Chamania, A.Jukan, *"Trading IP Routing Stability for Energy Efficiency: A Case for Traffic Offloading with Optical Bypass"*, ONDM 2011, Bologna, Italy, February 2011

[10] K. H. Liu et al. , *"Experimental Study of Dynamic IP Topology Reconfiguration in IP/WDM Networks,"* IEEE GLOBECOM, 2001.

[11] L. Cheng, et al.*"Network Engineering - Control of Dynamic Link Topology in User Networks,"* Bell Labs Technical Journal, Vol. 8, No 1, pages 207-218, July 2004.

[12] M.Chamania, M.Caria, A.Jukan, *"A Comparative Performance Analysis of IP Traffic Offloading Schemes over Dynamic Circuits"*, INFOCOM Mini Conference, 2011

[13] Gurobi Optimizer 3.0, http://www.gurobi.com

[14] B. Mukherjee, *Optical WDM Networks*, Springer, 2006, pp.422

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