Energy-Efficient Resilient Optical Networks: Challenges and Trade-offs

Yabin Ye, Felipe Jiménez Arribas, Jaafar Elmirghani, Filip Idzikowski, Jorge López Vizcaíno, Paolo Monti, Francesco Musumeci, Achille Pattavina, and Ward Van Heddeghem

Yabin Ye and Jorge López Vizcaíno are with Huawei Technologies Duesseldorf GmbH.

Felipe Jiménez Arribas is with Telefónica I+D.

Jaafar Elmirghani is with the University of Leeds.

Filip Idzikowski is with Poznan University of Technology.

Paolo Monti is with KTH Royal Institute of Technology.

Francesco Musumeci and Achille Pattavina are with Politecnico di Milano.

Ward Van Heddeghem is with Ghent University — iMinds.

ABSTRACT

Energy efficiency and resilience are two well established research topics in optical transport networks. However, their overall objectives (i.e., power minimization and resource utilization/availability maximization) conflict. In fact, provisioning schemes optimized for best resilience performance are in most cases not energy-efficient in their operations, and vice versa. However, very few works in the literature consider the interesting issues that may arise when energy efficiency and resilience are combined in the same networking solution. The objective of this article is to identify a number of research challenges and trade-offs for the design of energy-efficient and resilient optical transport networks from the perspective of long-term traffic forecasts, short-term traffic dynamics, and service level agreement requirements. We support the challenges with justifying numbers based on lessons learned from our previous work. The article also discusses suitable metrics for energy efficiency and resilience evaluation, in addition to a number of steps that need to be taken at the standardization level to incorporate energy efficiency into already existing and well established protocols.

INTRODUCTION

Traffic in core networks has been growing at a yearly rate of 45 percent since 2004 [1]. The ever increasing traffic levels have forced operators to focus on ways to increase the network capacity accordingly. However, such an increase is expected to bring higher energy consumption. The worldwide electricity consumption of communication networks was estimated to be 350 TWh in 2012 (or nearly 2 percent of the worldwide electricity consumption), showing an average annual growth rate of 10 percent since 2007 [2]. Therefore, energy efficiency is becoming one of the key design parameters for planning and operating today's telecommunication networks. In this regard, optical transport solutions might be beneficial because they reduce the number of opto/electrical/opto (O/E/O) operations, which are very costly from power consumption point of view.

A critical issue when dealing with optical transport networks is the amount of traffic that can be disrupted by a network fault, an aspect that needs to be addressed properly if certain quality of service (QoS) levels are to be guaranteed to the end user. This makes it necessary to reserve some redundant resources to ensure *network resilience* (i.e., the ability of the network to provide and maintain an acceptable level of service in the face of different faults). These redundant resources are then used to reroute traffic, bypassing the failed network element(s) over the so-called backup path(s).

In general, resilient schemes can be divided into two categories: *restoration* or *protection*. Their main difference lies in the way they compute the backup paths. In a restoration approach, the backup paths are computed on the fly (i.e., only after a failure occurs), while in a protection approach the backup paths are precomputed, with the backup resources already reserved and ready to be used in the occurrence of a failure. Each scheme has its pros and cons. Restoration allows for more efficient use of network resources, but has longer recovery time and does not provide a 100 percent recovery guarantee against failures. Protection schemes, on the other hand, have a faster recovery time and can guarantee 100 percent recovery from the failure scenario for which they were designed, but require more energy in absolute terms. Protection schemes can be further categorized in the way backup resources are used. *Dedicated schemes* do not allow sharing of backup resources among multiple backup paths. *Shared schemes*, under specific conditions, allow some backup paths to use the same wavelength resources. Dedicated protection (DP) can be implemented in a 1+1 or 1:1 fashion. In the former, the traffic is duplicated over two disjoint paths, and one of the paths is selected by the destination node, whereas in the latter, the disjoint backup path may carry traffic that can be preempted in the case of a failure happening in the working path.

Operators tend to implement the resilience concepts presented above following three high-level mechanisms. The first one is network overprovisioning, that is, from the duplication of a number of line cards in a node, up to the dupli-

 0163-6804/15/$25.00

cation of completely geographically separated routers, nodes, and fiber links. In the latter case, the end result will be a dual plane network at one or several layers. The second mechanism is to use 1+1 protection at the optical layer, which can provide fast switchover time (< 50 ms) when a failure occurs. On the other hand, when the backup path carries the switched-over traffic (and the failed element along the primary path is under repair), this traffic is in a vulnerable state: it is unprotected against any additional failures. For this reason, while the traffic is carried by the backup path, restoration is used to replace the old (failed) working path. This is the third option (i.e., dedicated protection plus restoration). After a new backup path is computed, the network is again protected against any single (link or node) failure scenario, with the original backup path serving as the current working path, and the repaired (old) working path now acting as the backup path. Resilient networks rely on a number of duplicated resources, which are unused most of the time. This is obviously not the most energy-efficient solution.

Core networks have always been designed focusing on minimization of the network equipment and deployment cost, that is, capital expenditures (CAPEX). Given a certain request of data traffic, the definition of a network topology, and the set of equipment to be installed, a design procedure determines the processing capability of the switching/routing nodes on one hand and the capacity of the inter-node transmission links on the other hand. However, reducing operational expenditures (OPEX) is becoming increasingly important for operators, and energy consumption is one of the key contributors. Despite numerous studies focused on resilience vs. cost or QoS, energy consumption cannot simply be interchanged with cost and QoS, so previous techniques cannot be directly applied to the problem of energy-efficient resilient optical network design. First, energy efficiency has to be traded against QoS. A simple cost minimization approach may save energy by shaping traffic over fewer routes, but on the other hand reduce network reliability due to the increased impact of a cut in a consolidated route. Second, network devices implementing sophisticated energy saving functionalities could initially be more expensive than the conventional ones. Finally, backup (and working path) routes may be designed to use time-varying renewable energy sources efficiently and/or adapting the transmission to the instantaneous traffic demand, which reduce the network carbon footprint, but may not reduce or minimize the costs directly.

Two facts should be considered for energy-efficient resilient design:
- Core networks are dimensioned assuming peak traffic levels.
- The type of protection (i.e., service guarantees) can be differentiated and adapted to the specific traffic type.

Both observations provide opportunities to reduce energy consumption. As for peak traffic dimensioning, real traffic demands vary over time (i.e., the traffic demands at night are usually much lower than those during the day). This means that unused resources (line cards in nodes or even whole nodes) can be put to sleep to save energy. As for the service guarantees, not all traffic types need the same level of protection; therefore, resource redundancy for service guarantees can be restricted to a subset of traffic types, thus leading to further energy savings. In other words, in order to improve the energy efficiency of resilient optical core networks, there is a clear need for design and provisioning strategies specifically tailored to reduce network energy consumption.

Certainly, optical network design is complicated and needs to balance a number of metrics (e.g., cost, energy, resilience, and scalability). The purpose of this article is to identify the major research challenges in the relatively new field of energy-efficient resilience in optical transport networks. In this respect, our objective, rather than presenting uneqivocal solutions, is to highlight and reason about the nature of the potential options and the challenges that they present; that is, their performance trade-offs.

Research Topics/Open Problems

Core networks are generally designed and provisioned in response to three main inputs: long-term traffic forecasts, short-term traffic dynamics, and service level agreement (SLA) requirements. These factors influence the design and choice of protection and restoration mechanisms as well as the energy efficiency performance. In this section, we provide an overview to address these long-term static network architecture design choices, the protocols and hardware functionalities needed for adaptation to the short-term traffic dynamics, and the strategies needed to meet the SLA requirements in an energy-efficient fashion. Most of the presented challenges are justified by numerical results achieved in our previous work (although in different network and traffic scenarios).

Long-Term Static Architecture Network Design Choices

Network Architecture — Current optical networks are based on wavelength-division multiplexing (WDM) technologies, which operate over fixed channel spacing defined by the International Telecommunications Union —Telecommunication sector (ITU-T). WDM network architectures may consider single line rate (SLR) transmission in all the wavelength channels, or mixed line rate (MLR) transmissions, where channels transmitted with different rates coexist on the same fiber. Elastic optical networks (EONs), which allow for flexible-bandwidth transmissions and adaptive modulation, have emerged as the future technology for the optical network. The choice of a particular network architecture over another, together with the adopted resilient scheme (i.e., restoration vs. protection, or dedicated vs. shared protection), will have clear effects on energy consumption. For instance, changing the protection scheme from DP 1+1, the most reliable and more energy-consuming scheme, to more energy-efficient ones such as DP 1:1 or shared protection (SP), will result in different energy savings depending

Core networks are generally designed and provisioned in response to three main inputs: long-term traffic forecasts, short-term traffic dynamics, and service level agreement requirements. These factors influence the design and the choice of protection and restoration mechanisms as well as the energy efficiency performance.

Core networks are typically organized as multilayer networks where each layer can be viewed as a single network. Higher layers rely on the resources and services provided by the lower layers, and resilience can be provided at any layer.

on the network architecture. For example, adopting SP or DP 1:1, the energy consumption of the optical layer can be reduced up to 49 percent with respect to DP 1+1 in EON and up to 42 percent in WDM with a SLR of 40 Gb/s [3].

Embodied Energy — The energy consumed during the whole lifetime of an installed device needs to be considered (including, e.g., manufacturing and decommissioning) in the design of optical transport networks. In fact, the embodied energy accounts for approximately 70 percent of the total energy consumption of the network (i.e., operational plus embodied energy consumption) [4]. The fiber actually has the highest impact in terms of embodied energy [4] due to the large amount of material used to protect the fibers in a cable and the hundreds of thousands of kilometers of fiber cable in a core network. Therefore, from an energy perspective, it is especially important to minimize the number of redundant fiber cables in resilient networks.

Resilience at Different Layers — Core networks are typically organized as multilayer networks where each layer can be viewed as a single network. Higher layers rely on the resources and services provided by the lower layers, and resilience can be provided at any layer. Figure 1a illustrates 1+1 protection in the IP layer, whereas Fig. 1b shows the same traffic demand protected at the optical layer. Note the subtle difference of the latter with the employment of optical bypass in combination with 1+1 protection at the IP layer shown in Fig. 1c. If protection is provided simultaneously at multiple layers and there is no coordination between the layers, parallel recovery actions may take place, which can have a significant impact on the overall network stability and lead to suboptimal resource usage. Multilayer recovery schemes have been explored in earlier research; however, it is not yet clear how these strategies perform from an energy efficiency point of view. Intuitively, the best option is to provide protection at the (most energy-efficient) bottom layer, perhaps augmented with a bottom-up escalation strategy where recovery starts in the lowest detecting layer and escalates upward. However, this strategy comes with the drawback of poorer handling of higher-layer failures, and the impact of escalation timings on the overall recovery time has to be assessed. Furthermore, network operators often duplicate their backbone networks and simultaneously operate two nearly identical networks. With 1+1 protection often used at the optical transport network (OTN) layer as shown in Fig. 1d, this means in effect that each requested demand results in four times the capacity at the optical layer. Energy-wise, this is clearly not an optimal situation.

Topologies — Looking at a single layer, the network topology should keep a relatively high number of alternative paths for the traffic demands coming from the higher layers. Although a full-mesh highly overprovisioned topology is ideal from the perspective of protection, the number of deployed links and their capacities should be traded against energy consumption. This trade-off was considered in [5] for physical topology design. The constraint on minimum nodal degree equal to 2 (securing at least one alternative path) leads to negligible increase of network power consumption (0.5 percent increase with respect to a network with minimum nodal degree equal to 1 using a multihop bypass approach and symmetric traffic demand for the NSFNET network). Furthermore, the size of the network will also play an important role when trading energy efficiency and resilience. For instance, if some links become highly occupied due to traffic growth, backup paths may become much longer and require the use of signal regeneration or a more robust modulation format, with the consequent increase in energy consumption.

Open Issues in Adaptation to Short-Term Traffic Dynamics

Novel Equipment Features — The nominal power consumption of devices/equipment considered for installation must be assessed during the network design phase. Emerging equipment features, such as sleep modes and dynamic reconfiguration of modulation format and transmission reach, could provide new opportunities to reduce the power consumption of backup links. For example, more energy-efficient protection schemes could be devised to exploit sleep modes at the expense of slower recovery. However, to the best of our knowledge, there is no current information of any commercial device implementing these innovative features, meaning all related work reporting results is based on assumptions. New protection schemes could also allow multi-rate transponders to fall back to lower transmission rates depending on the SLA and reduce energy consumption. This is related to the next section, where we discuss quality of protection (QoP) differentiation. Moreover, there are open issues which must be addressed:

1 Are complex devices more prone to failures?
2 How quickly can the devices be brought up and down to react to traffic variations?
3 How can the energy consumption information of the device and electricity cost be accurately monitored?

Trade-offs in Dynamic Adaptation to Temporal Variation of Traffic — Traffic in core networks is usually higher during the day and lower during the night. Putting idle devices into sleep or energy-saving mode (e.g., adaptation of transmission rates) can effectively reduce the energy consumption of the backup resources (e.g., adapting the backup transmission in DP 1+1 can save up to 22 percent of energy consumption at the optical layer during the low-traffic hours of the weekend [3]). Another strategy relies on concentrating backup paths into separate fibers in order to be able to put the devices on these links into sleep mode without being constrained by the presence of working paths. This is applied to DP [6] and SP [7] with power savings of up to 35 percent with respect to non-sleep-mode scenar-

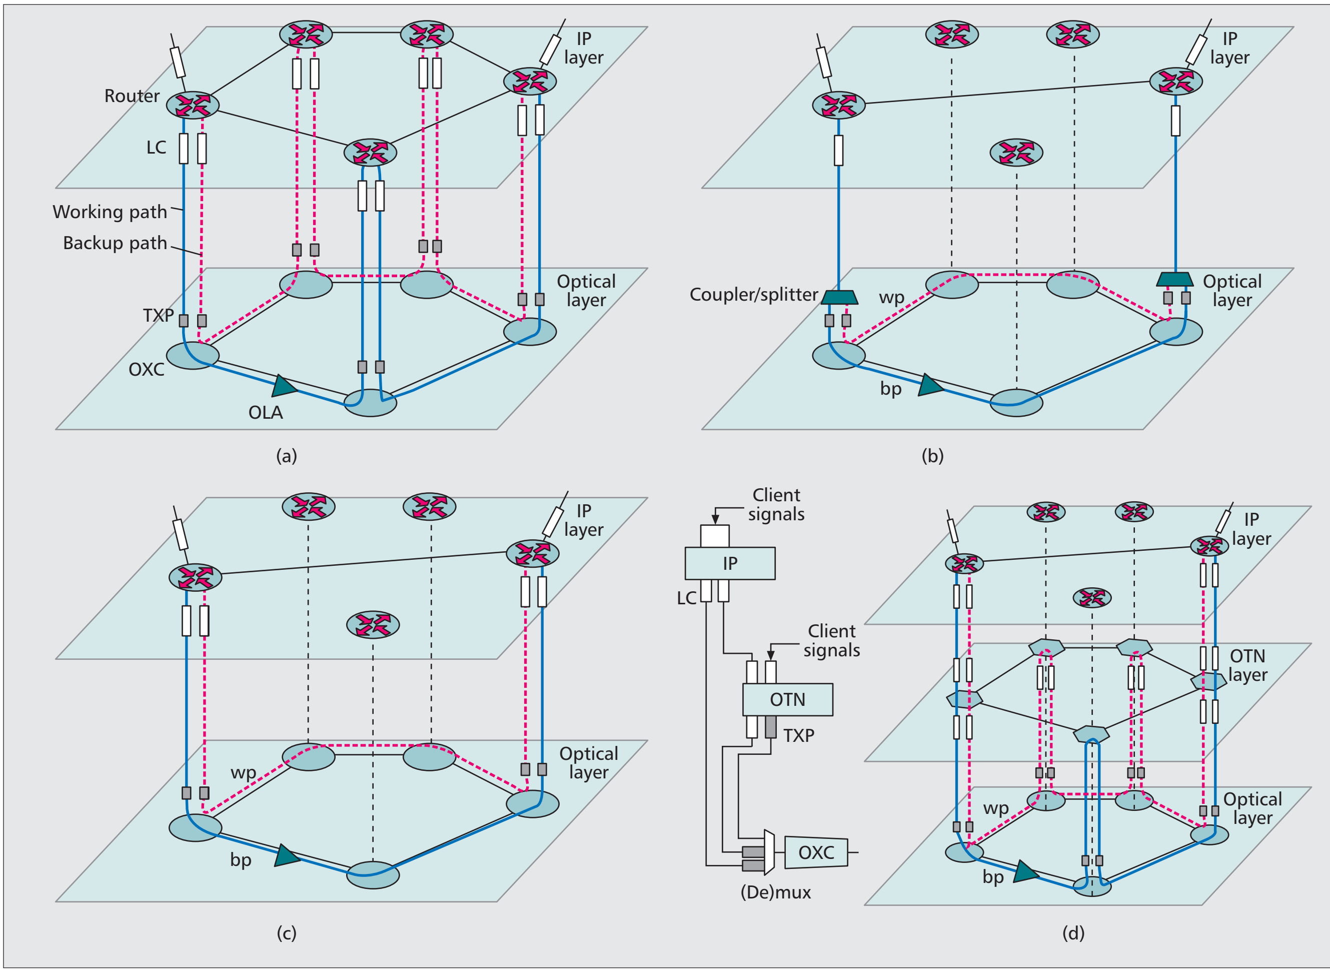

Figure 1. a) 1+1 protection at the IP layer; b) 1+1 protection at the optical layer; c) employing optical bypass at the IP layer; d) grooming at the OTN layer. LC: line card, TXP: transponder, OXC: optical cross connect, OLA: optical line amplifier, wp: working path, bp: backup path.

ios. Moreover, the sleep mode functionality can also be exploited considering grooming at the electrical layer with energy savings of up to 60 percent in both the optical and IP layers [8]. Nevertheless, for these mechanisms, it will be important to determine the best trade-off between resilience, energy efficiency, and low reconfiguration costs.

Geographical Traffic Distribution — Geographical traffic distribution in networks covering large areas may be diverse, especially in networks spanning several time zones. Is it possible to use idle resources from one time zone as protection resources in another time zone? Traffic distribution can also suffer strong variations due to external influences like natural disasters or wide audience events. Extra redundancy can be provided to face the potential traffic changes, but it will be costly. Therefore, how to handle this potential variation from the protection point of view is an open question. There are some works that differentiate traffic in different time zones [5] or use traffic data originating from measurements [9]. However, to the best of our knowledge, there is no related work focused on geographical traffic distribution and green networking.

Granularity of Traffic Demands — The size of a single traffic demand determines the possibility of grooming it into a set of lightpaths and links to utilize the available network resources in an efficient way. However, grooming of multiple traffic demands over a limited number of network resources poses a challenge from the resilience point of view, since some traffic demands need to traverse multiple hops from their origin to their target. This, in turn, increases the number of devices on their path, each being prone to failures. This leads to important questions related to the maximum granularity of traffic demands (with respect to network resources) so that traffic grooming for energy saving pays off without violation of protection constraints. While an initial exploration on this issue [9, Fig. 7] indicates that grooming can be more energy-efficient when the granularity of the traffic demands is below half of the line rate, the cited work does not look at the impact of grooming on availability (in terms of relative uptime per year).

Strategies to Meet the SLA Requirements

Physical Impairments — Network design techniques have to consider physical impairments (e.g., fiber loss, dispersion, or nonlinear effects)

Preemption is based on the intuition that low-priority services can be provisioned with the option to be discarded should a connection with a higher criticality level need to use the network resources. This could happen after a failure concatenation, or some unexpected or very infrequent event.

and their impact on the resulting quality of transmission (QoT). Some energy saving techniques can exacerbate the effect of some physical impairments (e.g., cross-phase modulation and crosstalk) [10], as they tend to concentrate most of the traffic on a few links to put unused devices into sleep mode. This problem can be overcome by introducing design techniques that are both energy- and impairment-aware. The work [10] shows that such techniques are able to achieve the same energy savings of conventional (i.e., non-impairment-aware) green strategies while providing QoT levels close to those provided by the impairment-aware design strategies (i.e., only a small percentage of the lightpaths cannot be established due to insufficient QoT). Even though the work in [10] considers an unprotected scenario, we can expect that this aspect will be even more critical when designing protected networks, where the backup paths are on average longer than their respective working paths. Power transient phenomena will also have to be considered when setting network elements into sleep mode and adequate mechanisms (e.g., disabling the transponder electronics while keeping the injected optical power) must be provided. Moreover, the continuous development of high-speed systems comes at the cost of reduced reach. Consequently, a higher number of O/E/O regenerations (highly energy consuming processes) would be needed for a generalized deployment. Employing more robust modulation formats for long backup paths can be an alternative to regeneration, but at the expense of requiring higher energy per bit [3].

Differentiated Quality of Protection — Commonly, a single policy is applied to all aggregated traffic demands as shown in Fig. 1. An alternative to this policy would be to apply differentiated QoP by assigning different resilience levels to demand "subsets" with different SLA needs [11]. These subsets could be allocated, for instance, to different residential services or to different performance classes of a corporate service (e.g., optical virtual private networks, VPNs). The application of differentiated QoP allows for a reduction of protection resources, which may result in improved energy efficiency with respect to the conventional DP 1+1 scheme (energy efficiency per gigahertz, EEPG, improvements of up to 300 percent are reported in [11]). Another possibility is to adapt the reliability performance of a given protection scheme to the reliability requirements of the provisioned traffic, that is, the differentiated reliability (DiR) concept. Thus, if for some traffic demands the backup path does not always need to be available for any possible failure scenario, it will be possible to selectively assign protection resources only to those demands that need them the most. This approach can lead to significant energy savings (i.e., up to 25 percent [12]) with respect to conventional protection strategies that are always 100 percent survivable.

EON/BVT — EON technologies provide an extra degree of freedom in assigning traffic demands to QoP levels. For example, a sliceable bandwidth variable transponder (BVT) based on a multicarrier approach could be used to assign a number of carriers to a specific QoP. This can allow the optimizing of resource utilization and providing a finer granularity on QoP mapping. Besides, a BVT can adapt the transmission rate (modifying the bandwidth and/or modulation format) to the traffic variations, and thus reduce energy consumption. EON is shown to provide much higher EEPG than traditional fixed-grid WDM technologies with any protection scheme (e.g., EEPG can be improved up to 200 percent with respect to MLR with SP [3]).

Network Virtualization — Network resources can be partitioned over different operators and/or services. In that respect, the physical network would comprise a number of virtual networks, each receiving a different degree of resilience according to the amount of resources allocated to each service/operator. The virtualization of the network components combined with the creation of segregated virtual topologies with different architectures (including different routing and resilient protocols) can help support variable QoP levels and thus reduce energy consumption by sharing the resources among different users/virtual networks. The differentiation can be achieved in one or more dimensions including time delay, availability, protection scheme/protection level, QoT, and the network layer selected (i.e., electrical and/or optical). Despite numerous works that consider realization of a virtual network over a physical one, only a few of them consider network virtualization and protection. For example, 28 percent of power can be saved by allowing virtual links to bypass physical nodes in an IP-over-WDM network (multihop bypass compared to non-bypass approach for NSFNET network at 6:00 pm [5, Fig. 5], providing protection in the physical layer by keeping minimum nodal degree equal to 2. However, protection is not the focus of [5].

Data Preemption — Preemption is based on the intuition that low-priority services can be provisioned with the option to be discarded should a connection with a higher criticality level need to use the network resources. This could happen after a failure concatenation, or some unexpected or very infrequent event. In this case, lightpath preemption levels are used to decide which signal should be allocated first, releasing the resources in use by lightpaths with lower preemption levels. We can consider this mechanism as a subset of a QoP scheme. This, of course, has to be done without compromising the required QoS level of each connection. Preemption-based strategies can improve energy efficiency, since powering-on extra resources is not always necessary. For instance, in a scenario where preemption is used to dynamically provision the service subject to different QoP requirements, redundant resources used for protection could be used to convey extra traffic if the carried traffic is able to tolerate a disruption, should one or more failures occur in the network. This will be accomplished without unnecessarily turning on extra resources in the network, with evident benefits in terms of energy efficiency. How effective the use of such energy-efficient preemption techniques is still an open question, which, to the best of our knowledge, has not been studied so far. The answer will depend on a series of factors such as the protection mechanism used and the traffic composition, just to name a couple. Higher power reduction can be expected in the presence of traffic with highly diverse survivability needs, as well as with protection techniques that are less stringent in their requirements. Finally, there is one last aspect to consider. If backup resources are used to route preemptable traffic, they need to stay in an active state, potentially nullifying the benefits of some energy-efficient protection techniques. There is obviously a trade-off that needs to be assessed here to understand under which conditions preemption will be beneficial from an energy reduction point of view.

METRICS AND STANDARDS

This section considers how issues related to energy efficiency and resilience can be evaluated and what this would mean in terms of standardisation requirements.

Metrics — The energy efficiency of a network with various resilience schemes can be assessed with common metrics such as Watts per bit per second or Joules per bit, or the inverse of both. With this approach, less protective schemes (e.g., SP) will result in more energy-efficient designs than more protective schemes (e.g., DP 1+1). While this metric is a good indicator of the actual power or energy required to transport a bit of information, it does not take into account the level of protection. To do so, it might make sense to express the energy efficiency normalized by the protection factor, where the protection factor could be 2 for 1+1 DP, 1.5 for SP, and so on. This metric provides a fairer indication of the energy required for a given level of protection. Finally, with the potential evolution toward more flexi-grid equipment, the common Watts per bit per second metric (or its derivatives) might be insufficient to capture the efficiency in utilizing the available spectrum. For example, a BVT can make use of adaptive bandwidths depending on the required transmission rate and distance. Therefore, it might make sense to take into account the spectral efficiency of transmission with an energy efficiency per gigahertz metric, as proposed in [3].

Standards — Putting unused network devices into sleep mode is one way to save power. However, current Internet protocols operate based on the assumption that network elements are always on. The application and service reinitializations when these devices wake up again would potentially result in a non-negligible amount of signaling overhead [13]. Modifications to the control plane considering link or node removal should provide the ability to choose the level of redundancy available after the network topology has been trimmed. The complete removal of nodes or links from the network topology has several impacts on the control plane that must be considered [14]. For example, it is essential to modify the network topology so that the removed links or devices are not used to forward traffic remembering that such links exist, possibly including the neighbors and destinations reachable through those links or devices. One solution to this sleep mode problem could be based on the use of a proxy [13]. Before going into sleep mode, a node delegates its functionalities to a proxy, which will then respond to routine network traffic on behalf of the sleeping node and will wake the node up when needed. The protocols and procedures for proxy operation such as discovery, selection, delegation, and wake-up have to be defined. Another example would be to require that nodes can negotiate timeouts (in protocols that make use of timeouts), so a node might be able to go into sleep mode or attempt to synchronize periodic messages across a number of protocols. Thus, all these messages fall into a certain timeframe, and in between the node can sleep. The issues described above can also be addressed by designing sleep-compatible protocols or extending existing protocols (where possible) to include the ability to distinguish sleeping elements from failed ones. Some extensions required in existing generalized multiprotocol label switching (GMPLS), Open Shortest Path First (OSPF) routing, Resource Reservation Protocol (RSVP) signaling, and Link Management Protocol (LMP) are proposed in [15] to support energy-efficient traffic engineering technology.

If backup resources are used to route preemptable traffic, they need to stay in an active state, potentially nullifying the benefits of some energy-efficient protection techniques. There is obviously a trade-off that needs to be assessed.

CONCLUSION

Both energy efficiency and resilience in telecommunications networks are well established topics in the research community. However, the combination of both for energy-efficient resilient optical network design is still a relatively new research field. In this article, we identify the corresponding major research challenges and performance trade-offs for core networks from three different aspects: long-term traffic predictions (including network architecture, embodied energy, and resilience at different layers and topologies), short-term traffic dynamics (including novel equipment features, trade-offs in dynamic adaption to temporal variation of traffic, geographical traffic distribution, and granularity of traffic demands), and SLA requirements (including physical impairments, differentiated QoP, EON/BVT, network virtualization, and data preemption). All these factors need to be considered as they influence not only the design and the choice of resilience mechanisms, but also the energy efficiency performance.

New metrics need to be used for energy efficiency resilience evaluation by either considering the energy efficiency with a protection factor or the utilized optical spectrum. From a standardization point of view, existing protocols need to be extended for energy-efficient resilient optical networks using sleep mode devices.

ACKNOWLEDGMENTS

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 257740 (TREND) and grant agreement no. 318137 (DISCUS), and

New metrics need to be used for energy-efficiency resiliency evaluation by either considering the energy-efficiency with a protection factor or the utilized optical spectrum. From a standardisation point of view, existing protocols need to be extended for energy-efficient resilient optical networks using sleep mode devices.

from the Polish National Science Center (decision DEC-2014/12/S/ST7/00415).

References

[1] E. Le Rouzic, "Network Evolution and the Impact in Core Networks," *Proc. ECOC*, Turin, Italy, Sept. 2010.

[2] S. Lambert *et al.*, "Worldwide Electricity Consumption of Communication Networks," *Optics Express*, vol. 20, no. 26, Dec. 2012, pp. B513–24.

[3] J. López Vizcaíno *et al.*, "Protection in Optical Transport Networks with Fixed and Flexible Grid: Cost and Energy Efficiency Evaluation," *Opt. Switching and Networking*, vol. 11, part A, Jan. 2014, pp. 55–71.

[4] X. Dong *et al.*, "Energy-Efficient Core Networks," *Proc. ONDM*, Colchester, U.K., Apr. 2012.

[5] X. Dong *et al.*, "On the Energy Efficiency of Physical Topology Design for IP Over WDM Networks," *IEEE/OSA J. Lightwave Technology*, vol. 30, no. 12, June 2012, pp. 1931–42.

[6] A. Jirattigalachote *et al.*, "Dynamic Provisioning Strategies for Energy Efficient WDM Networks with Dedicated Path Protection," *Optical Switching and Networking*, vol. 8, no. 3, July 2011, pp. 201–13.

[7] A. Coiro *et al.*, "Reducing Power Consumption in Wavelength Routed Networks by Selective Switch Off of Optical Links," *IEEE J. Selected Topics in Quantum Electronics*, vol. 17, no. 2, Mar. 2011, pp. 428–36.

[8] F. Musumeci *et al.*, "Energy-Efficiency of Protected IP-over-WDM Networks with Sleep-Mode Devices," *J. High Speed Networks*, vol. 19, no. 1, 2013, pp. 19–32.

[9] W. Van Heddeghem *et al.*, "A Power Consumption Sensitivity of Circuit-Switched versus Packet-Switched Backbone Networks," *Computer Networks*, 2014; http://dx.doi.org/10.1016/j.comnet.2014.09.015.

[10] C. Cavdar *et al.*, "Design of Green Optical Networks with Signal Quality Guarantee," *Proc. IEEE ICC*, Ottawa, Canada, June 2012.

[11] J. López *et al.*, "Energy- and Cost-Efficient Protection in Core Networks by a Differentiated Quality of Protection Scheme," *Proc. ECOC*, London, U.K., Sept. 2013.

[12] A. Muhammad *et al.*, "Reliability Differentiation in Energy Efficient Optical Networks with Shared Path Protection," *Proc. IEEE GreenComm*, Oct. 2013.

[13] R. Winter, S. Jeong, and J. H. Choi, "Towards an Energy-Efficient Internet," http://tools.ietf.org/id/draft-winter-energy-efficient-internet-01.txt, Oct. 2012.

[14] A. Retana, R. White, and M. Paul, "A Framework and Requirements for Energy Aware Control Planes," http://www.ietf.org/id/draft-retana-rtgwg-eacp-00.txt, July 2012.

[15] S. Okamoto, "Requirements of GMPLS Extensions for Energy Efficient Traffic Engineering," http://tools.ietf.org/id/draft-okamoto-ccamp-midori-gmpls-extension-reqs-01.txt, Feb. 2012.

Biographies

YABIN YE (yeyabin@huawei.com) received his B.E. and Ph.D. degrees from Tsinghua University, Beijing, China, in 1997 and 2002, respectively. From 2002 to 2004 he was a research scientist with the Institute for Infocomm Research (I2R), Singapore. From 2004 to 2008 he was a senior researcher with Create-Net, Italy. Currently, he is a senior researcher with Huawei Technologies European Research Center, Munich, Germany. His research interests include optical networking, high-speed transmission technologies, and hybrid optical/wireless access technologies.

FELIPE JIMÉNEZ ARRIBAS (felipej@tid.es) holds an electronic physics degree from UCM. He joined Telefónica I+D in 1994 on a Ph.D. student scholarship, becoming part of the staff in 1995. He started working on the development of the Telefónica Spain nationwide IP network, and since then has been involved in many projects related to access and core architectures as well as in field trials of advanced fiber technologies. He has also worked on relevant EU projects, like the FP6 MUSE, FP7 Accordance, TREND, Discus, and Idealist projects.

JAAFAR ELMIRGHANI [F] (J.M.H.Elmirghani@leeds.ac.uk), professor at University of Leeds, is a Fellow of the IET and Fellow of the Institute of Physics. He received the IEEE Comsoc 2005 Hal Sobol award, the IEEE Comsoc Signal Processing and Communication Electronics outstanding service award, 2009, and an IEEE ICC 2013 Best Paper award. He has been awarded £22 million in grants, published over 400 papers, and has research interests in communication networks. He is an IEEE Comsoc Distinguished Lecturer.

FILIP IDZIKOWSKI (filip.idzikowski@put.poznan.pl) is an assistant professor at Poznan University of Technology (PUT), Poland. He received his Ph.D. (Dr.-Ing.) from Technical University of Berlin (TUB), Germany, in 2014. His main research interests include power consumption of multi-layer core networks, traffic modeling, and routing. Part of this work was carried out when he was with TUB and Sapienza University of Rome, Italy.

JORGE LÓPEZ VIZCAÍNO (jorge.vizcaino@huawei.com) received his Bachelor's degree in telecommunications engineering from the Technical University of Madrid in 2008 and his M.Sc. degree in telecommunications from the Technical University of Denmark in 2011. Since 2012, he has been a researcher at the European Research Center of Huawei Technologies in Munich, Germany, and a Ph.D. candidate at the Technical University of Dortmund, Germany. He has been involved in the EU project TREND related to energy efficiency. His research interests are in network planning, network protocols, energy efficiency, and optical networking.

PAOLO MONTI (pmonti@kth.se) received a Ph.D. (2005) from the University of Texas at Dallas and is now an assistant professor at KTH Royal Institute of Technology. His main research interests are within the networking aspects of all-optical networks, while part of his work has recently focused on the energy efficiency of telecommunication networks with particular emphasis on current and future converged (wired + wireless) broadband access solutions.

FRANCESCO MUSUMECI (francesco.musumeci@polimi.it) received his Ph.D. degree in information engineering from Politecnico di Milano, Italy, in 2013. He is currently a research assistant at the Department of Electronics, Information and Bioengineering of Politecnico di Milano. He is an author of more than 20 papers published in relevant international journals and conference proceedings in the field of optical communication networks. His research interests focus on energy-saving switching and networking, optical core network architectures, content delivery networks, and network resilience.

ACHILLE PATTAVINA (achille.pattavina@polimi.it) received his Dr. Eng. degree in electronic engineering in 1977 and since 1995 has been a full professor at Politecnico di Milano," (Italy). He has been an author of more than 200 papers and two books in the area of communications networks. He has been a coordinator of national and international research activities, including European Union funded projects. His current research interests are in the areas of green ICT, optical networks, switching theory, traffic modeling, and broadband convergent access/metro networks.

WARD VAN HEDDEGHEM (ward.vanheddeghem@intec.ugent.be) received his Master's degree in applied engineering electromechanics from Hogeschool Gent, Belgium (1999), and his M.Sc. degree in computer science engineering from Vrije Universiteit Brussel, Belgium (2009). In 2014 he obtained a Ph.D. degree at the IBCN research group at INTEC UGent/iMinds. His research interests are in the field of environmental impact of ICT and energy-efficient network architectures.