

Design of low-margin optical networks

Yvan Pointurier, *Senior Member, IEEE*

(Invited Paper)

Abstract—We review margins used in optical networks and review a formerly proposed margin taxonomy. For each category of margins, we review techniques that the network designer can use, in order to increase the capacity of optical networks, extend their life, decrease deployment cost (CAPEX) or total cost of ownership over their life duration. Both green field (new network deployments) and brown field techniques (used after initial network deployment) are discussed. The technology needed to leverage the margins and achieve the aforementioned gains are also reviewed, along with the associated challenges.

Index Terms—Networks Optimization; Optical Networks; Circuit-Switched Networks.

I. INTRODUCTION

MARGINS are mandatory to ensure that optical networks support the planned demand capacity at initial deployment (green field) or when deploying a new service or a light path (brown field) virtually error-free during operation over the full network life, which may span several decades. Margins are both demanded by the operators, who want to ensure proper operation of the network, and the network equipment providers, who want to ensure that the delivered network will indeed provide the guaranteed performance, e.g., to meet a Service Level Agreement. However, as in any industry, margins come with a cost as they incur network over-dimensioning, both for green field and brown field scenarios, and thus margin should be minimized.

At the optical layer, the margin of a light path may be quantified as the difference between the actual quality of transmission (QoT) metric (e.g., electrical or (optical) signal to noise ratio (SNR/OSNR), Q^2 -factor, reach, bit error rate) of the signal supporting the light path, and the threshold above which the signal is deemed recoverable “error-free” (i.e., the FEC (Forward Error Correction) limit).

In his seminal paper [2], Augé proposed the following margins taxonomy:

Unallocated margins (*U-margins*) encompass both the capacity and reach margins, i.e., the difference of capacity/reach between the demand and that of the equipment, in particular the optical transponder (TRX) really deployed. U-margins result from the discrete datarate and reach granularity of commercial transmission equipment.

System margins (*S-margins*) account for time-varying network operating conditions. S-margins include fast varying impairments such as polarization effects, and slow varying impairments; the latter are due to either increasing channel

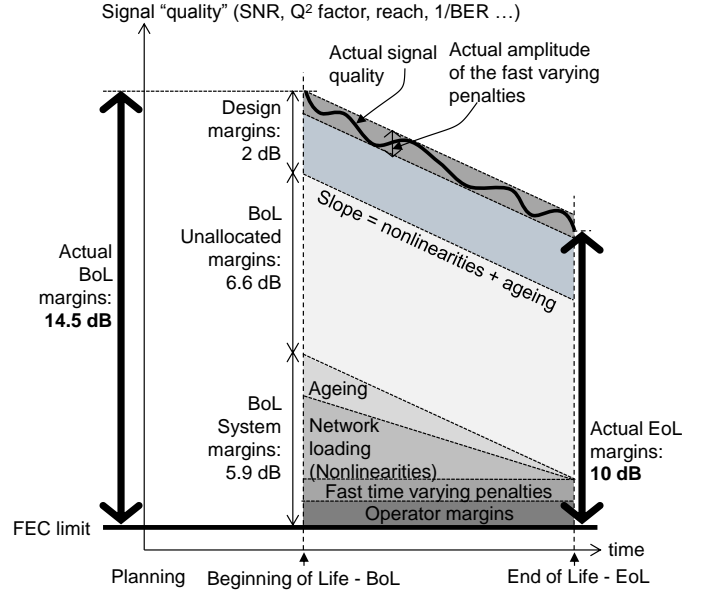


Fig. 1. Types of margins and their evolution. Sample margin values for the example given in Fig. 2 and Table I.

loading during the network life, which translates into additional nonlinearities, or to network equipment ageing: increasing fiber losses due to splices to repair fiber cuts, degrading amplifier noise factor, and detuning of the lasers leading to misalignment with optical filters in the intermediate nodes such as ROADMs (Reconfigurable Add/Drop Multiplexers). S-margins may include an additional operator margin [2]. Operator margins are required by the operator; as such, no effort is made here to minimize or leverage them. S-margins define the minimum quality value of the signal to be met at network Beginning of Life (BoL). Adding U- and S-margins to the FEC limit defines the planned BoL value of the signal quality metric.

Design margins (*D-margins*) are the difference between the planned BoL value and the real value of the quality metric, and are due to the inaccuracies of the design tool used to evaluate the QoT of all signals during network planning, which stem from two main sources: the inaccuracy of the QoT model itself, and the inaccuracy of the QoT tool inputs.

In this paper, after giving a short example on what values each margin may reach for a sample light path in Section II, we review in turn how each of the unallocated (Section III), system and unallocated combined (Section IV), and design (Section V) margins may be reduced or leveraged by the network designer or operator. In Section VI, we recap the technological challenges in designing low design margins; we summarize the techniques presented in this paper in

Y. Pointurier is with Nokia Bell Labs, Nozay, France; e-mail: mailto:yvan@ieee.org.

This paper is an extended version of [1].

Manuscript submitted in June 2016.




Fig. 2. Sample light path with 3 route-and-select nodes including the “add” node (TX) and the “drop” node (RX), linked by three 100 km fiber spans.

TABLE I
MARGIN TYPES AND TYPICAL VALUES

Margin type	SNR Margin
Unallocated margins (U)	Several dB
Design margin (D)	<2 dB [2]
Nonlinearities (S)	1.5–3dB [3]
Amplifier NF ageing (S)	0.7 dB [4]
Fiber ageing (cuts) (S)	1.6e−3 dB/km/year (linear SNR) [5]
Nodes ageing (filtering) (S)	0.05 dB / filter [6], [7]
Transponder ageing (S)	0.5 dB [7]
Fast variations (S)	0.4 dB [8]

Section VII.

II. SAMPLE VALUES FOR THE MARGINS

Margin evolution with time is illustrated in Fig. 1 and typical values for each margin can be found in Table I.

Consider for example the 600 km long light path depicted in Fig. 2, which carries a 100 Gb/s PDM-QPSK (Polarization Division Multiplexed - Quadrature Phase Shift Keying) signal with standard 12% soft decision FEC. Assume the light path is active for 10 years in a network with route-and-select nodes (i.e., 2 wavelength selective switches and hence 2 filters per intermediate node), 100 km fiber spans of standard single mode fiber with no in-line dispersion compensation, 1 node every 300 km, and amplifiers with noise factor (NF) of 4.5 dB that compensate exactly the fiber losses of 0.22 dB/km. We compute typical margins (in SNR) here.

Assuming, as in Table I, 3 dB (SNR) margins for the nonlinearities, amplifier NF ageing of 0.7 dB, fiber span ageing of 1.1 dB¹, filter ageing of 0.05 dB per filter i.e. 0.2 dB for 4 filters (1 filter at the add side, 2 filters at the intermediate node, and 1 filter at the drop side), TRX ageing of 0.5 dB, and 0.4 dB for the fast varying effects, the S-margin is 5.9 dB at BoL².

Further assuming a completely unloaded system at BoL and using the model in [9] and accounting for penalties of 1 dB for TRX and 0.03 dB per filter, the reach of such a system is 7500 km, resulting in a combined U- and S-margin of $10 \log_{10}(7500/600) = 12.5$ dB, i.e. a U-margin of 6.6 dB.

¹This corresponds to 4.3 fiber cuts per 1000 miles per year, 0.3 dB additional loss per splice needed to repair the fiber, resulting in a linear SNR penalty of 1.6e-3 dB/km/year due to increased ASE noise; for a line consisting of homogeneous 100 km spans and homogeneously spread cuts, the linear SNR penalty over 10 years is 1.6 dB. For the optical power that minimizes BER, penalties due to linear effects are twice the penalties due to nonlinear effect, such that the SNR penalty corresponding to a linear SNR penalty α is $2\alpha/3$.

²We neglect any additional operator margins, which are not discussed here.




Fig. 3. Accommodating traffic growth by leveraging U-margins. (a) Present mode of operation; (b) Leveraging U-margins: when more capacity is needed, the light path is upgraded to 200 Gb/s.

Including a D-margin of 2 dB, the BoL SNR margin is 14.5 dB.

Assuming that components have aged as planned (typically a worst case) and that the network is fully loaded at EoL, the S-margins ideally amount to 0.4 dB (i.e., fast varying effects) at network End of Life (EoL) while the D- and U-margins have not changed, yielding a total EoL margin of 9 dB.

In the following, we will examine how all of those margins can be reduced and leveraged.

III. UNALLOCATED MARGINS REDUCTION

As explained in [2], U-margins are known before network deployment, since U-margins come from the mismatch between demand (capacity or reach) and equipment capability (capacity or reach), which are known at design time. Hence, U-margins can be leveraged both in green field or brown field (network upgrades) scenarios. U-margins could be completely removed if, for each demand (the combination of a datarate and of a reach), a transponder that could transmit exactly the requested datarate (and not more) over the requested reach (and not more) was deployed. Such transponders cannot exist, since the relation between datarate and reach is dictated by physics and information theory (and quality of implementation). For this reason, U-margins are unavoidable.

However, reduction of the U-margins may be achieved through the utilization of a rate-flexible transponder (flex-TRX), which adjusts its datarate to the targeted reach. Coarse granularity flex-TRX (e.g., 100/200/400 Gb/s) will use only part of the U-margin, while fine granularity flex-TRX relying for instance on time hybrid QAM (Quadrature Amplitude Modulation) or 4-D modulation formats [10]–[12] further reduce the U-margins.

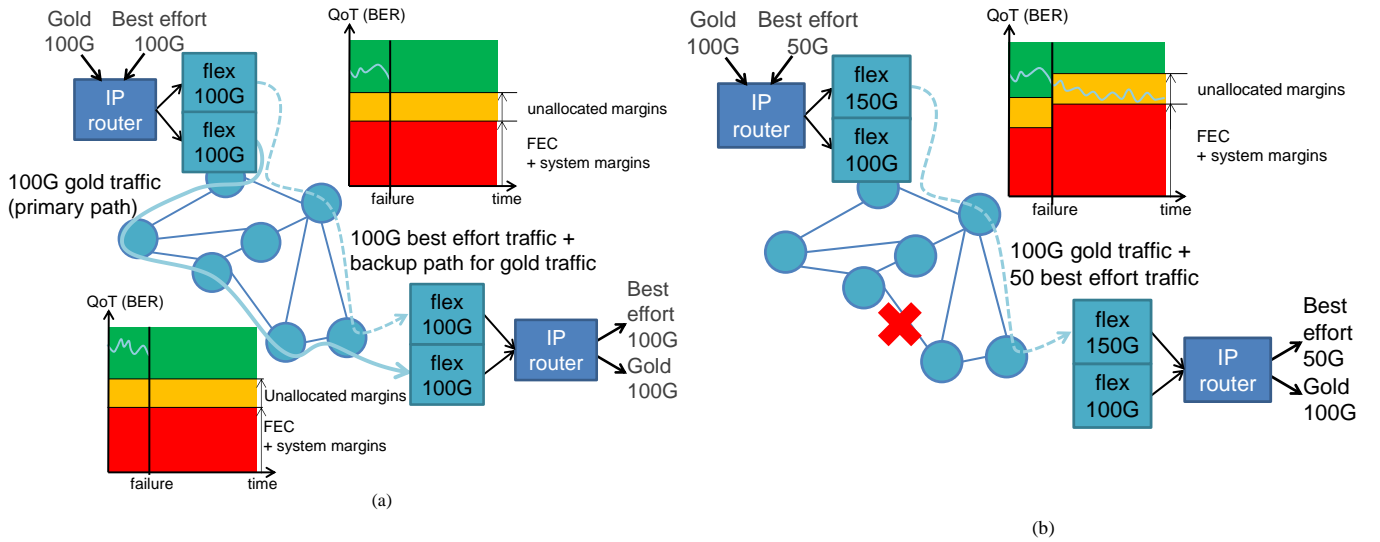


Fig. 4. Optical restoration scenario: (a) before failure; (b) after a link failure (red cross).

A. Network deployment and upgrade

Some U-margins remain even with arbitrarily fine granularity rate-flexible transponders, since, by trading-off additional capacity for reduced reach (while still meeting the reach constraint), extra capacity is available in the network, which is one form of U-margins. Such extra capacity, although not needed to meet the demand for which the network is designed, can be leveraged to decrease the total cost of ownership of the network over multiple years or periods [13].

As an illustration, consider Fig. 3(a), where a 100 Gb/s, 600 km demand is allocated a PDM-QPSK 100 Gb/s light path, leaving a large capacity margin (200 Gb/s could be transmitted over 600 km using PDM-16QAM) or reach margin (the reach of a standard 100 Gb/s signal is several thousands of km, as seen in Section II). Note that, in addition to the 100 Gb/s optical TRX, a 100 Gb/s client port (e.g., an IP router or OTN switch port) is also typically needed, since optical light paths are fed with data coming from an electronic network. Assuming that the deployed optical TRX is rate-flexible and 200 Gb/s-capable: when the demand grows, for instance to 200 Gb/s, the U-margin is leveraged and the extra capacity can be allocated without replacing the optical TRX, by simply changing its modulation format from PDM-QPSK to PDM-16QAM (Fig. 3(b)). This requires, nevertheless, the deployment of a second client (IP or OTN) 100 Gb/s port, which is needed anyway when the demand grows. Such multi-layer (electronic and optical) network operation requires specific support from the control plane, and online network capacity re-allocation.

B. Optical restoration

As a second example for U-margins reduction, consider the restoration scenario depicted in Fig. 4(a). A 100 Gb/s gold traffic demand (requiring protection in case of the failure of a network element) is allocated a first light path (plain line); a second light path (dashed line) is also allocated to protect this gold traffic. In this restoration scenario, the gold traffic backup path is shared with a 100 Gb/s best effort traffic demand. The

gold demand is 600 km long and the best effort demand is 800 km long and rate-flexible optical TRX are used, such that large reach or capacity U-margins are available.

The QoT evolution of each signal is also depicted in Fig. 4(a). The red zone accounts for the S-margins and thus depicts operation forbidden in typical network operation. The orange zone further includes the U-margins; operation in the orange zone is allowed, but not planned. The green zone depicts planned QoT operation.

During regular (no failure) operation, the QoT of each light path is within the green zone. Suppose a link failure occurs, as depicted in Fig. 4(b). If U-margins were not leveraged, the 100 Gb/s best effort traffic would be completely lost to enable the protection path QoT to remain in the green zone.

Fig. 4(b) further depicts how the U-margins can be leveraged. Instead of losing 100 Gb/s of best effort traffic, the protection TRX adapts its rate from 100 Gb/s to 150 Gb/s (e.g., from PDM-QPSK to PDM-8QAM) and the extra 50 Gb/s of capacity can be used to carry part of the best effort traffic³. Since the TRX increases its data rate, the QoT needed for the carried signal to remain above the FEC threshold increases, such that the red and orange zones are shifted up. The light path now operates in the orange zone, closer to the FEC limit, but still above the S-margins. The U-margins were leveraged to allow the increase of the signal data rate and the minimization of lost best effort traffic. This scenario is studied more in depth in [14], where it is shown, for a national network, that the total number of optical TRX can be reduced by 35%.

The U-margins reductions outlined in this section are generalized in the next section, and can be further combined with S-margins reduction to obtain significant network extra capacity or cost gains.

IV. UNALLOCATED AND SYSTEM MARGIN REDUCTION

Unlike the U-margins, S-margins are known after network deployment, and may vary with time. Fast time varying effects

³200 Gb/s PDM-16QAM is not possible over 800 km hence some best effort traffic has to be lost.

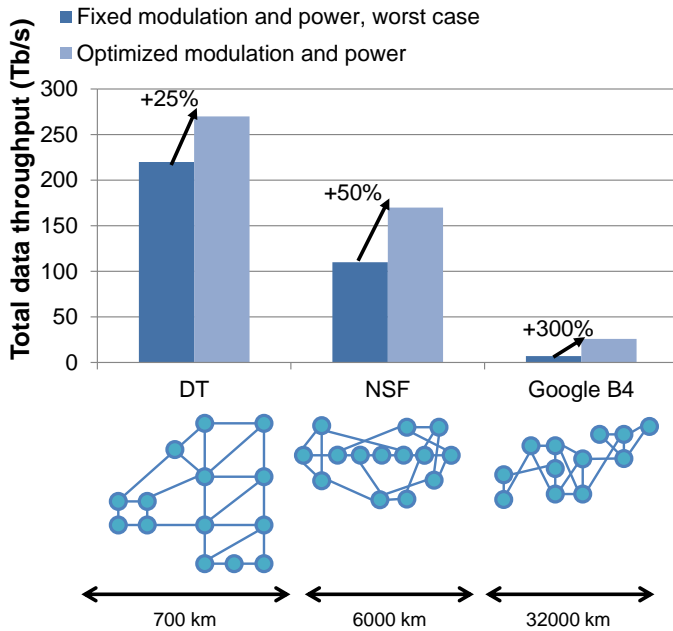


Fig. 5. Reduction of system and unallocated margins with flexible modulation format and power (from [3]).

that are not directly mitigated through TRX digital signal processing may be translated into capacity only at the expense of reduced network resilience, or to time-varying transported capacity which may temporarily be below the demand.

Slow varying effects such as nonlinear effects, which increase as new channels are lit, and component ageing, are more predictable, and may be leveraged when upgrading the network. To fully leverage S-margins stemming from network loading (i.e., from nonlinear impairments), careful power allocation is required; in fact, each light path may have its own modulation and power, leading to a RMSPA (routing, modulation, spectrum, power allocation) network design problem. At the light path level, [15] shows that, when accounting for the exact link load, the reach for a standard PDM-QPSK signal is almost double at BoL compared with EoL; the highest reach gains are achieved in lightly loaded networks when nonlinearities are smaller, i.e., in the early stages of the life of the network.

The following two techniques rely on multi-rate transponders and show how to leverage both S- and U-margins simultaneously; the first technique minimizes nonlinear effects to allocate more capacity in a network. The second technique leverages margins to delay equipment deployment and benefit from equipment cost erosion to minimize the network cost over a long period of time. Note that, in either case, perfect knowledge of the network is assumed to conduct the study. Such knowledge is generally not available; practically, measurements of the actual state of the network through dedicated monitoring would be needed in the studies outlined below.

A. Network capacity maximization

The impact of optimizing light path modulation format, power and spectrum to accommodate more capacity in a network was studied in depth in [3]. Results are shown in

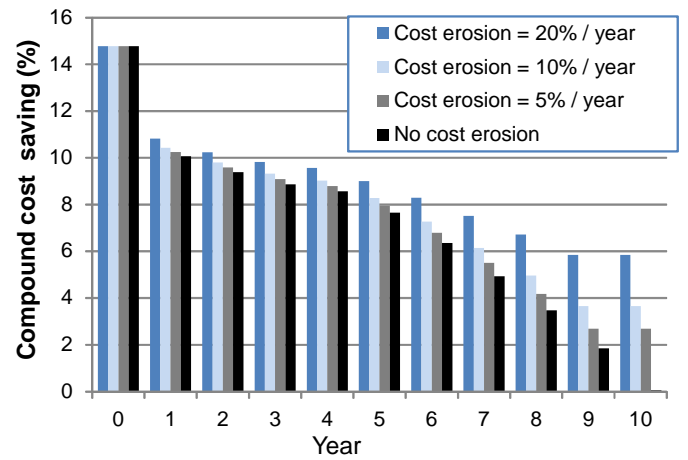


Fig. 6. Cost savings for the delayed deployment of rate-flexible transponders in a European national network (from [7]).

Fig. 5 for 3 networks with very different sizes: a national network (diameter of a few hundreds of km), a continental network (diameter of a few thousands of km) and a global network (diameter of a few tens of thousands of km). The capacity for the optimized dimensioning is compared with the case of worst-case dimensioning (fixed modulation format and power, all links fully loaded). The additional capacity that can be carried with the optimized design for the 3 aforementioned network topologies is found to be 30%, 50%, and 300%, respectively. Note that the larger the network, the larger the gain, as it is the network diameter that drives the modulation format in the fixed-rate, reference case, and the larger the diversity in route lengths, the higher the benefit of multi-rate transponders. Authors in [16] find similar results for continental networks.

As mentioned above, S- and U-margins are not known until the network is deployed. Hence, the gains demonstrated here may only be fully leveraged to provision extra capacity after the network starts operating and the margins are actually known. However, since gains may be leveraged early in the network life, the network capacity maximization technique described here can be considered as a green field optimization.

B. Multi-period network cost optimization

With the second technique [7], the equipment is used at the maximum capacity at network BoL, to benefit from the high S-margins. The extra capacity enabled by rate adaptation is used to carry additional traffic, such that U-margins are also used. As the network elements age and the network load (nonlinearities) increases, the system margin decreases and the capacity that can actually be carried in the network also decreases. New equipment is then provisioned to carry the capacity that is lost due to equipment ageing and increasing nonlinearities. Since this new equipment is deployed well after the initial network commissioning, it has benefited from cost erosion (i.e., it is less expensive than if it had been bought at network commissioning time), and the total cost of the network over its lifetime also decreases compared with the case where all network equipment had been deployed at

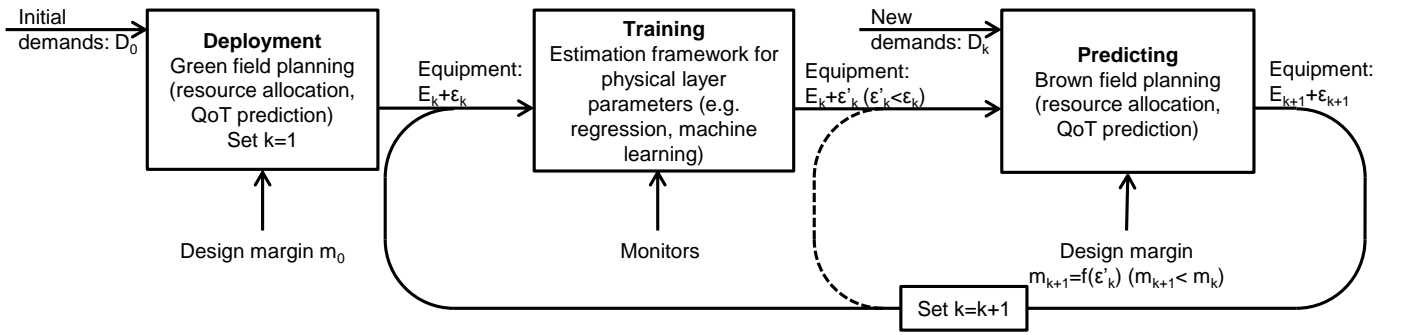


Fig. 7. Reduction of D-margins with monitoring and QoT tool inputs estimation.

BoL. This technique is used for brown field scenarios, where networks are already deployed.

Note that, although promising, the gains are offset by the increase of the network demand, which requires new equipment deployment with time. Planning then becomes a multi-period planning problem, where equipment is deployed in several stages, for instance, every 1 or 2 years, and each new piece of equipment is used to compensate at the same time both ageing and increased demands. The impact of such multi-year network planning is further studied in [7], where slow variations only account for component ageing (typically 2 dB in a national network) and not network loading (increasing nonlinearities), which is taken as a worst case (full load even at BoL). The results are depicted in Fig. 6, for several values of the cost erosion (from 0 to 20% per year), and for a demand compound average growth rate of 20% per year. The study considers only the cost of the rate-flexible transponders (100/150/200 Gb/s) and was carried for an Italian-sized national network, the network life duration was assumed to be 10 years, and compound cost savings (total savings from BoL up to 1, 2 ... 10 years) are reported. Other assumptions can be found in [7]. The largest savings of around 15% are made near the network beginning of life; as the network ages and traffic increases, the savings decrease until reaching 6% at network EoL assuming the largest investigated cost erosion of 20% per year. Including finer-granularity TRX, flexible grid, multi-layer dimensioning and the effect of increasing nonlinearities during the network life is the topic of ongoing work, and will likely improve the savings.

V. DESIGN MARGINS REDUCTION

D-margins come from the uncertainty of both the QoT estimation tool, and of the inputs of the tool: network equipment characteristics (link attenuation, chromatic dispersion map, amplifier noise factor, filter alignment...). Those two effects are fundamentally difficult to separate and their compound impact is only known at deployment time.

Reducing the margin due to the inaccuracy of the QoT tool itself is the topic of much research, which involves the design of accurate physical models suitable for realistic network deployments [9], [17]–[20].

It is additionally possible to reduce the QoT tool inputs uncertainty through monitoring and estimation techniques, in order to make more accurate QoT prediction for new

light paths during network upgrades, or in dynamic networks subject to new light path arrivals. Consider the network design cycle depicted in Fig. 7. The left-hand side of the figure depicts the usual green field planning using initial demand set D_0 , which results in a network with imperfect knowledge of the deployed equipment $E_0 + \epsilon_0$ (where E_0 quantifies the true physical characteristics of the deployed equipment — e.g. link lengths, dispersion map, amplifiers noise factors, transponder rates... and ϵ_0 quantifies the uncertainty on E_0). Uncertainties are accounted for in the resource allocation algorithm via a pre-defined D-margin m_0 ; these margins can be pre-defined or dependent on each light path, e.g. longer light paths may be associated with higher margin as in [21].

D-margins can then be measured and are thus known; to avoid a new light path to be subject to the original D-margins (upper-bounded by m_0), it is possible to benefit from the wealth of path-level monitoring information made available by coherent receivers almost for free, including received power, residual chromatic dispersion, noise level, and polarization [22].

On the k^{th} upgrade, i.e. at the arrival of demand (set) D_k , the following two-step process is used. During the first, “training” step, monitoring information is used to feed an “estimation framework for physical layer parameters”, which goal is to refine the knowledge of the underlying physical layer and thus decrease uncertainty ϵ_k to $\epsilon'_k < \epsilon_k$.

Then, during the second, “prediction” step, the network planner will then be able to use lower D-margins $m_{k+1} = f(\epsilon'_k) < m_k$ when establishing a new light path. Lowering the D-margin will in turn decrease the amount of equipment $E_{k+1} + \epsilon_{k+1}$ to be deployed when accommodating new demands. The process then repeats for each new arrival. Optionally, the training process may be bypassed for some arrivals — as depicted with the dashed line in Fig. 7 — and run only periodically, for instance if arrivals are fast with respect to the training time.

The physical layer parameters estimator is a generic estimator. As an example, observe that QoT estimators typically require link-level characteristics while coherent receivers yield path-level measurements; link-level metrics may be estimated via correlation techniques such as network kriging [23] when the characteristics are linearly additive (i.e. addition of link-level metrics yields path-level metric), or more advanced techniques such as machine learning, which are better adapted




Fig. 8. Reduction of over provisioning through minimization of the design margins thanks to monitoring and QoT tool inputs estimation (from [25]).

to nonlinear network characteristics [24].

Although we focused here on the refinement of uncertainty reduction, note that it is also possible to fully replace the QoT estimation tool with machine learning/“big data” techniques, by learning the transmission model itself through measurements and monitoring. Although such a method may seem attractive, it suffers from the lack of insight brought by physical modeling. For this reason, physical modeling, advanced estimation frameworks such as machine learning, and monitoring are complementary and should be used together to build accurate QoT estimation tools.

The impact of the reduction of D-margins thanks to monitoring was assessed in [25] for a national network topology subject to dynamic arrivals (and tear-downs) of network demands. The total number of regenerators (extra optical transponders used to regenerate the optical signal on light paths that are longer than the signal reach) needed in the network was assessed for the following 3 cases: assuming fully loaded links (this overestimates the nonlinear impairments), assuming perfect knowledge of the QoT inputs (impractical ideal case), and by estimating the inputs of the QoT estimator using monitoring and the network kriging framework mentioned above upon the arrival of each new demand. It is shown in Fig. 8 that the number of regenerators needed with monitoring and estimation of the missing QoT parameters approximates well the number of regenerators that would be needed if the physical layer was perfectly known; compared with the worst-case design used today, 20 to 50% fewer regenerators are needed depending on the network load. A similar concept was implemented in a 6-node testbed, where it experimentally shown that OSNR monitoring can reduce the D-margins to 1.5 dB [26].

VI. CHALLENGES

Although designing low-margin networks can result in a clear network capacity increase, translation into CAPEX gains

will prove challenging for the following reasons. First, margins are highly fragmented, despite the total margin reaching an appealing 10 dB or even more, as shown in the example in Section II. The U-margins (several dB) are easier to leverage, essentially requiring flexible transponders. Within the S-margins, the fast variable component (a fraction of a dB) will be difficult to leverage, as either a fast variable transponder, or a fast reconfigurable network infrastructure, would be needed. Ageing excluding nonlinearities may reach several dB, which can only be exploited with proper monitoring. Network loading (nonlinearities) may amount to another 3 dB, but exploiting them requires fine, difficult per-wavelength power tuning. Mitigating D-margins, which account for up to 2 dB, requires advanced monitoring and control plane support for information correlation [27].

Hence, in order to achieve the gains allowed by margins reduction, the following elements are needed:

- rate-flexible transponders and nodes (essentially, filters in ROADMs),
- online network re-optimization,
- optical performance monitors,
- per-light path power management,
- better physical models,
- an adapted control plane.

Most technology is available today and is facilitated by centralized network management in Software Defined Networks (SDN); note that frequent transponder rate adjustment may be needed during network re-optimization after each period, possibly leading to data loss during reconfiguration. Such losses can be almost completely prevented using for instance the transponder described in [28]. Much optical monitoring capability is already deployed within the coherent TRX; whether additional, dedicated optical power monitors would further help decreasing margins is the topic of ongoing research.

Deploying flexible interfaces and varying their capacity with time, means that the interfaces client and line sides should be independent. Indeed, demands (on the client side) should be met even when the interface datarate changes to adjust to a varying margin. Multilayer nodes that are able to dynamically map electronic resources to optical resources are thus needed. Electronic switching, in addition to the optical transmission equipment, should therefore be provisioned appropriately. This calls for multilayer (electronic and optical), multi-year (accounting for foreseeable ageing such as nonlinearities) routing, spectrum, modulation format, power allocation algorithms relying on monitoring information to constantly adjust network capacity and capacity prediction to the true network state.

Finally, per-light path power management is considered as tricky by the operators, since light path power not only depends on the launch power, but also on the amplifiers’ response and on equalization at the intermediate nodes. In addition, changing the power of a light path also impacts on the QoT of already established light paths. Whether per-light path power management will be adapted or not will depend on the trade-off between network ease of operation (OPEX) and the gain in capacity (CAPEX).

TABLE II
SUMMARY OF TECHNIQUES ALLOWING A LOW MARGINS NETWORK DESIGN

Margin	Cause	Mitigation approach	Enabling technology	Expected improvement
Unallocated (Section III-A)	TRX capacity or reach granularity vs. demand mismatch	Routing, spectrum, rate allocation	Rate-flexible TRX and ROADM, online reconfiguration, adapted control plane	Higher capacity / better spectrum utilization: longer network life, delayed or decreased investment (green and brown field)
Unallocated (Section III-B [14])	TRX capacity or reach granularity vs. demand mismatch	Class-of-service-aware restoration	Rate-flexible TRX and ROADM, online reconfiguration, adapted control plane	Higher capacity / better spectrum utilization: longer network life, delayed or decreased investment (green and brown field)
System, Unallocated (Section IV-A [3])	Worst-case light path power allocation	Routing, modulation, spectrum, power, rate allocation	Rate-flexible TRX and ROADM, control plane, per-light path power allocation, optical monitoring	Higher capacity / better spectrum utilization: longer network life, delayed or decreased investment (green and brown field)
System, Unallocated (Section IV-B [7])	Ageing	Multi-period allocation	Rate-flexible TRX and ROADM, online reconfiguration, adapted control plane, optical monitoring	Higher capacity / better spectrum utilization: longer network life, delayed or decreased investment on network upgrades (brown field)
Design (Section V [9], [17]–[20])	QoT tool inaccuracy	Refined QoT estimation tool	Physical insight on impairments	Decrease QoT prediction inaccuracy: deploy less equipment at BoL and upgrades (green and brown field)
Design (Section V [25])	QoT inputs uncertainty	Power, spectrum, rate allocation	Optical monitoring, online reconfiguration, adapted control plane	Decrease QoT prediction inaccuracy: deploy less equipment on upgrades (brown field)

VII. SUMMARY AND CONCLUSIONS

Network margins, although plentiful, require a variety of technologies to be fully exploited and translated into additional capacity or decreased network cost (CAPEX and OPEX), typically in brown field scenarios, sometimes in green field scenarios. The variety of the techniques presented here to leverage margins is summarized in Table II.

Overall, the margins available in an optical network are the sum of 3 main components: unallocated, system, and design margins, each of which can actually be the sum of even smaller components. Tackling each component separately may not yield significant benefits; in fact it is necessary to tackle all components to fully leverage the optical margins and achieve substantial gains in capacity, network life duration, CAPEX, or total cost of ownership.

We also reviewed the key components needed to leverage margins. Although many of those blocks already exist — at least at the research level,— using them in real networks is still considered as a tremendous operational but promising challenge today.

ACKNOWLEDGMENT

The author would like to thank his Nokia and Nokia Bell Labs colleagues: T. Zami, J. Pesic, P. Ramantanis, P. Jennevé, N. Rossi, C. Delezoide and S. Bigo for their helpful insights and comments.

This work was partly funded by the European Commission through the H2020 ORCHESTRA project.

REFERENCES

- [1] Y. Pointurier, “Design of low-margin optical networks,” in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2016, paper Tu3F.5 (invited).
- [2] J.-L. Augé, “Can we use flexible transponders to reduce margins?” in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2013, paper OTu2A.1.
- [3] D. J. Ives, P. Bayvel, and S. Savory, “Routing, modulation, spectrum and launch power assignment to maximize the traffic throughput of a nonlinear optical mesh network,” *Springer Photonic Network Communications*, vol. 29, no. 3, pp. 244–256, Jun. 2015.
- [4] L. N. Binh, *Optical Fiber Communications Systems*. CRC Press, 2011.
- [5] M. To and P. Neusy, “Unavailability analysis of long-haul networks,” *IEEE J. Sel. Areas Commun.*, vol. 12, no. 1, pp. 100–109, Jan. 1994.
- [6] T. Rahman, A. Napoli, D. Rafique, B. Spinnler, M. Kuschnerov, I. Lobato, B. Clouet, M. Bohn, C. Okonkwo, and H. de Waardt, “On the mitigation of optical filtering penalties originating from ROADM cascade,” *IEEE Photon. Technol. Lett.*, vol. 26, no. 2, pp. 154–157, Jan. 2014.
- [7] J. Pesic, T. Zami, P. Ramantanis, and S. Bigo, “Faster return of investment in WDM networks when elastic transponders dynamically fit ageing of link margins,” in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2016, paper M3K.2.
- [8] L. E. Nelson, G. Zhang, M. Birk, C. Skolnick, R. Isaac, Y. Pan, C. Rasmussen, G. Pendock, and B. Mikkelsen, “A robust real-time 100G transceiver with soft-decision forward error correction,” *IEEE J. Opt. Commun. Netw.*, vol. 4, no. 11, pp. B131–B141, Nov. 2012.
- [9] A. Bononi, N. Rossi, and P. Serena, “On the nonlinear threshold versus distance in long-haul highly-dispersive coherent systems,” *Optics Express*, vol. 20, no. 26, pp. B204–B216, Dec. 2012.
- [10] X. Zhou, L. E. Nelson, and P. Magill, “Rate-adaptable optics for next generation long-haul transport networks,” *IEEE Commun. Mag.*, vol. 51, no. 3, pp. 41–49, Mar. 2013.
- [11] J. Renaudier, O. Bertran-Pardo, A. Ghazisaeidi, P. Tran, H. Mardoyan, P. Brindel, A. Voicila, G. Charlet, and S. Bigo, “Experimental transmission of Nyquist pulse shaped 4-D coded modulation using dual polarization 16QAM set-partitioning schemes at 28 Gbaud,” in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2013, paper OTu3B.1.
- [12] M. A. Mestre, J. M. Estaran, P. Jennevé, H. Mardoyan, I. Tafur Monroy, D. Zibar, and S. Bigo, “On adaptive transponders for optical slot switched networks: Nyquist-QAM and CO-OFDM experimental comparison,” *J. Lightw. Technol.*, vol. 34, no. 8, pp. 1851–1858, Apr. 2016.
- [13] Ciena, “Transforming margin into capacity with liquid spectrum,” 2015, white paper.
- [14] A. Morea, G. Charlet, and D. Verchère, “Elasticity for dynamic recovery in OTN networks,” in *Proc. Asia Communications and Photonics Conference (ACP)*, Nov. 2014, paper AW4E.2.
- [15] A. Bononi, P. Serena, A. Morea, and G. Picchi, “Regeneration savings in flexible optical networks with a new load-aware reach maximization,” *Elsevier Optical Switching and Networking*, vol. 19, no. 3, pp. 123–134, Jan. 2016.
- [16] A. Lord, P. Wright, and A. Mitra, “Core networks in the flexgrid era,” *J. Lightw. Technol.*, vol. 33, no. 5, pp. 1126–1135, Mar. 2015.

- [17] E. Seve, P. Ramantanis, J.-C. Antona, E. Grellier, O. Rival, F. Vacondio, and S. Bigo, "Semi-analytical model for the performance estimation of 100Gb/s PDM-QPSK optical transmission systems without inline dispersion compensation and mixed fiber types," in *Proc. European Conference and Exhibition on Optical Communication (ECOC)*, Sep. 2013, paper Th.1.D.2.
- [18] A. Carena, G. Bosco, V. Curri, Y. Jiang, P. Poggiolini, and F. Forghieri, "EGN model of non-linear fiber propagation," *Optics Express*, vol. 22, no. 13, pp. 16 335–16 362, Mar. 2014.
- [19] P. Serena and A. Bononi, "A time-domain extended gaussian noise model," *J. Lightw. Technol.*, vol. 33, no. 7, pp. 1459–1472, Apr. 2015.
- [20] R. Dar, M. Feder, A. Mecozzi, and M. Shtaf, "Inter-channel nonlinear interference noise in WDM systems: modeling and mitigation," *J. Lightw. Technol.*, vol. 33, no. 5, pp. 1044–1053, Mar. 2015.
- [21] T. Zami, A. Morea, and F. Leplingard, "A new method to plan more realistic optical transparent networks," *Bell Labs Technical Journal*, vol. 14, no. 4, pp. 213–226, Winter 2010.
- [22] K. Christodoulopoulos, P. Kokkinos, A. Di Giglio, A. Pagano, N. Argyris, C. Spatharakis, S. Dris, H. Avramopoulos, J.-C. Antona, C. Delezoide, P. Jennevé, J. Pesic, Y. Pointurier, N. Sambo, F. Cugini, P. Castoldi, G. Bernini, G. Carrozzo, and E. Varvarigos, "ORCHESTRA - optical performance monitoring enabling flexible networking," in *Proc. International Conference on Transparent Optical Networks (ICTON)*, Jul. 2015, paper We.C1.2.
- [23] Y. Pointurier, M. Coates, and M. Rabbat, "Cross-layer monitoring in transparent optical networks," *IEEE J. Opt. Commun. Netw.*, vol. 3, no. 3, pp. 189–198, Mar. 2011.
- [24] D. Zibar, M. Piels, R. Jones, and C. G. Schäffer, "Machine learning techniques in optical communication," *J. Lightw. Technol.*, vol. 34, no. 6, pp. 1442–1452, Mar. 2016.
- [25] I. Sartzetakis, K. Christodoulopoulos, C. P. Tsekrekos, D. Syvridis, and E. Varvarigos, "Estimating QoT of unestablished lightpaths," in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2016, paper Tu3F.2.
- [26] S. Oda, M. Miyabe, S. Yoshida, T. Katagiri, Y. Aoki, J. C. Rasmussen, M. Birk, and K. Tse, "Demonstration of an autonomous, software controlled living optical network that eliminates the need for pre-planning," in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2016, paper W2A.44.
- [27] N. Sambo, Y. Pointurier, F. Cugini, L. Valcarenghi, P. Castoldi, and I. Tomkos, "Lightpath establishment assisted by off-line QoT estimation in transparent optical networks," *IEEE J. Opt. Commun. Netw.*, vol. 2, no. 11, pp. 928–937, Nov. 2010.
- [28] A. Dupas, P. Layec, E. Dutisseuil, S. Bigo, S. Belotti, S. Misto, S. Annoni, Y. Yan, E. Hugues-Salas, G. Zervas, and D. Simeonidou, "Hitless 100 Gbit/s OTN Bandwidth Variable Transmitter for Software-Defined Networks," in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2016, paper Th3I.1.



Yvan Pointurier (S02M06SM12) received a Ph.D. from the University of Virginia, USA in 2006. Between 2006 and 2009, he was a postdoctoral fellow at McGill University in Montréal and then a senior researcher at AIT, Greece. In 2009, Dr. Pointurier joined Alcatel-Lucent (now Nokia) Bell Labs as a research engineer, and is now the head of the "Dynamic Optical Networking and Switching" department at Nokia Bell Labs. His team is working on circuit and optical packet switched networks, with activities ranging from the physical layer to planning

algorithms.

Dr. Pointurier has authored or co-authored more than 15 European and US patents, and over 80 technical papers in leading journals, key conferences (OFC, ECOC, ACM Internet Measurement Conference, IEEE Infocom, ICC, Globecom), and book chapters. He received a best paper award at the IEEE ICC conference in 2006 and an IEEE Communication Letters Exemplary Reviewer award in 2014 and 2015 (top 3% of the reviewers). Dr. Pointurier has been a TPC member for IEEE ICC, the flagship IEEE conference on communications, since 2007. He is currently the French coordinator for the CELTIC+ SENDATE-TANDEM project, a 36-month, 2200 person-month project led by Nokia Bell Labs.