Impairment-Aware Traffic Engineering Using Cross-Layer Protocols

F. Fidler(1,2), P. J. Winzer(2), M. K. Thottan(2), K. Bergman(1)

(1) Dept. of Electrical Engineering, Columbia University New York, NY, USA, *ffidler@ee.columbia.edu*
(2) Optical Networks Research, Bell Labs, Holmdel, NJ, USA, *peter.winzer@ieee.org*

Abstract *Traffic engineering based on cross-layer communications to counter-act inter-channel impairments in optical networks is shown to yield significant throughput gains for a variety of traffic models.*

Introduction

Today's wavelength division multiplexed (WDM) transport networks allocate sufficient margins for high physical-layer reliability, with very low outage probabilities (10^{-4} to 10^{-8}) and bit-error ratios (BERs) below 10^{-15} after physical-layer forward error correction (FEC). At the same time, optical mesh networking functionalities are constantly becoming more flexible and dynamic, leading to an increasingly wide range of WDM channels with different propagation histories, bit rates, and modulation formats sharing common fiber paths. The support of alien wavelengths and the possibility of rate-negotiating "optical modems" further enhance the variety of signals found on the physical layer. In order to guarantee adequate transmission performance despite this tremendous increase in networking flexibility, one can either design the physical layer for the worst case, which due to the variety of scenarios is likely to result in restrictive engineering rules (and hence in overly large margins), or one can think of relaxing transmission margins by introducing cross-layer protocols to enable dynamic, impairment-aware traffic management[1,2].

In this work, we quantitatively assess the throughput performance of optical networks in the presence of dynamic intra-channel impairments, both with and without the use of general cross-layer communication protocols. We find significant throughput gains for a variety of different traffic models.

Network topology and cross-layer functionality

We consider the example network shown in Fig. 1 and assume an optical network supporting a maximum of $N \in [100,250]$ WDM channels per link of mixed 10-Gb/s (N_{10}) and 40-Gb/s (N_{40}) per-channel bit rates, with $N = N_{10}+N_{40}$. The ratio $r = N_{40} / N_{10}$ is a variable parameter in our simulations; for fixed N, the overall network capacity increases with increasing r as more and more 40-Gb/s wavelengths are deployed. We assume reconfigurable optical add/drop multiplexers (ROADMs) at all nodes and make use of optical wavelength routing whenever possible. Opto-electronic regeneration (OEO) is required every 1500 km for both bit rates, and re-grooming among wavelengths may take place whenever OEO is performed. As for the physical-layer impairments, we assume a simple statistical model for the impact of cross-channel fiber nonlinearities, whereby we let a 10-Gb/s channel degrade the BER of a 40-Gb/s channel with probability p, provided that the two channels co-propagate on the same fiber. We then assume the 40-Gb/s channel to experience an outage, unless cross-layer protection strategies are used; in that case, we allow a traffic engineering entity integrated into the network's global control plane to identify the problematic 10-Gb/s channel by correlating the existing physical-layer connections running over a particular fiber with the pre-FEC BER read-outs from those transponders that route over that fiber. Traffic priorities permitting, the problematic 10-Gb/s channel is then switched off well before packets are lost on the 40-Gb/s link.

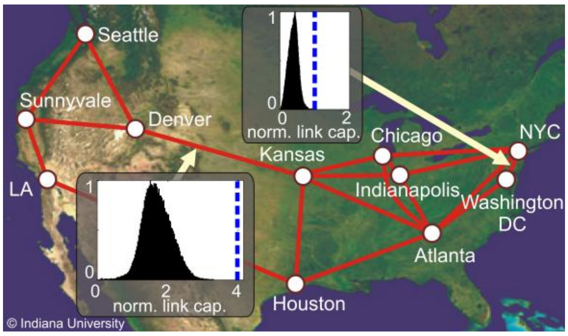

Fig. 1: Abilene network[3]. Insets: traffic histograms and provisioned capacities (dashed) for two representative links.

Capacity allocation and traffic matrices

Following Ref. [4], we use a linear programming (LP) approach to predict the required link capacities for a given class of traffic matrices ***D*** under shortest-path routing, and we provision the network capacities accordingly. The elements d_{ij} of the symmetric

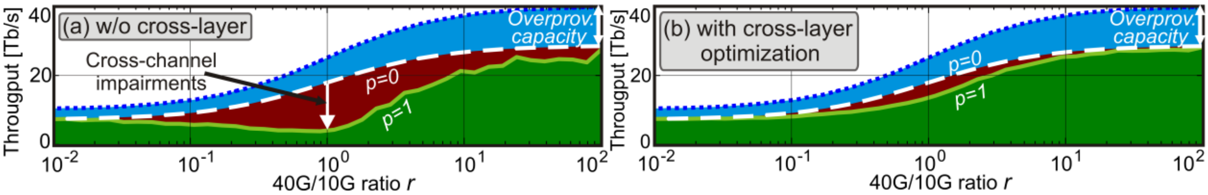

Fig. 2: Throughput vs. 40G/10G ratio r; (a) no cross-layer communication; (b) cross-layer communication. (Dotted: provisioned capacity; dashed: throughput w/o cross-channel impairments (p = 0), solid: throughput w/ cross-channel impairments (p = 1))

($d_{ij} = d_{ji}$) traffic matrices $\boldsymbol{D}$ describe the point-to-point demands between node i and node j. We then allocate wavelengths with the desired 40G/10G ratio r to each link, using N wavelengths for the highest-capacity link.

As for the traffic matrices, we investigate

- the class of *hose-constrained matrices*[5] ($0 \le d_{ij} \le 1$, $d_{ij} = 0$, $\sum_i d_{ij} = \sum_j d_{ij} = 1$; constant ingress/egress capacities, and randomly chosen individual demands with exponential statistics);
- the class of *gravity-based matrices*[6], where the mean value of d_{ij} is chosen proportional to the product of the (human) population at the node locations; a statistical spread (10% and 100%) is added to account for traffic variations.

We use extensive *Monte Carlo* simulations to generate ensembles of random traffic matrices within the respective class. The resulting demands are individually assigned to the deployed wavelengths in random order. Insets in Fig. 1 show, for two representative links, the provisioned capacities (dashed) and histograms of the occurring link capacities for 10,000 routed hose matrices. The gap between histogram and provisioned capacity reflects the required amount of over-provisioning for flexible routing of the entire class of hose matrices[4]. More restrictive classes of traffic matrices require less over-provisioning.

Network throughput and cross-layer performance

Figure 2 shows the over-provisioned capacity (dotted) as a function of the 40G/10G ratio r, and the mean network *throughput* (dashed) in the absence of cross-channel impairments ($p = 0$) for 10,000 random hose-constrained traffic matrices. Both curves increase monotonically as more 40G channels are deployed.

In the presence of cross-channel impairments (solid) and without cross-layer traffic engineering (Fig. 2a), the network throughput drops significantly. The worst case in throughput reduction occurs for an equal number of 10-Gb/s and 40-Gb/s channels ($r = 1$), unless specific traffic engineering rules are used (e.g., 'fill all 40-Gb/s channels first'). The green area below the solid line represents the amount of *guaranteed* throughput, i.e. the transport capacity that can be used to carry high priority ('*class A'*) traffic despite the presence of cross-channel impairments. The traffic in the red area may then only be carried as best-effort *'class B'* traffic.

Figure 2b shows the case where those 10-Gb/s channels that otherwise would cause a 40-Gb/s channel to experience an outage are dynamically switched off by the cross-layer entity, resulting in a significantly higher throughput for 'class A' traffic (green area). A small throughput loss remains (red area), as the interfering 10-Gb/s channels are removed.

Figure 3a shows the throughput gain from cross-layer aware networking (mean and standard deviation) as a function of the impairment probability p and $r = 1$. Little variation is found when varying the number of wavelengths N between 100 and 250. Figure 3b gives the throughput gain and its standard deviation (shaded areas) as a function of r and for $p = 1$, emphasizing the benefits of cross-layer communication protocols in networks with maximally mixed bit rates (blue curves, $r = 1$). The green curves apply to a wavelength assignment using the above mentioned traffic engineering rule, whereby 40-Gb/s channels are filled first. Cross-layer protocols can still achieve a throughput gain of ~100% in this case.

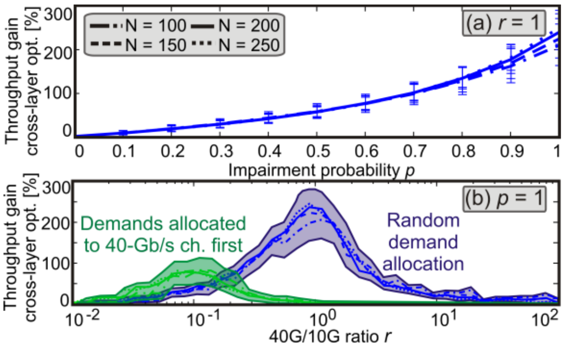

Fig. 3: Throughput gain vs. (a) impairment probability p and (b) channel ratio r (N: max. number of wavelength channels).

Finally, we investigate the impact of different traffic models. Table 1 compares key results for hose-constrained and gravity-based matrices (the latter with a maximum deviation of the individual demands of 10% and 100%). Even though most overall network parameters (such as the required amount of over-provisioning) change with the traffic model, the gain (mean and standard deviation) arising from cross-layer communication is largely *independent* of the assumed traffic model.

Tab. 1: Comparison: Hose vs gravity-based traffic matrices, showing mean throughput gain ± standard deviation.

Random channel filling	**Hose-constr.**	**Gravity-based**	
		dev. 100%	**dev. 10%**
Overprov. [%] w/o impairments	40	33	6
Throughput gain w/ cross-layer (r=1, p=1) [%]	227 ± 63	219 ± 69	241 ± 54

Conclusions

Using Monte Carlo simulations we showed a throughput gain of up to ~300% for cross-layer communication based protection mechanisms in optical networks suffering from cross-channel impairments. The throughput gains were found to be largely independent of the underlying traffic model.

We acknowledge valuable discussions with A.Lometti.

References

1 I.Tomkos et al., Proc. ECOC'08, We.3.D.1 (2008)
2 O.Gerstel, Proc. OFC'09, OML7 (2009)
3 [Online], www.internet2.edu/abilene, (2009)
4 F.B.Shepherd et al., JON **5**, 320 (2006)
5 N.G.Duffield et al., Trans.Networking **10**, 679 (2002)
6 I.Juva, Proc. ICC'07, 562 (2007)