Reduction of Delay Overfulfillment in IP-over-DWDM Transport Networks

Uwe Bauknecht, Tobias Enderle(✉), and Arthur Witt

Institute of Communication Networks and Computer Engineering, University of Stuttgart, Stuttgart, Germany
{uwe.bauknecht,tobias.enderle,arthur.witt}@ikr.uni-stuttgart.de

Abstract. The traffic in today's transport networks is increasing dramatically due to more demanding applications like video on demand and improving access technologies like 5G. Additionally, the quality of service requirements are becoming more stringent while network operators are seeking new ways for revenue creation. We propose a multi-layer network reconfiguration approach that reduces the overfulfillment of service delay requirements. In that way it provides an incentive to customers with low-delay services to acquire a more expensive service class for their traffic. Additionally, it relieves highly utilized links for services with very strict delay requirements. We provide an ILP formulation that solves this multi-layer network problem by performing a cross-layer optimization. Further, we evaluate our approach for two nationwide backbone networks. We show that a reduction of service delay overfulfillment is possible and how that affects other network metrics.

Keywords: Delay · ILP · Network reconfiguration · Optimization · Overfulfillment · QoS

1 Introduction

Since its inception, the Internet has changed significantly from a best-effort network of low-rate data exchange services into a ubiquitous and indispensable service delivery infrastructure for individuals and businesses alike. Trends like software as a service and video on demand, fueled by more potent access technologies such as FTTx and 5G, allow Internet-based business models such as cloud and content delivery services to thrive. These so-called over-the-top (OTT) services rely on the availability and quality of the underlying connection services rendered by Internet service providers (ISPs). Therefore, OTT providers require connection services that provide higher data rates and also adhere to guarantees on quality of service (QoS) metrics codified in service level agreement (SLAs). In order to ensure their fulfillment and avoid contract penalties, ISPs typically resort to overprovisioning QoS parameters on such SLA-based services. This practice translates to the ISPs' multi-layered core networks as well, where the packet forwarding layer and the optical transport layer are both subject to overdimensioning and further safety margins.

Published by Springer Nature Switzerland AG 2020
A. Tzanakaki et al. (Eds.): ONDM 2019, LNCS 11616, pp. 528–539, 2020.
https://doi.org/10.1007/978-3-030-38085-4_45

Matching the continuous growth in traffic volume of currently 26% annually[1] [5] becomes increasingly expensive, especially due to improvements of core network transmission equipment approaching the limits of the deployed fiber infrastructure [6]. To avoid or at least defer costly investments for commissioning new fibers, ISPs try to reduce operational costs and explore new avenues of revenue creation. Important means towards cost reduction are improving the utilization of existing equipment and reducing safety margins. Several works have been published which address this and related issues. For example, the authors of [8] introduce an approach which increases the network utilization by exploiting the fact that some service requests can temporarily tolerate lower bandwidth. This frees resources in high load situations and allows more services to be accepted in the long run. In [1] we presented a latency-aware network reconfiguration approach that minimizes the amount of active network hardware. In [11], ways to reduce safety margins that compensate physical layer impairments and aging are discussed. Techniques like these can be employed by a software-defined networking (SDN) approach which exploits global knowledge of the network state to dynamically adjust the network's configuration such that light path and traffic routing precisely fit the QoS requirements of the services to be realized and thereby reduce overprovisioning [13].

In addition to cost reduction, a more differentiated service level agreement (SLA) portfolio can help increase an ISP's revenue [14]. On the side of individual end customers with low-rate best-effort contracts, ISPs have already instated so-called "data caps", where the bandwidth is throttled by rate shaping if a customer exceeds a predetermined data volume. This way, they capitalize on heavy users by offering uncapped services at increased rates [12]. For business clients who use their connections beyond the limits of their SLA contracts up to the maximum overprovisioning, rate shaping is also an option to control the data rate. Other QoS parameters are harder to limit with inexpensive measures, leaving business customers with little to no incentive to acquire contracts with more stringent guarantees, given that they are well aware of their contracts' overfulfillment. Among these parameters especially delay requirements are imperative to the aforementioned services.

Sanctioning customers by actively limiting delay overfulfillment could be realized on both, the packet layer and the optical layer. On the packet layer, incoming traffic could simply be queued in buffers for a preconfigured amount of time. However, since the buffer size necessary for delaying scales directly with the data rate and typical line cards have limited buffer memory, additional hardware would be needed given the large data rates of some business customer connections. Moreover, this also raises concerns when considering availability requirements since additional active hardware also harbors additional points of failure. On the optical layer, latency can be artificially introduced by simple fiber delay lines. Due to the low cost and deterministic delay, this approach seems superior, but also requires custom delay lines for each service and a more complicated fiber routing to match services to delays.

[1] Compound Annual Growth Rate 2017 to 2022.

Rather than introducing additional hardware to artificially reduce quality, a differentiated service routing on the existing infrastructure enables reduced operational cost, provides means of sanctioning customers and allows for increased revenue services by exploiting the overprovisioning of existing hardware. Traffic of customers under sanction can be shifted to physically longer and therefore delay-inducing paths where switching equipment is otherwise underutilized while the shorter routes can be dedicated to premium services increasing revenue.

The main part of this work is organized as follows. In Sect. 2, we formulate the considered network architecture and reconfiguration problem. In Sect. 3, we describe our proposed path generation and optimization approach. Simulation setup and results are presented in Sect. 4. Eventually, we conclude the work in Sect. 5.

2 Problem Statement

We consider a two-layer transport network which receives aggregated client traffic demands from earlier stages of the ISP's aggregation network. We assume a controller of global knowledge on all state of the transport network. This could be realized by an SDN-controlled IP-over-DWDM architecture. Our prototypical network therefore consists of a circuit-oriented optical layer and a packet-oriented electrical layer on top. The physical topology is given as the directed graph $G = (V, E)$, where V is the set of nodes and E is the set of directed fiber links. Each node consists of an optical switch and a packet router connected to it. The optical switch acts as a contentionless, directionless and colorless reconfigurable optical add-drop multiplexer (ROADM), such that arbitrary optical circuits can be established between any fiber, which connects to another node, and any of the tunable ports in the line cards of the router. Each fiber link can provide a capacity in terms of data rate, defined by the maximum number of wavelengths multiplied by their individual data rates. We focus on the single line rate case.

Further, we are given a set of demands K that need to be routed in the network. Each demand $k \in K$ is an aggregation of services. It connects a node pair $(s,t) \in U = \{(u,v) \in V \times V : u \neq v\}$ with a fixed data rate h_k and a maximum allowed end-to-end delay $\tau_{k,\max}$, also referred to as target delay. A node pair can also be connected by multiple demands.

Demands are routed in the virtual topology which in turn is defined by optical circuits in the optical layer. An optical circuit requires one line port at the router of its source node and one at the router of its target node. Intermediate nodes are bypassed, i.e., the circuit is switched in the ROADM without conversion to the electrical domain. Depending on the path on which a demand k is routed through the network it experiences a certain delay τ_k. We denote the difference between $\tau_{k,\max}$, the maximum delay a demand is allowed to experience, and its actual delay τ_k as *delay overfulfillment*. Further we denote

$$\Delta\tau_k = \frac{\tau_{k,\max} - \tau_k}{\tau_{k,\max}} \tag{1}$$

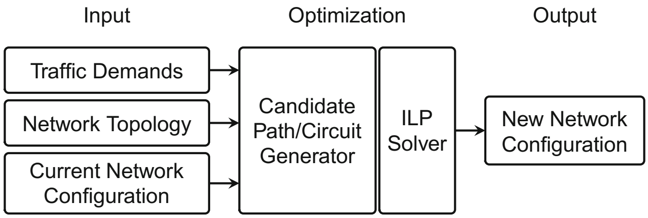

Fig. 1. Overview of the solution approach

as *relative delay overfulfillment*. We focus on core networks that span large geographical distances which in turn lead to significant propagation delays. In the core, these are the largest contributor to overall delay [15], followed by queuing delay. While queuing delay in routers can generally result in large delay spikes, these occur only rarely in the core where client traffic had already been subject to prior rate shaping. Together with traffic engineering sufficient link capacity to avoid excessive congestion can be ensured, such that we assume that queuing delays can be controlled in regular operation. Since existing studies show typical queuing delays to vary significantly between different networks from a few micro- [3] to a few milliseconds [2], which typically remains below the propagation delay in our scenario, we choose to focus on the deterministic part of the delay by modeling a demand's delay τ_k on a path by its propagation delay.

In [7], we already presented a delay-aware routing approach that minimizes the amount of active hardware in the network. This approach yields short paths for most demands which results in high delay overfulfillment. Our goal is to find a routing approach that reduces the delay overfulfillment compared to a resource optimization as shown in [7]. To be more precise, we find a network configuration, i.e., a set of circuits and a corresponding routing, for a given set of demands K that minimizes the average relative delay overfulfillment for a certain subset $K_{\mathrm{S}} \subseteq K$ of those demands. We study the effects on the blocking behavior for a continuous reconfiguration of the network with demands that change over time. Furthermore, we consider the effects on the amount of network hardware that is needed. We focus on router line cards because they are the lion's share of the operator's capital expenditure (CAPEX) [10].

3 Solution Approach

We solve the problem introduced in the previous section with the help of an integer linear program (ILP). The ILP is based on a link-path formulation, i.e., it selects optimal routes out of a precomputed set of candidate paths. The whole process is depicted in Fig. 1. As input the algorithm receives the network topology, the set of demands as well as the current configuration of the network. In the first step, candidate paths are generated for each of the demands. Subsequently, the demands are routed using the ILP. The result is a new network configuration. We employ a make before break migration approach meaning that if the

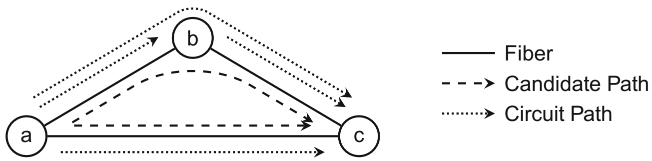

Fig. 2. Candidate path generation for a demand from node a to node c

routing of a demand changes from the previous to the current configuration, then the new path must be completely set up before the old one is torn down. Consequently, we need to ensure that enough capacity is available to realize old and new path simultaneously. The details are explained in the following sections.

3.1 Path Generation

The ILP selects optimal routes out of a set of precomputed candidate paths. The generation of these paths consists of two steps. In the first step, we find paths for the demands in the physical topology. In the second step, we create corresponding circuit realizations for each of the previously found paths. With this approach, a joint optimization of electrical and optical layer is possible.

For every demand $k \in K$ that connects nodes s and t we use the first m shortest simple paths between s and t in the physical layer $G = (V, E)$ as candidate paths. If the demand allows a maximum delay of $\tau_{k,\max}$, then only the subset of the m paths that fulfill this delay requirement is used. We denote the set of these valid paths by L_k.

We represent a candidate path $l \in L_k$ as a tuple of physical links, i.e., $l = (e_1, e_2, \ldots)$ where $e_i \in E$. In order to realize such a candidate path l one or more consecutive optical circuits need to be set up. An optical circuit can span one or more physical links; therefore, if we split l into several sub-tuples we obtain a sequence of optical circuit paths that realize l. We find all possible sequences of optical circuit paths and denote them by $R_{k,l}$. Sequences containing circuit paths that are longer than a predefined transparent optical reach are excluded. We denote the set of all individual circuit paths by C.

An example is given in Fig. 2 where we assume a demand from node a to c. Two simple paths, a–c and a–b–c, connect the two nodes. The path a–c can only be realized by a circuit connecting nodes a and c directly. The path a–b–c can be realized by a single circuit that bypasses node b or by two consecutive circuits, a–b and b–c.

3.2 ILP Formulation

The goal of the ILP is the joint minimization of demand blocking, delay overfulfillment and required router line cards. To this end, we define variables that encode the selected candidate paths as well as the resulting amount of circuits and hardware that is required. Specifically, we define

- $g_{k,l,r} \in \{0,1\}$ $(k \in K, l \in L_k, r \in R_{k,l})$ as path selector, i.e., $g_{k,l,r}$ is one if demand k is routed over path l using circuit realization r;
- $b_k \in \{0,1\}$ $(k \in K)$ as indicator for a blocked demand;
- $w_c \in \mathbb{N}$ $(c \in C)$ as the number of parallel circuits on circuit path c where C is the set of all circuit paths in the network;
- $w'_c \in \mathbb{N}$ $(c \in C)$ as the number of circuits on circuit path c taking migration into account;
- $p_{v,t} \in \mathbb{N}$ $((v,t) \in U)$ as the number of ports at v for connections with t;
- $q_v \in \mathbb{N}$ $(v \in V)$ as the number of line cards at node v;
- $i_e \in \{0,1\}$ $(e \in E)$ as indicator for a highly utilized link.

The objective function is

$$\min \left(\alpha \sum_{k \in K} b_k + \beta \sum_{v \in V} q_v + \frac{\gamma}{|K_\mathrm{S}|} \sum_{k \in K_\mathrm{S}, l \in L_k, r \in R_{k,l}} g_{k,l,r} \cdot \varDelta\tau_{k,l} + \frac{\mu}{|E|} \sum_{e \in E} i_e \right). \quad (2)$$

The first term penalizes blocked demands while the second term is responsible for the minimization of active line cards. The third term penalizes the average relative delay overfulfillment of the demands in $K_\mathrm{S} \subseteq K$. The parameter $\varDelta\tau_{k,l}$ is demand k's relative delay overfulfillment if it takes path l. The last term penalizes links that are highly utilized. This is important for the make before break migration approach because it encourages the optimizer to leave some free link capacity for later migration steps.

The following constraints complete the ILP.

$$\sum_{l \in L_k, r \in R_{k,l}} g_{k,l,r} = 1 - b_k \quad \forall k \in K \quad (3)$$

$$\sum_{k \in K, l \in L_k, r \in R_{k,l}} \rho_{c,k,l,r} \cdot h_k \cdot g_{k,l,r} \leq \xi \cdot w_c \quad \forall c \in C \quad (4)$$

$$\sum_{k \in K, l \in L_k, r \in R_{k,l}} \rho_{c,k,l,r} \cdot h_k \cdot g_{k,l,r} + \sum_{k \in \widehat{K}, l \in L_k, r \in R_{k,l}} \rho_{c,k,l,r} \cdot h_k \cdot \widehat{g}_{k,l,r}$$
$$- \sum_{k \in K \cap \widehat{K}, l \in L_k, r \in R_{k,l}} \rho_{c,k,l,r} \cdot h_k \cdot g_{k,l,r} \cdot \widehat{g}_{k,l,r} \leq \xi \cdot w'_c \quad \forall c \in C \quad (5)$$

$$\sum_{c \in C} \delta_{e,c} \cdot \xi \cdot w'_c \leq \pi_e \quad \forall e \in E \quad (6)$$

$$\frac{1}{\pi_e} \sum_{c \in C} \delta_{e,c} \cdot \xi \cdot w_c \leq \psi + i_e \quad \forall e \in E \quad (7)$$

$$\sum_{c \in C} w_c \cdot \varphi_{c,v,t} \leq p_{v,t} \quad \forall (v,t) \in U \quad (8)$$

$$\sum_{c \in C} w_c \cdot \varphi_{c,t,v} \leq p_{v,t} \quad \forall (v,t) \in U \quad (9)$$

$$\sum_{t \in V \setminus \{v\}} p_{v,t} \leq \kappa \cdot q_v \quad \forall v \in V \quad (10)$$

In these constraints the set $\widehat{K}$ contains the demands of the previous reconfiguration step. Additionally, we define the parameters

- $\xi \in \mathbb{R}$ as the capacity of a single optical circuit;
- $\pi_e \in \mathbb{R}$ as the capacity of link e;
- $\psi \in [0, 1]$ as the utilization threshold for highly utilized links;
- $\kappa \in \mathbb{N}$ as the number of router ports a line card can maximally hold;
- $\rho_{c,k,l,r} \in \{0, 1\}$ as indicator showing whether circuit realization r of demand k's candidate path l uses circuit path c;
- $\delta_{e,c} \in \{0, 1\}$ as indicator showing whether circuit path c traverses link e;
- $\varphi_{c,u,v} \in \{0, 1\}$ as indicator showing whether circuit path c connects nodes u and v with u as source node;
- $\widehat{g}_{k,l,r} \in \{0, 1\}$ as the optimal path selector of the previous configuration.

Equation (3) ensures that a demand is either routed on exactly one path or blocked. Constraints (4) and (5) ensure that enough circuits are installed depending on the chosen paths. In contrast to Constraint (4), which only considers the demands of the current reconfiguration step, Constraint (5) also takes the routing of the previous reconfiguration step into account allowing the make before break migration. (6) is a link capacity constraint and (7) triggers the indicator i_e if the utilization of link e exceeds a certain limit $\psi \in [0, 1]$. Constraints (8) and (9) reserve a sufficient amount of router ports to accommodate the optical circuits. Finally, (10) ensures that enough line cards are present to hold the ports.

4 Evaluation

In this section we evaluate our approach. We first present the topologies and traffic demands as well as the parameters we studied. Then we discuss our results.

4.1 Network Topologies and Architecture

We selected the two backbone networks Abilene and Géant, found in the SNDlib [9], to evaluate the behavior of our routing approach in different wide-area networks. We only show details for the Géant topology here since the results for Abilene are very similar.

Géant is a research network connecting various countries in Europe. As we neglect the connection to New York, we end up with 21 nodes and 68 directed links. The topology is shown in Fig. 3(a). We assume that the ports operate at a data rate of $\xi = 100$ Gbps with a transparent reach of 2500 km and that a router line card can hold a single port, i.e., $\kappa = 1$ [4]. Further, we assume that the number of wavelengths on a link is limited such that $\pi_e = 40 \cdot \xi$.

4.2 Traffic Demand

We synthetically generated traffic demands whose arrival and holding times follow a negative exponential distribution. We assume that the demands have wavelength granularity, i.e., the traffic value of a demand h_k equals the capacity ξ

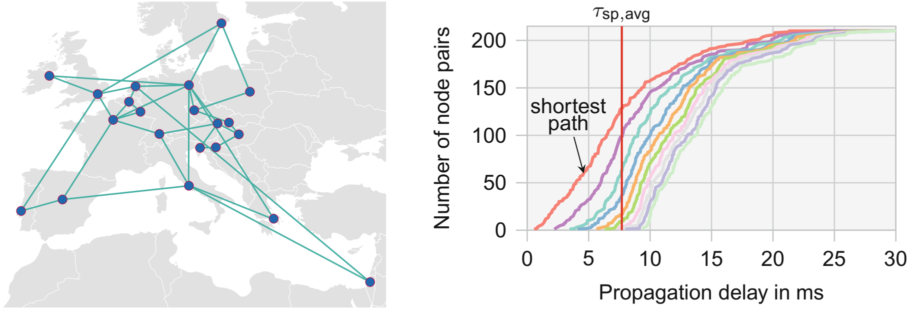

Fig. 3. Logical fiber topology (a) and propagation delay distributions of the 10 shortest paths (b) for the European research network Géant [9] without the node in New York. The path length index in (b) increases from left to right, i.e., the red curve on the left corresponds to the shortest path. The average delay of the shortest path at $\tau_{\mathrm{sp,avg}} = 7.7\,\mathrm{ms}$ is depicted with the red vertical line. (Color figure online)

of an optical circuit. We distinguish two traffic demand classes, namely *delay-sensitive* and *delay-insensitive* demands. The delay-insensitive demands do not have a delay requirement meaning that they can be routed on arbitrary routes. The delay-sensitive demands do have a certain delay requirement $\tau_{k,\mathrm{max}}$. In the ILP optimization they form the set K_{S}, i.e., their delay overfulfillment is minimized. Delay-sensitive demands were generated only between those node pairs for which the shortest connecting path satisfies the delay requirement. Varying the delay requirement therefore also changes the number of node pairs with delay-sensitive traffic. Figure 3(b) shows the delay distribution for the ten shortest paths between all node pairs. Additionally, the average shortest path delay $\tau_{\mathrm{sp,avg}}$ considering all node pairs is depicted. It is visible that for a delay requirement $\tau_{k,\mathrm{max}}$ equal to the average shortest path delay $\tau_{\mathrm{sp,avg}}$, 128 node pairs can be connected with a delay-sensitive demand. 98 of them have at least one alternative path. The remaining 30 node pairs can be connected on the shortest path only. Our approach tries to avoid the blocking of demands between those node pairs by routing demands with less strict delay requirements on circumjacent links.

4.3 Reconfiguration

The reconfiguration is triggered at regular intervals where the ratio of the interval duration and the mean holding time of the traffic demands equals 0.05. In each reconfiguration step we route the demands that will arrive in the following interval while taking the present configuration of the current interval into account. We employ a make before break migration to transition between two configurations. For the candidate path generation we consider the $m = 10$ shortest paths between each node pair. For each path we consider all possible circuit realizations. We consider a link to be highly loaded if its utilization exceeds a value of $\psi = 0.95$. For the objective function of the ILP we used the parameters $\alpha = 10\,000$ and $\mu = 1000$. In that way, avoiding unrouted demands is the most important optimization goal followed by the prevention of highly utilized links. By adjusting the ratio of β to γ the trade-off between the minimization of overfulfillment and active line cards can be controlled. We set $\beta = 0.0001$ and $\gamma = 10$ to focus on the maximum achievable overfulfillment reduction in this scenario.

4.4 Varied Scenario Parameters

For the evaluation of the introduced reconfiguration approach we adjusted a number of parameters relevant for the traffic generation.

The *offered load* describes the amount of traffic that needs to be transported through the network in relation to its capacity. Therefore, on the one hand, the offered load depends on the demand values themselves. But on the other hand it also depends on the path length of a routed demand. Demands between nodes with a physical connection can be routed single-hop whereas demands between distant nodes necessarily occupy resources on several links. We therefore introduce an offered load metric which, to a certain extent, takes these differences in the path length into account. For each demand k we compute the shortest possible path and count the number of physical links ζ_k it traverses. In our model, this value ζ_k is equivalent to the number of wavelengths the demand occupies when it is routed on the shortest path. We multiply the occupied wavelengths with the demand value h_k and relate it to the total network capacity, i.e., the sum of all link capacities. The offered load is then given by

$$A = \frac{\sum_{k \in K} \zeta_k \cdot h_k}{\sum_{e \in E} \pi_e}. \tag{11}$$

An additional parameter is the *share of delay-sensitive demands* in the total number of demands $\phi = |K_\mathrm{S}|/|K|$. Lastly, we vary the target delay $\tau_{k,\max}$ of the delay-sensitive demands. In order to compare networks with different geographical layouts we relate the target delay to the average shortest path delay $\tau_{\mathrm{sp,avg}}$ of the network. We introduce a *target delay factor* χ_τ which adjusts the target delay according to $\tau_{k,\max} = \chi_\tau \cdot \tau_{\mathrm{sp,avg}}, \forall k \in K_\mathrm{S}$.

Using the introduced parameters we have determined demand inter-arrival and holding times. Based on these, we generated a traffic demand series of 200 demand sets for each parameter combination of offered load A, share of delay-sensitive demands ϕ and target delay factor χ_τ as depicted in Fig. 4.

4.5 Results

The presented approach on the reduction of delay overfulfillment is compared to a resource optimization. The resource optimization routes the traffic load such that the amount of active line cards is minimized while the delay requirements are satisfied. This is realized by setting $\gamma = 0$ in the ILP's objective function.

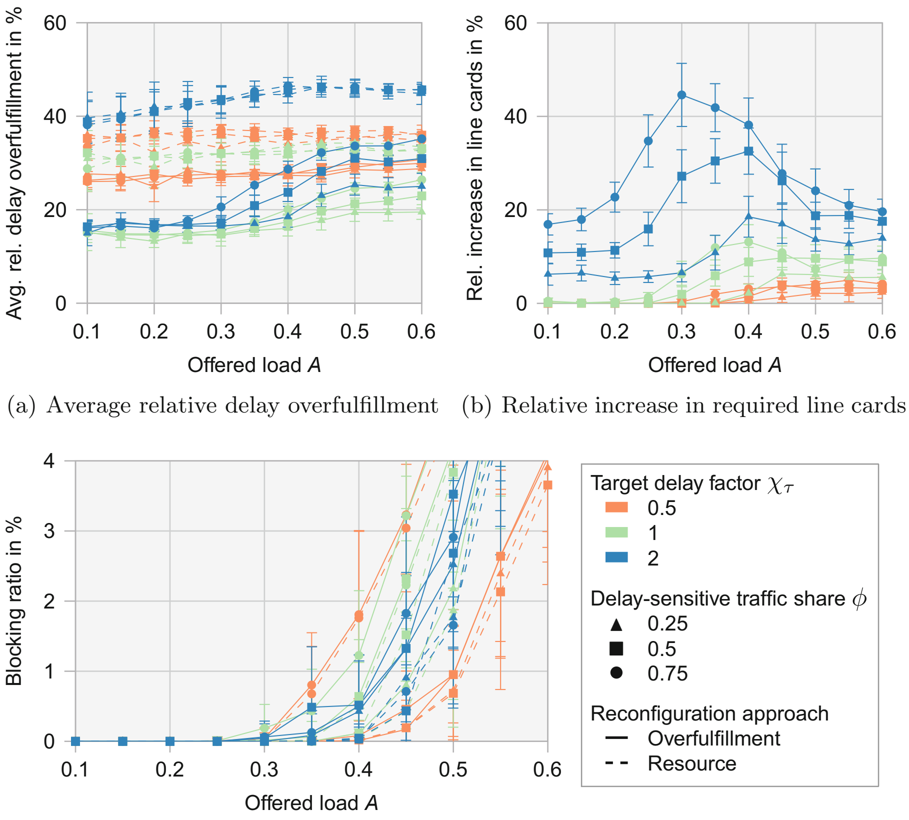

(a) Average relative delay overfulfillment (b) Relative increase in required line cards

(c) Blocking ratio

Fig. 4. Results of the introduced reconfiguration approach for wavelength demands with Poisson arrivals and departures. Shown metrics are the average relative delay overfulfillment (a), the relative increase in required line cards compared to a delay-aware resource optimization (b) and the blocking ratio (c). The error bars represent one standard deviation.

Figure 4 shows the results of our investigation as introduced for the Géant topology for an offered load ranging from 0.1 to 0.6 in steps of 0.05. Each data point represents the average of a series of 200 reconfigurations including migration. The error bars show one standard deviation. Figure 4(a) shows the relative delay overfulfillment averaged over all delay-sensitive demands K_S. As can be seen, our approach is able to reduce the relative overfulfillment by up to 60% compared to the resource optimization for the case that the target delay factor is 2. Even for a delay factor of 0.5, for which alternative paths for delay-sensitive demands are scarce, an overfulfillment reduction is possible. The share of delay-sensitive demands has only a small impact on the delay overfulfillment. Figure 4(c) shows the blocking ratio, i.e., the ratio of unrouted demands to the total number of demands K. Blocking occurs only if all wavelengths on an physical link are occupied and no alternative route is available. As can be seen, blocking occurs for all parameter combinations if the offered load exceeds a value of 0.4. This effect is also visible with the resource optimization which suggests that an offered load of 0.4 is already close to the capacity limit of the network. For values below 0.4 we can see that the overfulfillment minimization exhibits slightly higher blocking ratios. Figure 4(b) shows the relative increase in required line cards for the overfulfillment reducing optimization compared to the resource optimization. For a target delay factor of $\chi_\tau = 0.5$ and a traffic load below 0.3 no increase in the number of line cards is visible. Also for $\chi_\tau = 1$ and traffic loads up to 0.2 this number stays the same for both optimization approaches. This means that the delay overfulfillment can be reduced without the need to invest more line cards. For a target delay of twice the average shortest path delay additional line cards are necessary. However, we assume that this hardware overhead is within the margin of hardware overprovisioning which is typically found in transport networks. The amount of additional line cards depends both on the share of delay-sensitive demands and the offered load.

We have conducted the evaluation for the US network Abilene [9] as well. Compared to Géant, the Abilene network has a broader delay distribution resulting in less alternative paths for delay-sensitive demands. Hence, for a delay factor of 0.5 the overfulfillment cannot be reduced significantly compared to the resource optimization. For higher target delays the results are very similar to those presented for the Géant topology. In particular, they confirm that the delay overfulfillment can be significantly reduced by our approach while the blocking behavior is only slightly affected.

5 Conclusions

In this paper we presented a new routing and network reconfiguration approach which addresses the overfulfillment of delay requirements in transport networks. The approach routes traffic on paths other than the shortest path to reduce the delay overfulfillment of traffic demands. In that way, customers experience service delays closer to the limits specified in their SLA providing an incentive to acquire a more expensive service class if necessary. For the ISP this provides a new opportunity for revenue creation. As a secondary goal, our approach relieves highly utilized links for those demands that have no alternative to the shortest path due to their strict delay requirements.

We presented a parameterizable ILP that jointly optimizes routing in the optical and electrical layer of the transport network. We evaluated our proposed approach for two nationwide transport networks. The results show that the delay overfulfillment can be reduced by up to 60% without a significant increase in the blocking ratio compared to a resource optimization. The number of required line cards naturally increases because longer routes are taken. However, we assume it to remain within the margin of hardware overprovisioning by ISPs. Furthermore, the trade-off between the minimization of overfulfillment and required hardware depends on adjustable parameters which allows the ISP to adapt our approach to its business models and networking hardware.

Acknowledgements. This work has been performed in the framework of the CELTIC EUREKA project SENDATE-TANDEM (Project ID C2015/3-2), and it is partly funded by the German BMBF (Project ID 16KIS0458). The authors alone are responsible for the content of the paper.

References

1. Bauknecht, U.: Resource efficiency and latency in dynamic IP-over-WSON networks utilizing flexrate transponders. In: Photonic Networks; 18. ITG-Symposium, pp. 36–41, May 2017
2. Bozkurt, I.N., et al.: Dissecting latency in the internet's fiber infrastructure. arXiv preprint arXiv:1811.10737, November 2018
3. Choi, B.Y., Moon, S., Zhang, Z.L., Papagiannaki, K., Diot, C.: Analysis of point-to-point packet delay in an operational network. Comput. Netw. **51**(13), 3812–3827 (2007). https://doi.org/10.1016/j.comnet.2007.04.004
4. Cisco: CRS 1-port 100 Gigabit Ethernet Coherent DWDM Interface Module (2013)
5. Cisco: Visual Networking Index: Forecast and Trends, 2017–2022, November 2018
6. Ellis, A.D., McCarthy, M.E., Khateeb, M.A.Z.A., Sorokina, M., Doran, N.J.: Performance limits in optical communications due to fiber nonlinearity. Adv. Opt. Photonics **9**(3), 429–503 (2017). https://doi.org/10.1364/AOP.9.000429
7. Enderle, T., Bauknecht, U.: Modeling dynamic traffic demand behavior in telecommunication networks. In: Photonic Networks; 19th ITG-Symposium, pp. 18–25, June 2018
8. Lourenço, R.B., Tornatore, M., Martel, C.U., Mukherjee, B.: Running the network harder: connection provisioning under resource crunch. IEEE Trans. Netw. Serv. Manage. **15**(4), 1615–1629 (2018). https://doi.org/10.1109/TNSM.2018.2875103
9. Orlowski, S., Pióro, M., Tomaszewski, A., Wessäly, R.: SNDlib 1.0-survivable network design library. In: Proceedings of the 3rd International Network Optimization Conference (INOC 2007), Spa, Belgium, April 2007
10. Papanikolaou, P., Christodoulopoulos, K., Varvarigos, M.: Multilayer flex-grid network planning, pp. 151–156 (2015). https://doi.org/10.1109/ONDM.2015.7127290
11. Pointurier, Y.: Design of low-margin optical networks. IEEE/OSA J. Opt. Commun. Netw. **9**(1), A9–A17 (2017). https://doi.org/10.1364/JOCN.9.0000A9
12. Sen, S., Joe-Wong, C., Ha, S., Chiang, M.: A survey of smart data pricing: past proposals, current plans, and future trends. ACM Comput. Surv. **46**(2), 15:1–15:37 (2013). https://doi.org/10.1145/2543581.2543582
13. Velasco, L., Castro, A., King, D., Gerstel, O., Casellas, R., Lopez, V.: In-operation network planning. IEEE Commun. Mag. **52**(1), 52–60 (2014). https://doi.org/10.1109/MCOM.2014.6710064
14. Walrand, J.: Economic models of communication networks. In: Liu, Z., Xia, C.H. (eds.) Performance Modeling and Engineering, pp. 57–89. Springer, Boston (2008). https://doi.org/10.1007/978-0-387-79361-0_3
15. Xiao, P., Li, Z., Guo, S., Qi, H., Qu, W., Yu, H.: A K self-adaptive SDN controller placement for wide area networks. Front. Inf. Technol. Electron. Eng. **17**(7), 620–633 (2016). https://doi.org/10.1631/FITEE.1500350