Traffic Protection in Multilayer Core Networks by Optimum Thinning of MPLS Tunnel Capacities

Ronald Romero Reyes
Chair for Communication Networks
Technische Universität Chemnitz
Chemnitz, Germany
reyes@hrz.tu-chemnitz.de

Shkurte Esati
Chair for Communication Networks
Technische Universität Chemnitz
Chemnitz, Germany
shkurte.esati@s2019.tu-chemnitz.de

Thomas Bauschert
Chair for Communication Networks
Technische Universität Chemnitz
Chemnitz, Germany
thomas.bauschert@etit.tu-chemnitz.de

Abstract—**In this paper we study a cost-efficient multilayer restoration (MR) approach for IP/MPLS-Optical networks. In these networks, the capacity of the IP links is reduced by optical link failures that include fibre cuts and lightpath signal-to-noise ratio (SNR) degradations. For instance, when a fibre is cut, the lightpaths that traverse the fibre are rerouted. In most cases, these lightpaths must reduce their line-rates to adapt to the new path lengths, which reduces the capacity of the IP links carried by the lightpaths. This causes a failure state in the IP/MPLS layer characterized by multiple IP links with reduced capacities. When this occurs, the solution used today is to shut off these IP links and restore the traffic by applying IP/MPLS rerouting. This method is costly as it increases the network spare capacity. We study an alternative strategy based on adjustable robust optimization. By this approach, the affected IP links are not torn down. Instead, the MPLS tunnels routed over these links thin their capacities to adapt to the reduced IP link capacities. Then the stranded IP traffic is routed through MPLS tunnels set up on alternative paths. The thinning and routing of MPLS tunnels is planned for foreseen optical link failures. Robustness against unforeseen failures is provided by adjustable decision rules. The efficiency of the strategy is assessed for a German backbone network. The capacity requirements attained by the method are compared with selected MR strategies. The results show that the proposed approach yields capacity savings of least 22% w.r.t the traditional shut off IP links method.**

Keywords—Multilayer networks, traffic protection, IP/MPLS-Optical networks, robust optimization, flow thinning.

I. INTRODUCTION

Network planning with multilayer restoration (MR) is key for the design of low cost IP/MPLS-Optical networks with fast recovery capabilities. Today, multilayer networks use separate and uncoordinated restoration methods in the optical and the IP/MPLS layers. Pure optical restoration reduces the network capacity requirements, but it is slow and unaware of failures in the IP/MPLS layer (e.g. faulty routers/IP ports). IP/MPLS restoration, on the other hand, is faster and may recover from failures in the IP/MPLS and the optical layers. This approach, however, overprovisions the capacity of the IP links, which is needed to protect against optical failures. This results in costly solutions [1]. Fast restoration at minimum capacity costs can be achieved by MR, which combines the recovery capabilities of both layers.

In this paper we show that adjustable robust optimization [2] enables capacity-efficient MR in multilayer networks that use flex-grid wavelength division multiplexing (WDM) in the optical layer. In these networks, the capacity of the IP links is supplied by bandwidth variable transponders (BVTs) that set up lightpaths - in the WDM layer - between the routers of the IP links. Typical failure types in the WDM layer that affect the IP links are degradations of the signal-to-noise ratio (SNR) of the lightpaths and fibre cuts. When a failure of either type occurs, the BVTs may have to reduce their line-rates, thereby resulting in an IP/MPLS topology with multiple IP links with reduced capacities. (To recover from fibre cuts, the affected lightpaths are rerouted over paths which are usually longer than the working paths. Thus, BVTs reduce their line-rates to adapt their transparent reach to the new path length.) To recover from both failure types, the strategy applied today is IP/MPLS restoration, i.e. the IP links with reduced capacities are shut off and IP routing is used to offload the stranded traffic. This approach requires a lot of spare capacity [3]. We show that adjustable robust optimization enables the design of more efficient strategies that reduce the capacity costs while avoiding the disconnection of the affected IP links. Instead, the MPLS tunnels routed over these links thin their capacities to adapt to the reduced IP link capacities. Then the stranded IP traffic is routed over MPLS tunnels set up on alternative paths. The thinning and routing of the tunnels is planned for foreseen optical failures. Robustness to unforeseen failures is achieved by adjustable decision rules. By this approach, IP traffic can be protected from single/multiple fibre cuts and from SNR degradations that reduce the lightpath line-rates. This type of failures is studied in [4].

Extensive literature exists about the MR problem – see the work compiled in [5-10]. Most of this work is only applicable to networks in which the IP links totally lose their capacities in case of optical link failures, with single fibre cuts being the dominant use case. The case of IP/MPLS-Optical networks in failure states with multiple IP links with reduced capacities - i.e. partial link failures - remains nearly unexplored. Recently, adjustable robust optimization has been applied in [11-12] for traffic recovery in networks with variable and continuous link capacities. The results in [12] show that the method efficiently protects wireless networks against varying weather conditions that reduce the link capacities. Motivated by these results, in this paper we extend the work in [11-12] by formulating a MR approach for IP/MPLS-Optical networks with modular link capacities, i.e. the capacity of the IP links is supplied by an integer number of BVT modules. The goal is to harness the flexibility offered by the BVTs [13-15] to protect IP traffic from optical failures that cause multiple IP links with reduced capacities. We focus our study on the capacity requirements of the approach and compare its capacity savings with the strategies in [3], where heuristic IP-routing algorithms are proposed for MR in IP/MPLS-Optical networks.

This paper is organized as follows. Section II outlines the problem statement. Section III formulates MR as an adjustable robust optimization problem. In Section IV implementation aspects of the approach are presented. Section V presents a case study and Section VI concludes the paper.

This work was performed in the framework of the Celtic-Plus project AI-NET-PROTECT, funded by the German BMBF (ID 16KIS1285).

978-3-903176-33-1 © 2021 IFIP

II. PROBLEM STATEMENT

For the sake of clarity, all variables, parameters and sets defined in this and the following sections are listed in Table I.

Consider an IP/MPLS network represented by the graph $G(N, E)$, where N is the set of IP/MPLS routers and E is the set of IP links. The network carries point to point demands d defined in a demand set D. MPLS multipath routing is applied to carry the traffic volume $h(d)$ of the demand d over paths p defined in the set of paths $P(d)$. Let s_0 be the state of the IP/MPLS network in failure-free mode of operation. This state is defined by the network topology and the capacities of the IP links. In this state, for each demand d, an MPLS tunnel with capacity $c_{dp}^{s_0}$ is established on each path p in $P(d)$ to carry the portion of $h(d)$ routed over p. The tunnel capacities fulfil the constraint: $\sum_{p\in P(d)} c_{dp}^{s_0} \geq h(d)$, which ensures that the traffic volume $h(d)$ is carried by the MPLS tunnels.

The IP links $e \in E$ are realized by lightpaths set up on a WDM optical network. In failure-free mode, i.e. in state s_0, an IP link e has nominal capacity $c_e^{s_0}$, which is supplied by BVTs that provision lightpaths between the pair of IP/MPLS routers interconnected by link e. (Note: An IP link may consist of one or multiple lightpaths which are usually routed over the same physical path on the WDM network.) The BVTs connect to IP/MPLS router interfaces as well as to reconfigurable optical add/drop multiplexers (ROADMs), which switch lightpaths in the optical domain. Let $\mathcal{R}(d, e)$ be the set of paths p in $P(d)$ that traverse IP link e. The nominal capacity $c_e^{s_0}$ of this link supplies the capacities $c_{dp}^{s_0}$ of the MPLS tunnels traversing the link, which means that: $\sum_{d\in D} \sum_{p\in\mathcal{R}(d,e)} c_{dp}^{s_0} \leq c_e^{s_0}$.

Consider the problem of dimensioning the capacity of the IP/MPLS network so that it copes with a given set $\mathcal{F}$ of failure events that stem from optical link failures (e.g. the elements in $\mathcal{F}$ may define fibre cuts and/or lightpath SNR degradations). When a failure in $\mathcal{F}$ occurs, the BVTs reduce their line-rates to adjust to the restoration path lengths (in case of fibre cuts) or to the new SNR constraints (in case of SNR degradations). As a result, the nominal capacities $c_e^{s_0}$ of the affected IP links are reduced. To model this, let s be the state of the IP/MPLS network right after the BVTs reduce their line-rates. This state is given by the capacities of the IP links and the topology of the IP/MPLS network. In state s, the affected IP links achieve capacities: $c_e^s = \alpha(e, s) \cdot c_e^{s_0}$, where $0 \leq \alpha(e, s) < 1$, is the capacity reduction factor of IP link e in state s. This factor is the ratio of the link capacity after WDM restoration to the link capacity in failure-free mode. In s_0, $\alpha(e, s_0) = 1$ for all IP links. Each failure in $\mathcal{F}$ causes state transitions $s_0 \rightarrow s$, in which the IP/MPLS network has one or multiple IP links with reduced capacities. The states s are defined in the set S.

TABLE I. DEFINITION OF MIP VARIABLES, PARAMETERS AND SETS

Notation	***Definition***
$e \in E$	IP link e, where E is the set of IP links
$d \in D$	Point to point IP demand d, where D is the set of demands
$p \in P(d)$	Path p in the IP topology, where $P(d)$ is the set of paths that can be used to set up the MPLS tunnels that carry demand d
$E(d,p)$	Set of IP links traversed by path $p \in P(d)$, $E(d,p) \subseteq E$
$\mathcal{R}(d,e)$	Set that contains the paths $p \in P(d)$ that traverse IP link e
$h(d)$	Traffic volume of demand d. This traffic can be split and carried over the MPLS tunnels set up over the paths $p \in P(d)$
$s \in S$	Failure state s of the IP layer. This state is reached after the WDM layer recovers (by adjusting the BVT line-rates) from optical link failures. In this state one or multiple IP links reduce their capacities. The set of failure states is S
$S(e)$	Set of failure states $s \in S$ that reduce the capacity of IP link e
$\alpha(e,s)$	Capacity reduction factor of IP link e in failure state s
$a_{ek} \in \mathbb{Z}^+$	Number of transponders of type $k \in K$ assigned to IP link e. K is the set of BVT transponder modules
$u_k \geq 0$	Maximum capacity of a transponder module (or BVT) of type $k \in K$, e.g. $u_k \in \{1.25\text{G}, 2.5\text{G}, 10\text{G}, 40\text{G}, 100\text{G}\}$
$c_e^{s_0}$	Nominal capacity of IP link e in failure-free mode, i.e in state s_0, it is calculated as $c_e^{s_0} = \sum_{k\in K} a_{ek} \cdot u_k$
c_e^s	Capacity of IP link e in failure state s, with $c_e^s = \alpha(e,s) \cdot c_e^{s_0}$
$c_{dp}^{s_0}$	Capacity in failure-free mode of the MPLS tunnel established on path p for demand d
c_{dp}^s	Thinned or reduced capacity in failure state s of the MPLS tunnel established on path p for demand d, where $c_{dp}^s < c_{dp}^{s_0}$
z_{dp}^e	Parameter that accounts for the contribution of IP link e on path $p \in P(d)$ to the thinning of the MPLS tunnel that reduces its capacity from $c_{dp}^{s_0}$ to c_{dp}^s in failure state s

While in state s, the MPLS tunnels set up over the affected IP links may have to thin their capacities $c_{dp}^{s_0}$ to adjust to the reduced link capacities c_e^s. Thus, the state transition: $s_0 \rightarrow s$ causes IP link capacity reductions: $c_e^{s_0} \rightarrow c_e^s$, which enforce the thinning: $c_{dp}^{s_0} \rightarrow c_{dp}^s$ of the capacities of the MPLS tunnels that traverse the affected links. Observe that $c_{dp}^s < c_{dp}^{s_0}$, where c_{dp}^s is the reduced capacity in state s of the MPLS tunnel of demand d routed through path p. For this scenario, the MR problem consists in optimizing the nominal IP link capacities $c_e^{s_0}$, the nominal and reduced tunnel capacities $c_{dp}^{s_0}$ and c_{dp}^s as well as the routing of tunnels to ensure that the traffic volumes $h(d)$ are carried in all states s. Multipath routing is allowed and restricted to the sets of paths $P(d)$. (Note that $P(d)$ and the demand set D are assumed to be known.) The capacities and the routing must be optimized to protect the traffic from all optical failure events in $\mathcal{F}$. These events are predefined and determine the capacity reduction factors $\alpha(e, s)$. The solution to the problem defines the recovery strategy of the IP/MPLS network (i.e. rerouting and thinning of the capacities of the MPLS tunnels) applied right after the WDM network recovers from the optical link failure. An example that illustrates this problem is depicted in Fig. 1 for the case of a fibre cut.

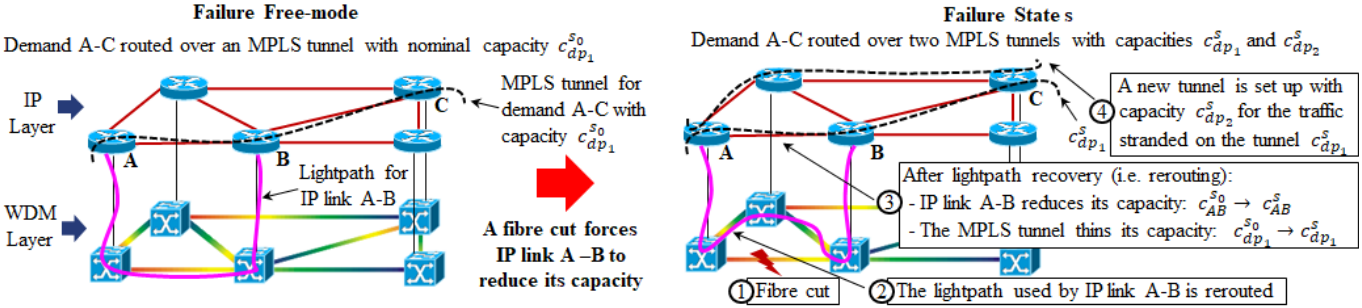

Fig. 1. Example of MR with adaptive routing and capacity thinning.

III. ROBUST NETWORK DIMENSIONING

A. Affine Mixed Integer Program (MIP) Formulation

The work in [11-12] studies flow thinning optimization in networks with variable and continuous link capacities. Based on that work, we formulate a Mixed Integer Linear Program (MIP) for IP traffic protection in IP/MPLS-Optical networks with modular link capacities. Let K be the set of BVT modules that supply the capacities of the IP links. A module of type $k \in K$ represents a BVT transponder with maximum capacity u_k. The number of modules of type k assigned to IP link e is a_{ek}. The nominal capacity of link e is: $c_e^{s_0} = \sum_{k\in K} a_{ek} \cdot u_k$. Assuming that a transponder module of type k has unit cost c_k, the optimization problem is defined as:

$$\text{Minimize: } \sum_{e\in E}\sum_{k\in K} a_{ek}\cdot c_k \qquad \text{(1a)}$$

Subject to:

$$\sum_{d\in D}\sum_{p\in \mathcal{R}(d,e)} c_{dp}^{s} \le \alpha(e,s)\cdot \sum_{k\in K} a_{ek}\cdot u_k\,, \quad e\in E, s\in S(e)\cup\{s_0\}, \qquad \text{(1b)}$$

$$\sum_{p\in P(d)} c_{dp}^{s} \ge h(d)\,,\quad d\in D, s\in S, \qquad \text{(1c)}$$

$$c_{dp}^{s} = c_{dp}^{s_0} - \sum_{e\in E(d,p)}[1-\alpha(e,s)]\cdot z_{dp}^{e}\,, \quad d\in D,\ p\in P(d),\ s\in S, \qquad \text{(1d)}$$

$$c_{dp}^{s_0}\ge 0,\ c_{dp}^{s}\ge 0,\ z_{dp}^{e}\ge 0,\ a_{ek}\in\mathbb{Z}^{+}, \quad d\in D,\ p\in P(d),\ e\in E,\ s\in S,\ k\in K, \qquad \text{(1e)}$$

The objective is to minimize the cost of the network capacity (1a). The capacity constraint (1b) ensures that IP link e has capacity to carry the MPLS tunnels in both failure-free mode (i.e. in s_0) and in any state $s \in S(e)$, where $S(e)$ is the set of states s in S that reduce the capacity of IP link e. For all states in $S(e)$, $\alpha(e,s) < 1$, whereas for the state s_0, $\alpha(e,s_0) = 1$. Constraint (1c) ensures that the MPLS tunnels have capacity in all states s to carry the traffic volume $h(d)$. Constraint (1d) provides robustness of the solution to states s not defined in S, i.e. states not caused by any failure in $\mathcal{F}$. This constraint defines the reduced capacity c_{dp}^{s} as the nominal capacity $c_{dp}^{s_0}$ minus the thinning factor: $\sum_{e\in E(d,p)}[1-\alpha(e,s)]\cdot z_{dp}^{e}$, which is a linear combination of the variables $z_{dp}^{e} \ge 0$, weighted by the coefficients: $1-\alpha(e,s)$. The summation runs over the set $E(d,p)$, which is the set of links e that traverse path p. Thus, the capacity reduction of a tunnel is modelled as dependent on the thinning of the IP links it traverses. Constraint (1e) defines the decision variables $c_{dp}^{s_0}$, c_{dp}^{s}, z_{dp}^{e} as positive real numbers and a_{ek} as positive integers. The problem consists in defining the values of these variables that minimize the capacity costs given that $h(d)$, $P(d)$, $\alpha(e,s)$, c_k and u_k are known. Notice that the optimization of $c_{dp}^{s_0}$ and c_{dp}^{s} implicitly defines the MPLS routing for each demand d.

B. Affinely Adjustable Decision Rule for Robust Recovery

The MIP (1) is an adjustable, *NP*-hard robust optimization problem. It provides robustness to states s not defined in S. For example, states that arise when an unforeseen pattern of fibre cuts occurs or when the network enters a state for which the actual factor $\alpha(e,s)$ differs from the one used to solve (1). To interpret the MIP, let $\mathbf{x}$ be a vector with components given by the variables $c_{dp}^{s_0}$ and a_{ek}, and let $\mathbf{y}$ be the vector whose components are defined by the variables c_{dp}^{s}. The vector $\mathbf{x}$ defines here-and-now decisions, i.e. variables which once optimized remain fixed regardless of the state s observed in network operation. On the contrary, $\mathbf{y}$ defines wait-and-see decisions, i.e. variables that once optimized can be adjusted to the actual states s observed in network operation. The solution to (1) is given by the vectors $\mathbf{x}^*$ and $\mathbf{y}^*$ that minimize the cost. The MIP (1) is an adjustable robust optimization problem which has the generic form [2]:

$$\min_{\mathbf{x},\mathbf{y}}\{\textstyle\sum_{e,k} a_{ek}\cdot c_k :\ \mathbf{A}(\boldsymbol{\alpha})\cdot\mathbf{x} + \mathbf{y}(\boldsymbol{\alpha}) \le \mathbf{d},\ \forall\boldsymbol{\alpha}\} \qquad (2)$$

where $\mathbf{A}(\boldsymbol{\alpha})$ and $\mathbf{d}$ are known and represent, respectively, the matrix and vector of coefficients of the optimization problem. The components of the vector $\boldsymbol{\alpha}$ are the factors $\alpha(e,s)$. In (1)-(2), the dependency of the problem on the states s is given by the factors $\alpha(e,s)$. Therefore, robustness to unknown failure states s is accomplished by providing robustness to unknown factors $\alpha(e,s)$. Notice that c_{dp}^{s} depends on $\alpha(e,s)$, which in (2) is pointed out by the functional dependency: $\mathbf{y} = \mathbf{y}(\boldsymbol{\alpha})$.

It is difficult to optimize (1)-(2) over the set of all possible functions $\mathbf{y}(\boldsymbol{\alpha})$. To circumvent this complexity, the theory of adjustable robust optimization shows that linear (i.e. affine) or quadratic parameterized function approximations yield good solutions [2]. Based on this, in (2), $\mathbf{y}(\boldsymbol{\alpha})$ can be approximated by the parameterized function:

$$\mathbf{y}(\boldsymbol{\alpha}) = \mathbf{y}^0 + \mathbf{Z}\cdot\ \boldsymbol{\alpha} \qquad (3)$$

which is linear (affine) in $\mathbf{Z} = [z_{dp}^{e}]$. The components z_{dp}^{e} of the matrix $\mathbf{Z}$ are the parameters of the affine function and the offset values are defined by the vector $\mathbf{y}^0$, whose components are the nominal capacities $c_{dp}^{s_0}$ of the MPLS tunnels. With the affine approximation (3), the MIP in (2) is reformulated as:

$$\min_{\mathbf{x},\mathbf{Z}}\{\textstyle\sum_{e,k} a_{ek}\cdot c_k : \mathbf{A}(\boldsymbol{\alpha})\cdot\mathbf{x} + \mathbf{y}^0 + \mathbf{Z}\cdot\ \boldsymbol{\alpha} \le \mathbf{d},\ \forall\boldsymbol{\alpha}\} \qquad (4)$$

which is the compact form of the MIP (1), where the problem is to calculate the vector $\mathbf{x}^*$ (i.e. the here-and-now variables $c_{dp}^{s_0}$ and a_{ek}) and the matrix of parameters $\mathbf{Z}^*$ that minimize the capacity cost. It is worth emphasising that $\boldsymbol{\alpha}$ is known and defined by the factors $\alpha(e,s)$ of the states in S. From $\mathbf{x}^*$ and $\mathbf{Z}^*$, the optimum vector $\mathbf{y}^*(\boldsymbol{\alpha})$ of wait-and-see variables (i.e. the thinned capacities c_{dp}^{s}) is calculated with (3). The explicit form of (3) is given by constraint (1d).

The solution to (1) is calculated offline and used for traffic protection in the states $s \in S$ visited in network operation. If an unforeseen failure causes the IP/MPLS network to enter a state $s' \notin S$, the recovery works follows. The WDM network performs optical restoration, adjusts the BVT line-rates, and informs the IP/MPLS network of the new factors $\alpha(e,s')$. These factors are used to calculate the capacities $c_{dp}^{s'}$ of the affected tunnels. For that, there is no need to solve (1) by including s' in S. Instead, the already optimized values of $c_{dp}^{s_0}$ and z_{dp}^{e} are used - along with $\alpha(e,s')$ - to calculate $c_{dp}^{s'}$ online with (1d). Thus, the capacities and routing of the tunnels are adapted according to (1d), which is an affine decision rule that provides robustness to unknown states. An in-depth study of affine and quadratic decision rules for networks with variable and continuous link capacities is found in [12].

IV. PRACTICAL ASPECTS OF MULTILAYER RESTORATION BY THINNING OF MPLS TUNNEL CAPACITIES

Multilayer recovery requires control planes that coordinate the exchange of control information between the WDM and the IP layers. This can be accomplished by software defined transport network (T-SDN) control [16]. To illustrate this, let us consider the multilayer network in Fig. 2a. The controller module represents either a single controller (i.e. monolithic T-SDN control) or separate IP/MPLS, WDM controllers with an orchestrator for multilayer control (e.g. hierarchical T-SDN). At time t_i, the network is sized with the capability to recover from a predefined set $\mathcal{F}$ of optical link failures. The network capacity and the demand routing is planned for traffic volumes $h(d)$ expected over a planning period T. The IP and the WDM network topologies are known. This problem is solved offline by the T-SDN controller in three steps.

First, for each IP link, a shortest path algorithm is applied to calculate a primary path – in the optical topology – to route the lightpaths of the IP link in failure-free mode. Among the BVTs with transparent reach longer than the path length, the BVT with the highest line-rate defines the maximum capacity u_{max} attainable by a lightpath on this path. In the second step, the set of failure states S is derived from $\mathcal{F}$. A failure event in $\mathcal{F}$ may consist of one or multiple faulty optical links that cause a state s. To characterize s, in case of fibre cuts, a shortest path algorithm is used to calculate a restoration path - in the WDM layer - for the lightpaths of the IP links affected. The state s is defined by the factors $\alpha(e,s)$, which for each affected IP link e is given by: $\alpha(e,s) = u_s/u_{max}$, where u_s is the maximum line-rate feasible for a BVT on the restoration path. For failure events in $\mathcal{F}$ due to SNR degradations, $\alpha(e,s)$ can be derived from the expected SNR reductions – see [4] as an example. In the third step, for each demand d, a shortest path algorithm is applied in the IP/MPLS layer to calculate the path sets $P(d)$. These sets and the output from the previous steps define the input for the MIP which is solved by the controller.

The network is installed at time t_i with the number a_{ek} of BVT modules calculated by the MIP. The control plane then configures these BVTs to set up the lightpaths on their primary paths, thereby provisioning the IP links with capacities $c_e^{s_0}$. Next, the MPLS tunnels are set up with the optimum nominal capacities $c_{dp}^{s_0}$. The traffic is routed over the tunnels and the network enters the failure-free state s_0. (The restoration paths in the WDM network are only set up when a fibre cut occurs.) In network operation (see Fig. 2b as an example), c_{dp}^{s} and z_{dp}^{e} are used for traffic protection. In particular, in case of a failure in $\mathcal{F}$, the capacity and routing of the tunnels are adjusted according to c_{dp}^{s}. For failures not defined in $\mathcal{F}$ (as seen in Section III), $c_{dp}^{s_0}$ and z_{dp}^{e} are used to calculate c_{dp}^{s} with (1d).

The MIP (1) enables online adaptability to relevant failure events. For example, if a fibre cut defined in $\mathcal{F}$ occurs and a pre-calculated restoration path is unable to admit an affected lightpath, an alternative path can be calculated. The factors $\alpha(e,s)$ that stem from this path define a new state s to which the network adapts with (1d). Another relevant case includes unforeseen SNR degradations. If the SNR reductions allow the BVTs to adapt their line-rates without tearing down the lightpaths, the resulting factors $\alpha(e,s)$ can be measured to thin the MPLS tunnel capacities with (1d). In this scenario, lightpath rerouting might not be required. Today, networks recover from these failures by shutting-off the affected IP links [4]. In general, the set $\mathcal{F}$ may define a mix of failure event types, e.g. some events that stem from single and/or multiple fibre-cuts, others from SNR degradations, etc. The states caused by these events must in all cases be defined by their associated factors $\alpha(e,s)$.

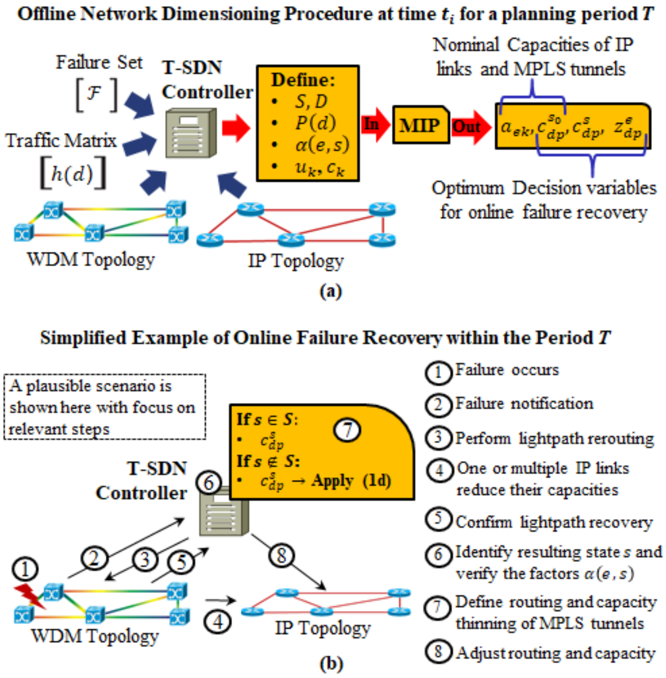

Fig. 2. a) Offline network dimensioning; b) Online recovery example.

V. CASE STUDY

A. Input Information

We consider the network in Fig. 3 which is inspired by a Deutsche Telekom's real multilayer network. The IP/MPLS network consists of 12 routers and 31 IP links. These IP links are realized by lightpaths set up over a WDM network with 12 ROADM nodes and 21 fibre links with the lengths in km shown in Fig. 3. Table 2 lists technical specifications of the BVT modules, i.e. maximum capacity u_k in Gb/s, modulation format, transparent reach (in km) and unit cost c_k in strongest cost units (SCU). An SCU is the cost of a 10 Gb/s BVT with transparent reach of 750 km [17].

B. Methodology

We apply the three-step offline dimensioning procedure outlined in Section IV considering a failure set $\mathcal{F}$ that includes all possible single fibre-cuts. Since the WDM network has 21 fibre links, $\mathcal{F}$ defines 21 failure events, which is the number of states s in S. The primary and the restoration paths in the WDM network are, respectively, the shortest length paths (in

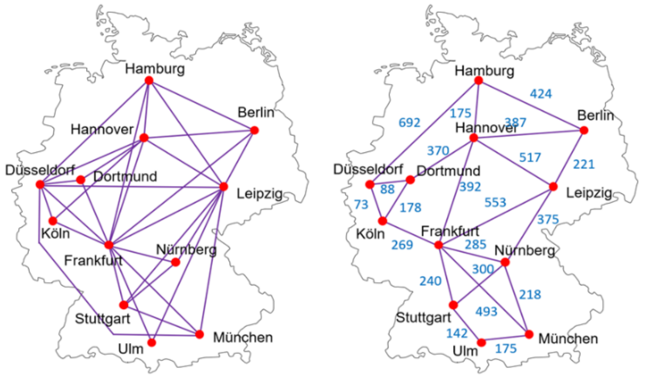

Fig. 3. German Network: IP topology (left) and WDM topology (right).

TABLE II. DEFINITION OF BVT MODULES

BVT Type k	**Technical Parameters**			
	Max. Capacity u_k (Gb/s)	***Modulation format***	***Transparent Reach (Km)***	***Unit cost c_k (SCU)***
Type 1	100	QPSK	2397	15.0
Type 2	150	8 QAM	1096	16.4
Type 3	200	16 QAM	537	16.8
Type 4	250	32 QAM	288	16.9
Type 5	300	64 QAM	135	17.0

km) available in failure-free mode and in the corresponding failure state. Having calculated these paths, for each fibre-cut in $\mathcal{F}$, the transparent reach and the capacity of the BVTs in Table 2 are used to calculate $\alpha(e,s)$ for each IP link e in state s. Moreover, a demand set D with 132 demands is considered. The sum of the traffic volumes is: $\sum_d h(d) = 8.2$ Tb/s. In the IP/MPLS topology, for each demand d, the set of paths $P(d)$ contains at most the four shortest paths calculated between the source and destination routers. The path cost is determined by the hop count. Furthermore, to assess the impact of the traffic load, evaluation scenarios are considered for demand sets D' that contain the demands in D with the traffic volumes scaled by a factor δ, such that: $h'(d) = \delta \cdot h(d)$. Scenarios are defined for $\delta = 0.5, 0.75, 1.0, 1.5, 2.0, 2.5$. Four alternative methods are considered to compare the efficiency of the MIP (1). The methods differ in the recovery strategy applied in the IP/MPLS layer right after the WDM layer reroutes the affected lightpaths:

- *Shut-off affected IP links* tears down the IP links with reduced capacities and offloads - on alternative paths - the IP traffic from the MPLS tunnels set up on those IP links. This is a common strategy used today.
- *No-Reroute* does not offload IP traffic from the MPLS tunnels set up on the affected IP links. Instead, these links are overprovisioned with enough capacity so that the MPLS tunnels cope with any capacity reduction.
- *Smallest Demand First (SDF)* sorts in ascending order - w.r.t the allocated capacity - the MPLS tunnels that traverse each affected IP link. The MPLS tunnels are rerouted from the affected IP links in ascending order (i.e. starting with the tunnel with lowest capacity) until the capacity of the IP links is sufficient to carry the remaining MPLS tunnels.
- *Multiple-Path Reroute* reroutes from the affected IP links the tunnels for which more than one alternative path exists between the source and destination routers. These tunnels are rerouted over the alternative paths which are not affected by the failure. All single-path tunnels remain routed over the affected IP links.

We have studied these approaches in [3], where *Multiple-Path Reroute* and *SDF* are proposed as heuristic strategies for MR from optical link failures. Both heuristics provide good spare capacity savings.

C. *Numerical Results*

Figure 4a shows the nominal network capacity calculated for the proposed approach - referred to as *Affinely-Adjustable MR (AAMR)* - and for the alternative methods. The nominal network capacity is determined by: $C = \sum_{e\in E} \sum_{k\in K} a_{ek} \cdot u_k$.

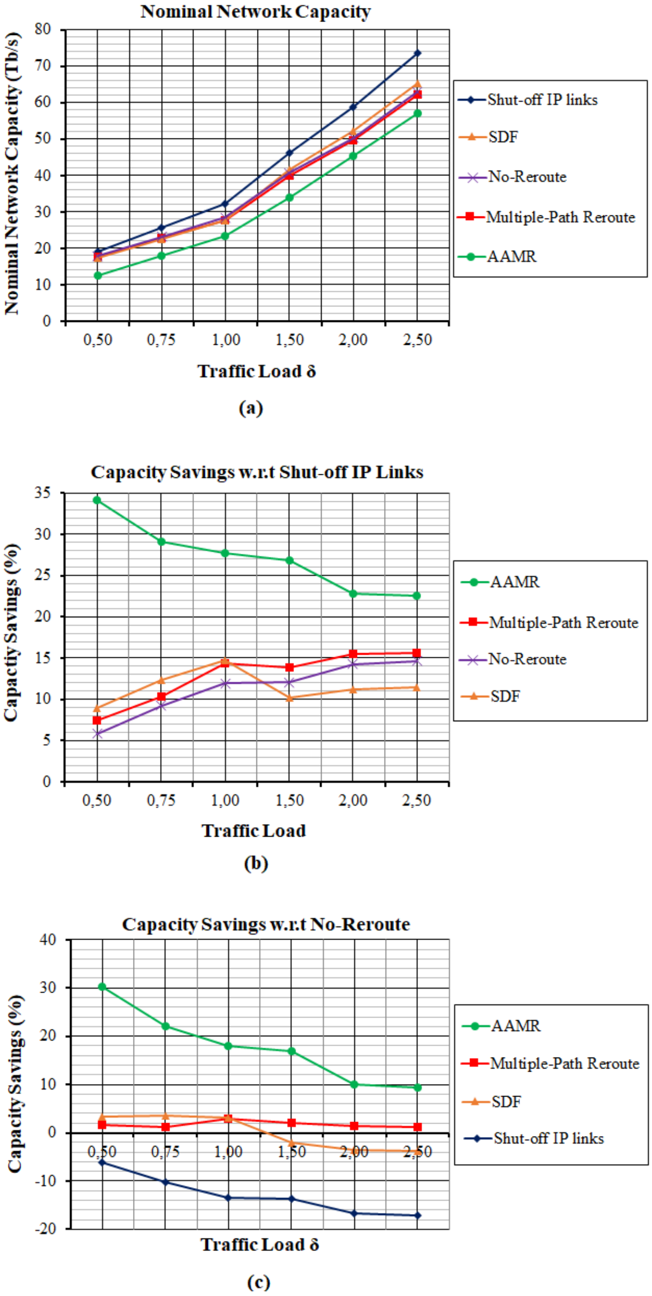

Fig. 4. a) Nominal network capacity; b) Capacity savings w.r.t *Shut-off IP links*; c) Capacity savings w.r.t *No-Reroute*.

For the alternative approaches, a_{ek} is obtained by solving the optimization algorithms in [3], while for the *AAMR* by solving the MIP (1). All problems are solved in CPLEX version 12.9 on an Intel i7-6700 machine at 3.4 GHz with 32 GB RAM.

As seen, across all traffic loads, *Shut-off affected IP links* is the least efficient approach while *AAMR* outperforms all methods. The second best MR strategy is *SDF* for $\delta \leq 1$, and *Multiple-Path Reroute* for $\delta > 1$. As argued in [3], for $\delta \leq 1$, *SDF* outperforms *Multiple-Path Reroute* because most traffic volumes are low enough that demand rerouting is not required in most cases, thereby requiring less spare capacity. As the traffic load grows (i.e. for $\delta > 1$), the traffic volumes grow in a way that *SDF* reroutes most demands carried on the affected IP links, which increases the spare capacity requirements. In this case, the rerouting strategy applied by *Multiple-Path Reroute* is more efficient. Nonetheless, Fig. 4a shows that *AAMR* significantly outperforms this approach and *SDF*.

Figure 4b depicts the capacity savings of the MR strategies w.r.t *Shut-off affected IP links*. The savings are calculated as: $100 \times [C_0 - C_i]/C_0$, with C_0 being the capacity of *Shut-off affected IP links* and C_i the capacity of the ith method (e.g. SDF). For $\delta \leq 1$, the savings attained by *AAMR* vary between 27% and 34%, which are significantly higher than those of *SDF* (the second best strategy) which vary between 9% and 15%. For $\delta > 1$, the savings are at least 22% for *AAMR*, which outperforms the maximum savings of the second best strategy, i.e. *Multiple-Path Reroute,* which in the best scenario ($\delta = 2.5$) achieves 16% savings. These results show the poor efficiency – in terms of capacity requirements – of applying pure IP/MPLS-based restoration, which is the case of *Shut-off affected IP links.* An alternative method, given its operational simplicity, is *No-Reroute*, which avoids routing at the expense of capacity overprovisioning. The savings incurred by the MR strategies w.r.t this method are shown in Fig. 4c. As seen, the capacity requirements of *No-Reroute* are comparable to those of *Multiple-Path Reroute* and *SDF*. For $\delta \leq 1$, the savings of these methods vary between 1.3% and 3.5%. For $\delta > 1$, *No-Reroute* outperforms *SDF*, while *Multiple-Path Reroute* is slightly better with savings around 1% for high traffic loads. (The negative savings in Fig. 4c refer to methods that need more capacity than *No-Reroute*.) On the other hand, *AAMR* significantly outperforms *No-Reroute* across all traffic loads, with savings that vary between 9% and 30%. This evinces the benefits of applying capacity thinning and rerouting of MPLS tunnels for multilayer restoration.

Robustness by the definition of decision rules such as (1d) may result in capacity requirements higher than those of the non-robust counterpart of the *AAMR*. To assess this, we have calculated the nominal capacity for the non-robust version of (1), which is formulated by replacing the affine constraint (1d) by: $c^{s}_{dp} < c^{s_0}_{dp}, d \in D, p \in P(d), s \in S$. The results show that the non-robust MIP yields the same capacity requirements as those of the *AAMR* in Fig. 4a. This statement, however, cannot be generalized and must be further validated for other network scenarios. Moreover, although the results prove the capacity-efficiency of the *AAMR,* two issues must be studied as future work. First, the *AAMR* guarantees enough capacity to operate in all states in S. For unexpected failures, the method adapts (through the affine rule) to carry the affected demands at its best. The performance of this mechanism must be assessed for unexpected failure scenarios. This includes the evaluation of metrics such as the traffic volume that cannot be served after the rule is applied. Another issue concerns the recovery delay, which is dependent on the control plane protocol applied for MR. In T-SDN networks the development of such protocols is work in progress. Protocols must be designed that minimize the delays of the method while considering the latencies added by the configuration of equipment in the network layers.

VI. CONCLUSION

The results in this paper lay ground work for the design of MR methods for IP/MPLS-Optical networks with partial link failures, i.e. failure states with multiple IP links with reduced capacities. This problem has remained nearly unexplored. We have shown that significant capacity savings can be achieved by applying adjustable robust optimization. By this approach, the MPLS tunnels adjust their routing and capacities to the capacity reductions observed in the optical layer. Besides, the network is endowed with decision rules to adapt to optical link failures not regarded in the optimization process. To further understand the properties of the method, we identify future research work in two directions. First, in the context of T-SDN control, protocols for MR must be designed that minimize the restoration delays of the method. The physical constraints of the network layers must be considered in the design. Secondly, the performance of the robustness mechanisms enabled by adjustable robust optimization must be assessed. This involves the evaluation of performance metrics (e.g. unserved traffic and capacity utilization) for affine and further decision rules in unexpected failure scenarios.

ACKNOWLEDGMENT

We thank Matthias Gunkel and Felix Wissel at *Deutsche Telekom Technik* for the valuable discussions and information provided to assess the case study presented in this paper.

REFERENCES

[1] O. Gerstel, C. Filsfils, T. Telkamp, M. Gunkel, M. Horneffer, V. Lopez, and A. Mayoral, " Multi-Layer capacity planning for IP-Optical networks," in IEEE Communications Magazine, vol. 52, pp. 44-51, January 2014.

[2] İ. Yanıkoğlu, B. Gorissen, and D. den Hertog, "A survey of adjustable robust optimization," European Journal of Operational Research, Vol. 277, pp. 799-813, September 2019.

[3] H. Saeed, R. Romero Reyes, T. Bauschert, M. Gunkel, and F. Wissel, "Spare capacity dimensioning for different re-routing strategies in transport networks," In 11th International Workshop on Resilient Networks Design and Modeling (RNDM), October 2019.

[4] R. Singh, M. Ghobadi, K.T. Foerster, M. Filer, and P. Gill, "RADWAN: rate adaptive wide area network," in Proceedings of the 2018 Conference of the ACM SIGCOMM, August 2018.

[5] J. Vasseur and M. Pickavet, "Network recovery: protection and restoration of optical, SONET-SDH, IP, and MPLS," Morgan Kaufmann, 2004.

[6] M. Pióro and D. Medhi D, "Routing, flow, and capacity design in communication and computer networks," Morgan-Kaufman, San Francisco, 2004.

[7] A. J. Piotr Cholda, "Recovery and its quality in multilayer networks," IEEE/OSA Journal of Lightwave Technology, vol. 28, 2010.

[8] J. Rak, "Resilient routing in communication networks," Springer International Publishing, Switzerland, 2015.

[9] J. Rak and D. Hutchison, "Guide to disaster-resilient communication networks," Computer Communications and Networks, Springer, 2020.

[10] M. Chiesa, A. Kamisiński, J. Rak, G. Rétvári, and S. Schmid, "Fast recovery mechanisms in the data plane," TechRxiv., May 2020.

[11] M. Pióro, Y. Fouquet, D. Nace, and M. Poss, "Optimizing flow thinning protection in multicommodity networks with variable link capacity," Operations Research Journal, Vol. 64, April 2016.

[12] I. Kalesnikau, M. Pióro, M. Poss, D. Nace, and A. Tomaszewski, "A robust optimization model for affine/quadratic flow thinning: A traffic protection mechanism for networks with variable link capacity," Networks Journal, Wiley, pp.420-437, February 2020.

[13] M. Gunkel, F. Wissel, W. Weiershausen, M. Franzke, V. Fuerst, and A. Mattheus, "Multi-layer interworking with rate-adaptive transmission technology - benefit and challenges of a new use case," In Proc. of ITG Symposion Photonic Networks, Leipzig, May 2015.

[14] M. Gunkel, F. Wissel, J. Blendin, D. Herrmann, M. Wichtlhuber, and D. Hausheer, "Efficient partial recovery of flexible-rate transceivers with SDN-based asymmetric multipath routing of IP traffic," ITG-Symposium Photonic Networks, Leipzig, May 2016.

[15] J. Blendin, D. Herrmann, M. Wichtlhuber, M. Gunkel, F.Wissel, and D. Hausheer, "Enabling efficient multi-layer repair in elastic optical networks by gradually superimposing SDN," In 12th Int. Conference on Network & Service Management, Montreal, November 2016.

[16] R. Alvizu, G. Maier, N. Kukreja, A. Pattavina, R. Morro, A. Capello, and C. Cavazzoni, "Comprehensive survey on T-SDN: Software-defined networking for transport networks," IEEE Communications Surveys and Tutorials, Vol. 19, pp. 2232-2283, Fourth quarter 2017.

[17] F. Rambach, B. Konrad, L. Dembeck, U. Gebhard, M. Gunkel, M. Quagliotti, L. Serra, and V. Lopez, "A multilayer cost model for metro/core networks," Journal of Optical Communications and Networking, vol. 5, no. 3, March 2013.