

# Power-Napping Networks: Towards Practically Usable Green Segment Routing for ISP Backbones

Daniel Otten<sup>°</sup>, Alexander Brundiers<sup>°•</sup>, Leonhard Brüggemann<sup>°</sup>, Nils Aschenbruck<sup>°</sup>

<sup>°</sup>Osnabrück University, Institute of Computer Science

Friedrich-Janssen-Str. 1, 49076 Osnabrück, Germany

Email: {daotten, brundiers, bruegemann, aschenbruck}@uos.de

<sup>•</sup>Deutsche Telekom Technik GmbH

Wolbecker Str. 268, 48155 Münster, Germany

Email: alexander.brundiers@telekom.de

**Abstract**—Amidst global warming and climate change, reducing the energy consumption of our telecommunication infrastructure becomes increasingly important. A popular approach to do so is switching off underutilized hardware components during low-load phases, supplemented by Traffic Engineering (TE) to free them from the remaining traffic. While there already is a body of work examining the potential of this Green TE concept, most of these approaches focus solely on maximizing the switched-off hardware while not overutilizing the network, thereby not considering other crucially important operational constraints and requirements (e.g., regarding latency). This renders the actual practical usability of these approaches questionable at best. To address this issue, this paper proposes two Segment Routing-based Green TE algorithms for maximizing energy savings while also adhering to a broad set of operational constraints. In extensive evaluations on a variety of networks, including recent data from a Tier-1 Internet Service Provider, we show that our approaches are able to maintain near-optimal power saving levels, while substantially reducing configuration effort, satisfying latency bounds, and adhering to other important constraints as well.

## I. INTRODUCTION

Driven by the need to continuously expand their network infrastructure to keep up with the ever-growing amount of data traffic, the power consumption of Internet Service Providers (ISPs) and other operators is reaching new landmarks every year. This not only results in increased operational costs, but, in times of global warming and climate change, becomes more and more of an ecological concern as well. Telefónica, for example, reported an energy consumption of 41 MWh per petabyte of traffic in 2023 resulting in  $CO_2$  emissions of around 337,119 tons [1].

A popular approach to improve the energy efficiency is *Green Traffic Engineering (TE)*. It aims at freeing up certain network components by steering traffic away from them, allowing them to be switched off. This is particularly promising for reducing energy consumption in off-peak/low-load phases in which large parts of the hardware are only sparsely utilised already (i.e., during the night).

Numerous works [2]–[5] have explored this concept, with some reporting energy savings of more than 50% [4]. However, basically all these approaches focus solely on computing TE configurations that maximize the switched-off hardware and, thereby, the energy savings while maintaining enough capacities to sufficiently route the respective traffic. The latter

is undeniably important and a necessary condition for any useful solution.

In the context of ISP backbones, for example, adhering to certain Quality-of-Service constraints (e.g., regarding the latency of certain traffic flows) is crucially important due to contractually guaranteed Service-Level-Agreements. Furthermore, Green TE solutions should be configurable with a reasonable amount of effort (i.e., as few changes as possible) and overhead to not put too much of a burden on network operation or the network itself. Unfortunately, existing approaches largely ignore these practical constraints (see Sec. III), making their practical usability questionable and potentially overestimating achievable energy savings.

This paper aims to address this research gap by studying the energy saving potential of Segment Routing (SR)-based Green TE approaches for backbone networks, when also adhering to a multitude of different operational constraints and requirements. Thereby, we make the following main contributions:

- We develop an Linear Program (LP)-based optimization algorithm for computing SR configurations that maximize the number of deactivated hardware components (e.g., linecards) while adhering to operational constraints.
- We are the first to apply SR Midpoint Optimization (MO) [6] to Green TE, significantly reducing the number of required SR policies and improving practical usability.
- In an extensive evaluation across various networks, including recent Tier-1 ISP backbone data, we show that our Green SR algorithms achieve energy savings comparable to conventional methods while respecting crucial operational constraints.
- We demonstrate that SR MO reduces the number of required SR policies by up to 99%, improving network maintainability and enabling faster configuration.

The remainder of the paper is structured as follows. First, Sec. II introduces fundamental concepts and background information required for understanding this paper, followed by a discussion of related work in Sec. III. After this, we present our novel Green SR approaches aiming to address the problem of incorporating crucial operational constraints into the energy saving optimization (Sec. IV). Sec. V then introduces our evaluation setup, and the respective evaluation results are presented in Sec. VI. Finally, Sec. VII discusses our main findings and limitations before Sec. VIII concludes this paper




Figure 1: Illustration of the traffic steering differences between MO and conventional E2E SR.

and its contributions, while also briefly discussing possible future research directions.

## II. BACKGROUND

This section introduces fundamental information relevant for the understanding of this paper, mainly focussing on the power consumption and related power-saving approaches in backbone networks, as well as a brief introduction to SR TE.

### A. Network Power Consumption

The power consumption of a backbone network primarily stems from the power usage of its individual components, with routers being the most significant contributors. Thus, steering traffic away from specific routers and turning them off during periods of low utilization might seem like a straightforward approach for reducing energy consumption. Unfortunately, the latter often is not implementable in practice for several reasons. First of all, routers do not only process traffic within the network, but also connect the backbone to smaller customer networks. Shutting down a router would sever such connections, leading to service disruptions or at least considerably reduced connectivity and, thus, fault tolerance. Additionally, current routers are designed for long uptimes and not meant to be switched on and off continuously. Rebooting them often takes around 30 minutes and can even lead to hardware failures.

Consequently, instead of deactivating entire routers, most Green TE approaches focus on alternative strategies aimed at minimizing router power consumption. Thereby, the latter can be attributed to three main components:

- 1) *Chassis and fans*: These provide essential functions such as power supply and cooling.
- 2) *Route Switch Processor (RSP)*: These processors handle the operating system and basic management tasks.
- 3) *Line cards*: These components provide physical endpoints for connections (i.e., ports). According to [8], [9], line cards consume the majority of power within a router and, thus, within the network itself.

Historically, many Green TE approaches, thus, focused on reducing power consumption by deactivating unused ports or

steering traffic away from the respective link to turn it off [10]–[14]. However, recent research [8], [9] shows that turning off only a few ports per line card yields minimal power savings. Meaningful reductions are only achieved when all ports on a line card are deactivated, thereby allowing for the entire card to be powered off. Hence, modern Green TE strategies should focus on switching off entire line cards to maximizing power savings.

### B. Segment Routing Traffic Engineering

Segment Routing (SR) [15] is a rather recent addition to the TE toolbox. It is based on the idea of controlling a packet's path through the network by adding waypoints (so called *segments*) to a packet, that the latter is then steered over in the given order. Contrary to similar concepts (i.e., Multiprotocol Label Switching (MPLS) with RSVP-TE [16]), SR tunnels (or *policies*) only need to be configured on the respective headend node with basically all other required information being carried in the packet itself. This offers considerable benefits regarding both overhead and scalability while still enabling the configuration of virtually arbitrary forwarding paths, which has lead to SR becoming one of if not the preferred technology amongst many operators [17].

However, basically all of the SR literature focuses on E2E SR which deploys SR policies as dedicated E2E tunnels for a single demand. While this offers detailed, per-flow traffic control, the intrinsic need to install a dedicated SR policy for every demand that needs to be detoured often causes a large number of policies being required to implement TE solutions [6], [18], resulting in (unnecessarily) high network complexity and overhead. To address this, recent works [6], [18] started to combine SR with more flexible traffic steering approaches (i.e., *IGP Shortcut*) that allow to detour multiple demands with just a single policy. It has been shown in different contexts that this facilitates a substantial reduction in the number of required SR policies (i.e., by up to 99%), thereby greatly improving the practical usability of SR. Since, later on in this paper, we also utilize this so called *Midpoint Optimization (MO)* concept to reduce the configuration complexity and overhead of our proposed GreenSR approach, Figure 1 further illustrates the idea behind the latter and the respective traffic steering

differences compared to E2E SR. For more details on the general MO concept, we refer to [6], [19].

### III. RELATED WORK

Reducing the power consumption of telecommunication networks has been an important objective for quite some time, with its societal relevance even further increasing in the face of climate change and global warming. As a result, there is a wide variety of publications [2]–[5], [10]–[14], [20]–[26] in the general area of Green TE. In the following, we focus on assessing these approaches regarding their actual practical usability. Thereby, we focus on three main aspects:

**1) The switched-off hardware:** Basically all Green TE approaches aim to reduce energy consumption by unburdening and switching off certain hardware components. However, there are considerable differences regarding both operational feasibility and actual energy saving potential when it comes to the type of the hardware that should be switched off. As already explained in Section II-A, shutting down whole routers is generally not really feasible/desirable from an operational perspective, and switching off individual ports – while at least practically feasible – only yields minimal power savings. Thus, the actual sweetspot for achieving relevant reductions in power consumption that are also implementable in practice, is to aim for disabling entire linecards.

**2) TE technology:** When it comes to steering traffic away from certain hardware components in order to switch them off, there are different TE technologies that can be used for this. Especially older approaches often still rely on Interior Gateway Protocol (IGP) metric tuning. While this is a long-standing and well understood TE concept, there are various limitations (e.g., its inability to precisely control the paths of individual demands). Thus, metric tuning is replaced by newer and more sophisticated TE technologies. One of those, which is frequently used in the Green TE context as well, is MPLS with Resource Reservation Protocol (RSVP)-TE. It offers immaculate traffic steering capabilities and is deployed in many large backbone networks across the globe. However, with networks growing more and more in size, its rather limited scalability becomes more and more of an issue. Thus, it is (again) replaced by recent technologies like SR. Since the vast majority of operators either already deploys SR or plans to do so soon [17], we argue that, to facilitate best compatibility, modern Green TE approaches should also be SR-based instead of relying on older technologies that are either already replaced or will be in the foreseeable future.

**3) Operational constraints & requirements:** Another aspect often overlooked in the Green TE literature are technical constraints and operational requirements that the respective TE solution has to adhere to in order to be actually useable in practice. While many of these are rather operator-specific and depend on the respective use case, there is a certain set of constraints that must be generally adhered to as they constitute technical necessities. In the context of SR, this includes the maximum number of segments that can be applied to a packet. From a theoretical perspective, the latter is basically unlimited,

allowing for virtually arbitrary forwarding paths to be implemented. In practice, however, it is limited by the *Maximum Segment Depth (MSD)* of the used hardware, which can vary substantially between different vendors and models. Hence, in order to ensure compatibility with a wide share of hard- and software, the number of segments per policy should be kept as low as possible. Another aspect to consider for both SR as well as MPLS is the fact that current routers only support splitting traffic flows equally into predefined fractions, which has to be incorporated in the respective optimization models. Furthermore, it has to be ensured that the respective TE configuration can be implemented with a reasonable number of changes (i.e., MPLS tunnels or SR policies), in order to keep the configuration efforts manageable and to reduce the resulting overhead (especially when using MPLS). Finally, a service-related requirement that is of utmost importance in most carrier-grade networks (i.e., ISP backbones) is the compliance with latency bounds for specific demands, resulting from contractual QoS agreements with customers (so called Service-Level-Agreements).

For reasons of space, we cannot discuss every existing Green TE approach and its compliance (or non-compliance) with the above criteria in detail. Instead, we provide a visual summary of the latter in Table I. It can be seen that very few approaches actually focus on switching off linecards, with the vast majority still aiming for either whole routers or individual ports, resulting in questionable practical applicability and suboptimal effectiveness, respectively. There also are completely different approaches, like the one presented in [2]. Here, the authors aim to direct traffic towards locations where green energy is available. While this is a novel and interesting strategy when it comes to reducing the ecological footprint of a network, we suspect that its economic impact (i.e., regarding the energy cost) is likely rather negligible since, in the European Union, all countries share a common energy market with electricity prices being determined by the most expensive generator required to meet demand. When it comes to the used TE technology, most existing approaches rely on rather outdated concepts that are already no longer in use or will be replaced in the foreseeable future, further limiting their practical usability. Finally, virtually none of the existing approaches incorporate any technical or operational constraints into their optimization models. While some of them discuss the *possibility* of including QoS constraints (e.g., [2]), the fact that they do not feature any of these in their evaluations, renders the respective results purely theoretical.

All in all, this shows that the actual practical applicability of most existing Green TE approaches is at least questionable, as they either focus on outdated TE technologies, unsuitable target functions, or do not include crucial operational constraints. In this paper, we aim to address this issue by proposing two SR-based Green TE approaches that minimize the number of active linecards in a network while also adhering to all of the above mentioned operational constraints and requirements.

Table I: Overview and comparison of existing work in the area of GreenTE regarding the hardware components that are switched off, the used TE technology, and whether they consider and adhere to important real-world constraints and requirements.

|                                 | Switched-off Hardware | TE Technology | SR MSD Limitations | Realistic Traffic Splitting | Reduction of Configuration Effort | Latency Constraints |
|---------------------------------|-----------------------|---------------|--------------------|-----------------------------|-----------------------------------|---------------------|
| [2]                             | Other                 | Metric Tuning | –                  | –                           | –                                 | ✗                   |
| [13], [20], [21]                | Ports                 | Metric Tuning | –                  | –                           | –                                 | ✗                   |
| [22], [23]                      | Router                | MPLS          | –                  | –                           | ✗                                 | ✗                   |
| [3], [10]–[12], [14], [24]–[26] | Ports                 | MPLS          | –                  | ✗                           | ✗                                 | ✗                   |
| [4], [5]                        | Linecards             | SR            | ✓                  | ✗                           | ✗                                 | ✗                   |
| <i>This Work</i>                | Linecards             | SR            | ✓                  | ✓                           | ✓                                 | ✓                   |

$$\min \sum_{v \in V} \left\lceil \frac{\sum_{p \in v} \pi_p}{k} \right\rceil \quad (1)$$

$$\text{s.t.} \quad \sum_{w \in V \setminus \{u\}} x_{uw}^w = 1 \quad \forall (u, v) \in V^2 \quad (2)$$

$$x_{u,v}^w \cdot (\text{del}(u, w) + \text{del}(w, v)) \leq \text{lb}(u, v) \quad \forall (u, v) \in V^2 \quad (3)$$

$$\sum_{(u,v) \in V^2} \sum_{w \in V \setminus \{u\}} g_{uv}^w(a) x_{uw}^w \leq \theta \sum_{p \in P(a)} \pi_p c_p \quad \forall a \in A \quad (4)$$

$$\sum_{p \in P(a_{uv})} \pi_p = \sum_{p \in P(a_{vu})} \pi_p \quad \forall a_{uv} \in A \quad (5)$$

$$x_{uw}^w \in \{0, 1\} \quad \forall (u, v) \in V^2 \quad (6)$$

$$\pi_p \in \{0, 1\} \quad \forall p \in \bigcup_{a \in A} P(a) \quad (7)$$

Problem 1: Green-Segment Routing (GSR) ILP formulation.

#### IV. ALGORITHMS

In this section, we propose two novel Green SR optimization models which are the first to incorporate all of the previously discussed aspects required for facilitating realistic and practically usable results (cf. Sec. III). Thereby, the first approach solely relies on conventional E2E SR, while the second one utilizes SR MO [6] to further reduce configuration efforts.

##### A. Fundamental Graph-based Network Model

We model the network as a directed graph  $G = (V, A)$ . The set of vertices  $V$  represents active routers in the network, and the set of arcs  $A$  represents its links. Note that we use a multigraph instead of a simple graph to allow for parallel links. As routers establish full duplex connections, it is necessary to model each link as two directed arcs, i.e., one link between router  $u$  and router  $v$  is represented by a set of two distinct arcs  $\{a_{uv}, a_{vu}\}$ . One such pair of arcs is made up of multiple ports on the endpoints, with the capacity of this connection being determined by the number and types of the respective ports. Let  $P(a)$  denote the set of ports establishing a connection, represented by arc  $a$ . It holds that

$$P(a_{uv}) = P(a_{vu})$$

as one port builds a full duplex connection. The capacity of a port is denoted as  $c_p$ . Thus, the capacity of an arc is given by

$$c(a_{u,v}) = c(a_{v,u}) = \sum_{p \in P(a_{u,v})} c_p.$$

While both directions have the same capacity due to them being established by the same ports, their utilization can differ since the routed traffic volume is generally not symmetric.

The amount of traffic that has to be routed from node  $i$  to  $j$  is denoted as  $t_{ij}$ . For every such traffic demand, there also is a *latency bound*  $\text{lb}(u, v)$  specifying the maximum acceptable E2E delay for the respective demand. In a similar context,  $\text{del}(u, v)$  denotes the Shortest Path Routing (SPR) delay between nodes  $u$  and  $v$ .

##### B. Green Segment Routing Optimization Model (GSR)

With our fundamental network model established, we can now describe our GSR approach for deducing a reduced network topology and a corresponding SR configuration that steers traffic away from the switched off hardware to allow for disruption-free deployment. This is achieved by solving a corresponding Integer Linear Program (ILP) (see Problem 1).

The latter relies on two types of binary decision variables. The  $x_{u,v}^w$  define the two segment routing (2-SR) policy<sup>1</sup> chosen

<sup>1</sup>We solely rely on SR policies with two segments to ensure highest compatibility with varying MSD limits in current hardware (cf. Sec. III).

for routing the demand from  $u$  to  $v$ , with  $w$  specifying the respective intermediate target. In this context, setting  $w = v$  resembles SPR. Furthermore, the  $\pi_p$  variables are used to indicate whether a port  $p$  is active or not. The latter information is required as a linecard can only be switched off if all of its  $k$  ports are deactivated. Thus, the objective function (Eq. 1) aims for minimizing the number of active linecards in the network by switching off groups of  $k$  ports per router.<sup>2</sup>

Equation 2 ensures that all traffic demands are satisfied by requiring an SR policy (or SPR) to be selected for each demand. The latter also prevents any kind of unrealistic traffic splitting by allowing at most one SR policy per demand. In order to guarantee the compliance of the solution with the specified latency bounds, Equation 3 eliminates routing policies that violate any of the latter. Equation 4 ensures that no link in the reduced network will be overloaded. Here, the function  $g_{uv}^w(a)$  calculates the amount of traffic from  $u$  to  $v$  routed through arc  $a$ . The left-hand side of the equation represents the total traffic traversing arc  $a$ , while the right-hand side specifies the link capacity. To avoid full utilization of links, an upper bound (commonly around 70% of the total capacity) can be imposed using the  $\theta$  parameter. Finally, the last equation guarantees that a link is fully deactivated only when both of its endpoints (i.e., ports) are turned off.

The ILP as presented in Problem 1 covers all the relevant real-world considerations discussed in Section III, apart from the configuration effort. As a result, solutions can feature a basically arbitrary number of 2-SR policies, making them hard or even impossible to configure. To address this, we employ a concept similar to that proposed in [27]. The core idea is to carry out a second optimization step that minimizes the number of policies required to implement the respective TE configuration. For this, the target function of the ILP is modified to prioritize minimizing the number of active policies instead by replacing it with

$$\min \sum_{w \in V \setminus \{u, v\}} x_{uv}^w.$$

### C. Green-MO: An Energy Optimization Model for SR MO

Like all other Green SR approaches, our GSR algorithm relies on conventional E2E SR. While the latter offers exceptional traffic control, it often results in rather high policy numbers, even when specifically minimizing the latter. In other contexts, utilizing the MO concept [6] has proven to be very effective in overcoming these issues (cf. Section II) by combining the SR paradigm with more flexible traffic steering mechanisms like IGP Shortcut [7]. With the respective TE features starting to become available in the most recent router hard- and software, widespread availability only seems to be a matter of time. Hence, we deem it worthwhile explore the potential benefits of SR MO in the context of Green TE.

For this, we present *Green Midpoint Optimization (G-MO)*, a novel Green SR optimization algorithm for the use with

<sup>2</sup>While we aim for linecard minimization here, this generic model can be adapted to other hardware components as well by adjusting the  $k$  parameter.

IGP Shortcut-based SR MO instead of conventional E2E SR. In simple terms, the latter is derived from adapting our previous E2E SR approach (Problem 1) with the general SR MO optimization model presented in [19]. The resulting ILP formulation is given in Problem 2. As it is conceptually similar to our E2E SR model, we will mainly focus on describing the respective differences and MO-related adaptations of the latter.

The biggest conceptual difference lies in the meaning of the  $x_{u,v}^w$  variables. For E2E SR, those were bound to a specific demand (i.e.,  $u \rightarrow v$ ) indicating which SR policy is installed for the latter. Now, they simply indicate whether an SR policy is installed between nodes  $u$  and  $v$  using intermediate node  $w$ , with this policy not being directly tied to any traffic demand at all. The latter means that multiple demands can be routed via a single SR policy and also that a single demand can pass through multiple policies on the way towards its destination. While this generally offers greater TE flexibility [19], it also induces so called *policy dependencies* [6]. SR policies can now influence and alter the traffic that passes through other SR policies in the network. This is no issue from an operational perspective, but it introduces major optimization-related challenges as the link utilizations resulting from the configuration of a policy cannot be efficiently precomputed anymore, rendering an efficient calculation of a corresponding LP formulation virtually impossible [6], [19]. To overcome this issue, we employ a technique proposed in [6] to prohibit the simultaneous configuration of policies that “influence each other” resulting in the addition of constraints 9 and 10. The former introduces a set of helper variables  $y_{u,v}$  that aggregate over all policies between a pair of nodes  $(u, v)$ , and the latter then uses those to prevent the configuration of any of the *influencing policies*  $\mathcal{I}_{u,v}$  for an already selected SR policy between nodes  $u$  and  $v$ . The remainder of the ILP formulation is conceptually identical to Problem 1, with Equations 11 and 12 constituting the latency and capacity constraints, respectively.

Lastly, to minimize the number of SR policies required to implement the respective TE solution, we utilize the same two-phased concept already employed for our GSR model.

## V. EVALUATION-SETUP

In this section, we provide an overview on our evaluation setup, introducing the used datasets, as well as the reference algorithms and respective parametrizations.

### A. Datasets

Our first dataset comprises of network traces collected in the backbone network of a globally operating Tier-1 ISP at 12 days during 2022 (one day per month). For each day, those contain the receptive network topology as well as a “worst-case” traffic matrix of the respective low-load period (i.e., between 01:00 and 09:00 [4]) of the respective day, for which we aim to switch off as many linecards as possible. The respective matrix

$$\min \sum_{v \in V} \left\lceil \frac{\sum_{p \in v} \pi_p}{k} \right\rceil \quad (8)$$

$$\text{s.t.} \quad \sum_{w \in V} x_{u,v}^w = y_{u,v} \quad \forall (u, v) \in V^2 \quad (9)$$

$$y_{u,v} + y_{i,j} \leq 1 \quad \forall (u, v) \in V^2, \forall (i, j) \in \mathcal{I}_{u,v} \quad (10)$$

$$\text{del}(i, j) + x_{uv}^w \cdot (\text{del}(u, w) + \text{del}(w, v) - \text{del}(u, v)) \leq \text{lb}(i, j) \quad \forall (i, j) \in V^2 \quad (11)$$

$$\text{spr}(a) + \sum_{(i,j) \in V^2} t_{i,j} \sum_{u,v,w \in V^3} \text{diff}_{i,j}^{u,v,w}(a) x_{u,v}^w \leq \theta \sum_{p \in P(a)} \pi_p c_p \quad \forall a \in A \quad (12)$$

$$\sum_{p \in P(a_{uv})} \pi_p = \sum_{p \in P(a_{vu})} \pi_p \quad \forall a_{uv} \in A \quad (13)$$

$$x_{u,v}^w \in \{0, 1\} \quad \forall (u, v) \in V^2 \quad (14)$$

$$y_{u,v} \in \{0, 1\} \quad \forall (u, v) \in V^2 \quad (15)$$

$$\pi_p \in \{0, 1\} \quad \forall p \in \bigcup_{a \in A} P(a) \quad (16)$$

## Problem 2: Green Midpoint Optimization (G-MO) ILP formulation.

is calculated by simply using the highest value observed for each demand during the respective time period:

$$t_{i,j}^* = \max_{(01:00, 09:00)} t_{i,j}.$$

The idea behind using such a matrix is to ensure the computed Green SR solution being capable of handling traffic throughout the whole period without causing overutilization [4].

The second dataset features problem instances from the publicly available Repetita dataset [28]. The latter contains a broad collection of topologies from the Internet Topology Zoo [29] for which it also provides artificial traffic matrices resembling high-utilization scenarios. Since many of the topologies in the Repetita dataset originate from rather old and small networks (e.g., the early ARPANET) which do not really reflect modern backbone architectures anymore, we limit our examinations only to those topologies that have at least 60 nodes, leaving us with a total of 30 problem instances. Furthermore, since Green TE generally focuses mostly on switching off hardware during low-load phases, the original, high-utilization traffic matrices provided by Repetita are not well suited for our evaluations. Thus, we scale them down to by 50% to better reflect low-load traffic [5]. Finally, the Repetita dataset also does not include any information on the respective hardware components. Hence, we assume all routers to be equipped with 8-port linecards (i.e., the commonly used A99-8X100GE-TR).

### B. Algorithms and Hyperparameters

In order to assess the energy-saving capabilities of our approaches, we compare them to 2SR-LC [5], a recent state-of-the-art Green SR approach which, however, does *not* adhere to most of the relevant real-world constraints (cf. Table I). For reference, we also examine the energy savings achievable without utilizing any explicit TE to steer traffic away from hardware, but by simply switching of ports until the link

utilization surpasses a predetermined threshold  $\theta$ . The latter is referred to as SPR in the following evaluation.

To specifically investigate the influence of the latency bound constraints on the energy saving potential, we examine two different variations of our optimization models. The default one, which includes *all* of the relevant constraints as described in the respective problem formulation (cf. Section IV), and a version *not* enforcing latency bounds, indicated by the “-ND” (i.e., “No Delay”) extension to its name.

Finally, all algorithms are configured to not surpass an Maximum Link Utilization (MLU) threshold of  $\theta = 70\%$  to leave sufficient network capacity to handle unexpected scenarios like traffic spikes or failures [5].


## VI. EVALUATION RESULTS

This section presents and discusses the results of our extensive evaluations of the GSR and G-MO approaches. Thereby, we, first and foremost, focus on the achievable power savings of each approach. Additionally, we also examine the compliance with relevant real-world constraints like latency bounds or the number of SR policies required to implement the respective TE configuration.


All computations are executed on a system equipped with two AMD EPYC 7452 CPUs, 512GB of RAM, under the 64-bit Ubuntu 20.04.1 operating system. LPs are solved using CPLEX [30] with a maximum timelimit of 24h.

### A. Tier-1 ISP Dataset

**Energy Saving Potential:** The first, most intuitive performance indicator to look at when examining Green TE approaches is the achievable energy savings. In this context, reporting actual power values proves to be difficult since those are heavily dependent on the respective hardware, and other network factors. Hence, research often focuses on the



(a) Percentage of switched off linecards.



(b) SR policies required to implement the respective TE configuration.

Figure 2: Evaluation results obtained on the ISP dataset.




Figure 3: Distribution of the number of violated latency bounds as well as the severity of their violation.

switched off hardware components (i.e., linecards) as a substitute measure instead [4], [5]. The respective results of our examinations on the ISP dataset are depicted in Figure 2a, showing the percentage of linecards that can be deactivated using different Green TE approaches. The first thing to notice is all TE-based approaches substantially outperforming the Green SPR baseline that only tries to switch off unused hardware without utilizing TE to actively steer traffic away from it. This rudimentary approach only manages to switch off less than 50% of the network’s linecards, while the TE-based approaches all achieve savings of around 75–80%. This impressively highlights the benefits of employing dedicated TE strategies to facilitate improved energy savings in large backbone networks.


Much more important, however, is the general observation that both of our novel Green SR models are able to keep up with the state-of-the-art (i.e., 2SR-LC [5]) in terms of power-saving capabilities. In addition, they integrate a variety of crucial operational constraints and requirements not covered by the latter (cf. Table I), thereby producing far more realistic results. In detail, it can be seen that, without the latency bound constraint, our GSR model basically perfectly matches 2SR-LC in terms of energy savings. When activating the latter, results only become marginally worse, with the deterioration

being basically negligible in practice (i.e., mostly just a single digit number of linecards in a network with over 500 linecards). A similar result can also be observed for our G-MO model. Here, it also becomes apparent that switching from E2E SR to MO to further reduce configuration efforts (cf. Sec. IV-C) comes at the price of small deteriorations in the achievable energy savings due to a further restricted solution space.


**SR Policy Numbers:** However, when looking at the policy numbers required to implement the respective solutions (see Figure 2b), the benefits of using SR MO become abundantly apparent. 2SR-LC requires multiple thousands of SR policies as it does not care about the latter but solely focuses on reducing power consumption. Such a high number of policies renders the solutions basically unusable, as they would take a considerable amount of time to configure, introduce non-negligible overhead, and also severely impact the clarity and maintainability of a network [6]. In contrast, our new GSR model allows to substantially reduce these numbers by employing a second optimization step to actively minimize the number of required policies (cf. Sec. IV-B), often achieving values in the low hundreds range.<sup>3</sup> While this already constitutes a considerably improvement over the current state-of-the-art, the policy numbers remain rather high, thereby still limiting the practical usability of the solutions. In this context, switching from conventional E2E SR to SR MO proves to be highly effective for further reducing policy numbers. Both G-MO variations consistently deliver solutions requiring fewer than 100 policies, with the integration of latency constraints even further reducing the latter. On average, the G-MO

<sup>3</sup>It has to be noted that the policy minimization step is very demanding in terms of computation times and often surpasses our 24h timelimit causing the optimization to be aborted prematurely. For 6 out of the 12 instances, we still obtained solutions that were reasonably close to the optimum (i.e., with a gap of less than 10% between the current solution and the lower bound). In these cases, the number of policies was reduced to just 200–300. For the other six instances, CPLEX failed to achieve similar solution quality, aborting optimization with a more than 60% gap, resulting in multiple thousands of policies still being required. The reported GSR policy numbers could, thus, be further reduced by allowing higher optimization times.





(a) Percentage of switched-off linecards.



(b) SR policies required to implement the respective TE configuration.

Figure 4: Evaluation results obtained on the Repetita dataset.

solutions require just 28 SR policies to be implemented, constituting a more than 99% reduction compared to 2SR-LC, the current state-of-the-art, while only causing a mostly negligible deterioration in the achievable energy savings.

**Latency Bounds:** Finally, to demonstrate the importance of actively including latency bound constraints into Green TE optimizations, we take a dedicated look at the number of bound violations resulting from *not* actively enforcing them during computation (i.e., 2SR-LC) as well as the severity of the resulting violations. The respective results are depicted in Figure 3. It can be seen that omitting such constraints leads to bound violations for around 20% or all demands, of which around half are of substantial severity with the respective bound being exceeded several times over. To put this into perspective, many business customers require connections with very low latency (i.e., a couple milliseconds) being guaranteed at any time during the day. In this context, exceeding such bounds by even a small margin can already induce considerable negative impact. Violations of up to factor 15, however, as they are observed in Figure 3, can result in the respective connection becoming basically unusable for the customer’s purposes. Thus, from an operational perspective, this constitutes a considerable limitation regarding the practical usability of the respective Green TE solutions in large carrier-grade networks, up to a point where they are basically not usable at all. Hence, integrating the respective constraints into Green TE optimization procedures is crucially important, and we show that it can be done with only negligible impact on the overall energy saving potential using our proposed optimization models (cf. Figure 2a).

#### B. Repetita Dataset Results

The results of an analogous evaluation carried out on the different problem instances in our Repetita dataset are depicted in Figure 4.<sup>4</sup> Overall, the results are qualitatively similar to what we have already seen for the ISP dataset. In general all Green TE approaches substantially outperform the SPR baseline, allowing nearly triple the linecards to be switched off. Furthermore, it can be seen again that our newly

proposed Green SR approaches are able to achieve energy savings on-par with 2SR-LC while integrating crucial real-world constraints and also reducing the number of SR policies required to implement the respective solutions by a substantial margin. Regarding the latter, the MO based G-MO model, again, proves to be exceptionally effective, achieving basically optimal power savings with only less than 1% of the policies required by 2SR-LC.

Overall, these results show that it is possible to integrate a variety of critical real-world constraints (e.g., realistic traffic splitting, MSD limitations, latency bounds, and policy number minimization) with basically negligible degradation of the energy saving potential compared to state-of-the-art approaches not considering the latter. This constitutes a significant step towards bridging the gap from theoretical observations towards actually practically usable Green TE.

## VII. DISCUSSION

Since our approaches are based on ILP, they are rather resource-intensive regarding both computation time as well as memory demand. For larger instances, computation of a suitable Green SR solution takes several hours and require several hundred gigabytes of RAM. However, with our work aiming at identifying a singular reduced night-time topology that is used for weeks or even months (i.e., until the network topology changes), investing such efforts is perfectly acceptable. Nevertheless, there might be use cases that aim for a more dynamic adaption of the network topology to facilitate energy savings even outside the low-load phases. For this, faster optimization approaches are needed (i.e., based on heuristic optimization). We plan to look into this in the future.

Furthermore, as already mentioned in Section III, there is a wide variety of (operational) constraints to consider when it comes to actually deploying Green TE in practice, many of those being highly use case dependent. Examining all of these is out of the scope of this paper. Hence, we focused on a subset of constraints that should be (more or less) universally applicable to most backbone networks. However, while we have shown that those can be integrated without notably deteriorating the achievable energy savings, there are further constraints to consider when it comes to deploying Green SR

<sup>4</sup>For the Repetita dataset, latency bound information is not available. Thus, we had to omit this aspect during our examinations.



solutions in practice. One of those aspects is the “port-to-linecard mapping” used in our models. We assume that, as soon as we are able to switch off  $k = 8$  ports of a router, a linecard can be powered off. This can, in theory, be achieved by adapting the port assignments on a router. In practice, however, the port allocation also has to adhere to other design principles, typically prioritizing reliability and fault tolerance. For example, network operators often distribute the ports for parallel links across different linecards to ensure that no single linecard failure can disrupt an entire connection between two nodes. As a result, the port-to-linecard mapping assumed by us is most-likely rather optimistic, thereby potentially overestimating the energy savings actually achievable when also adhering to realistic port assignment rules. This represents the primary barrier that currently prevents actual deployment of our computed solutions and requires further investigation once realistic port assignment data becomes available. We plan to further investigate this in the future, but this requires information on the actual port assignments of the examined networks.

### VIII. CONCLUSION

In the times of global warming, reducing the energy consumption of our telecommunication infrastructure and especially backbone networks becomes increasingly important, with Green TE constituting a promising approach to achieve this. However, most of the respective approaches proposed in the literature focus solely on the energy-saving aspect while completely ignoring various crucial operational constraints. This limits practical usability and may overestimate achievable energy savings. In this paper, we addressed this research gap by proposing two novel Green SR approaches aiming for reducing a networks energy consumption by switching off as many linecards as possible, while also adhering to a variety of important real-world constraints (e.g., regarding latency, configuration effort, or certain hardware limitations). In an extensive evaluation featuring real-world data from a Tier-1 ISP and a large set of instances from the publicly available Repetita [28] dataset, we are able to show that, despite being the first to enforce multiple crucial operational constraints, our approaches are able to match the energy-saving capabilities of the current state-of-the-art. Furthermore, they also reduce the configuration effort required to implement the respective Green TE solution (i.e., the number of required SR policies) by up to 99%, which constitutes a considerable advantage over conventional approaches regarding network clarity, maintainability, and overhead. All in all, the findings and contributions of this work constitute a big step forward on the way towards actually practically usable Green SR, thereby providing a solid foundation for further research in this area. The latter could, for example, look further into energy-aware port allocations or more dynamic reconfiguration strategies.

### IX. ACKNOWLEDGMENT

This work was supported in part by the German Research Foundation (DFG), Project No. 461207633.

### REFERENCES

- [1] Telefonica, “Management and Sustainability Report 2023,” accessed 11.02.2025. [Online]. Available: <https://www.telefonica.com/en/wp-content/uploads/sites/5/2024/03/management-sustainability-report-2023.pdf>
- [2] S. El-Zahr, P. Gunning, and N. Zilberman, “Exploring the benefits of carbon-aware routing,” in *Proc. of the 2023 CoNEXT*, 2023.
- [3] R. Carpa, O. Glück, and L. Lefevre, “Segment Routing based Traffic Engineering for Energy Efficient Backbone Networks,” in *Proc. of ANTS*, 2014.
- [4] D. Otten, A. Brundiers, T. Schüller, and N. Aschenbruck, “Green segment routing for improved sustainability of backbone networks,” in *Proc. of the LCN*, 2023.
- [5] D. Otten, M. Ilse, M. Chimani, and N. Aschenbruck, “Green traffic engineering by line card minimization,” in *Proc. of LCN*, 2023.
- [6] A. Brundiers, T. Schüller, and N. Aschenbruck, “Midpoint optimization for segment routing,” in *Proc. of the 2022 INFOCOM*, 2022.
- [7] J. Shen and H. Smit, “Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels,” RFC Editor, RFC 3906, 2004.
- [8] D. Otten, S. Neuner, and N. Aschenbruck, “Modeling the Power Consumption of a backbone Network,” in *Proc. of the ICC Workshops*, 2023.
- [9] R. Jacob, J. Lim, and L. Vanbever, “Does rate adaptation at daily timescales make sense?” in *Proc. of the 2nd Workshop on Sustainable Computer Systems*, 2023.
- [10] A. B. Vieira, W. N. Paraizo, L. J. Chaves, L. H. A. Correia, and E. F. Silva, “An SDN-based energy-aware traffic management mechanism,” *Annales des Télécommunications*, 2022.
- [11] K. S. Ghuman and A. Nayak, “Per-packet based energy aware segment routing approach for Data Center Networks with SDN,” in *Proc. of ICT*, 2017.
- [12] X. Jia, Y. Jiang, Z. Guo, G. Shen, and L. Wang, “Intelligent path control for energy-saving in hybrid SDN networks,” *Computer Networks*, 2018.
- [13] O. Okonor, S. Georgoulas, and M. Abdullahi, “Disruption/Time-Aware Green Traffic Engineering for Core ISP Networks,” in *Proc. of ICCCN*, 2016.
- [14] O. Okonor *et al.*, “Dynamic Link Sleeping Reconfigurations for Green Traffic Engineering,” *Int. J. Commun. Syst.*, 2017.
- [15] C. Filsfils *et al.*, “Segment Routing Architecture,” RFC Editor, RFC 8402, 2018.
- [16] D. O. Awduche, L. Berger, D. Gan, T. Li, V. Srinivasan, and G. Swallow, “RSVP-TE: Extensions to RSVP for LSP Tunnels,” RFC 3209, 2001.
- [17] R. Mota, “Segment Routing Survey,” ACG Research, White Paper, 2022.
- [18] A. Brundiers, T. Schüller, and N. Aschenbruck, “Fast Reoptimization with only a few Changes: Enhancing Tactical Traffic Engineering with Segment Routing Midpoint Optimization,” *IEEE JSAC*, 2025.
- [19] A. Brundiers, T. Schüller, and N. Aschenbruck, “An Extended Look at Midpoint Optimization for Segment Routing,” *IEEE OJComS*, 2024.
- [20] A. Cianfrani *et al.*, “An OSPF enhancement for energy saving in IP networks,” in *Proc. of INFOCOM WKSHPs*, 2011.
- [21] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, “An energy saving routing algorithm for a green ospf protocol,” in *Proc. of INFOCOM WKSHPs*.
- [22] G. C. Sankaran and K. M. Sivalingam, “Load-dependent power-efficient passive optical network architectures,” *JoCN*, 2014.
- [23] L. Chiaraviglio, M. Mellia, and F. Neri, “Reducing Power Consumption in Backbone Networks,” in *Proc. of ICC*, 2009.
- [24] N. Vasić, P. Bhurat, D. Novaković, M. Canini, S. Shekhar, and D. Kostić, “Identifying and using energy-critical paths,” in *Proc. of CoNext*, 2011.
- [25] K. Kurroliya, S. Mohanty, K. Kanodia, and B. Sahoo, “Grey Wolf Aware Energy-saving and Load-balancing in Software Defined Networks Considering Real Time Traffic,” in *Proc. of ICICT*, 2020.
- [26] M. Zhang, C. Yi, B. Liu, and B. Zhang, “GreenTE: Power-aware traffic engineering,” in *Proc. of ICNP*, 2010.
- [27] T. Schüller, N. Aschenbruck, M. Chimani, M. Horneffer, and S. Schnitter, “Traffic Engineering using Segment Routing and Considering Requirements of a Carrier IP Network,” *IEEE/ACM ToN*, 2018.
- [28] S. Gay, P. Schaus, and S. Vissicchio, “REPETITA: Repeatable Experiments for Performance Evaluation of Traffic-Engineering Algorithms,” *CoRR*, 2017.
- [29] S. Knight, H. Nguyen, N. Falkner, R. Bowden, and M. Roughan, “The Internet Topology Zoo,” *JSAC*, 2011.
- [30] IBM, “IBM ILOG CPLEX Optimization Studio 20.1.0,” <https://www.ibm.com/docs/en/icos/20.1.0>, 2020.