



Energy Efficiency Analysis of Multilayer Networks with Time-Varying Traffic

Downloaded from: <https://research.chalmers.se>, 2026-04-21 17:21 UTC

Citation for the original published paper (version of record):

Knapinska, A., Lechowicz, P., Spadaro, S. et al (2025). Energy Efficiency Analysis of Multilayer Networks with Time-Varying Traffic. International Conference on Transparent Optical Networks. <http://dx.doi.org/10.1109/ICTON67126.2025.11125245>

N.B. When citing this work, cite the original published paper.

© 2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, or reuse of any copyrighted component of this work in other works.

Energy Efficiency Analysis of Multilayer Networks with Time-Varying Traffic

Aleksandra Knapińska^{1,2}, Piotr Lechowicz^{1,2}, Salvatore Spadaro³, Krzysztof Walkowiak¹

¹Wrocław Univ. of Science and Technology, Poland ²Chalmers Univ. of Technology, Sweden ³Univ. Politècnica de Catalunya, Spain
{aleksandra.knapinska, piotr.lechowicz, krzysztof.walkowiak}@pwr.edu.pl; salvatore.spadaro@upc.edu

Abstract—Global power consumption is growing each year, with a significant contribution from the ICT sector. Although the number of primary devices, such as ROADMs and routers, is usually constant, the number of active transceivers depends on the routing and allocation policy and can be tuned. This work demonstrates how dynamic network optimization aided by traffic prediction leads to a 16% power saving, expressed as the number of active transceivers, and a 15% provisioned traffic increase.

Index Terms—Energy efficiency, multilayer network, time-varying traffic, machine learning.

I. INTRODUCTION

The global power consumption is accelerating each year, driven by the development of new technologies and artificial intelligence (AI). Energy efficiency is thus an essential issue, recognized also by the European Commission, which has identified the ICT sector as a relevant contributor to global energy consumption [1]. In the case of optical networks, various devices consume significant amounts of energy. Although the number of primary devices, such as ROADMs and routers, is usually constant, the number of active transceivers depends on the chosen routing and allocation policy, and can be tuned programmatically. It has been shown both theoretically and experimentally that transceivers use similar amounts of energy for different modulation formats and transmission distances [2], [3]. Thus, the number of active transceivers in the network is a good measure of the overall energy usage. In particular, the calculations in [3] derived that the difference in energy used by transceivers configured to operate for different modulation formats and transmission distances is relatively small when compared to the much more significant difference between their idle (currently not supporting connections) and active (currently supporting connections) state. In turn, for versatility, the energy consumption of networks operated using specific algorithms can be expressed as the number of active transceivers. At the same time, transceiver utilization contributes meaningfully to the network operational costs, thus also serving as a measure for efficiency estimation [4], [5]. Moreover, we assume that the energy consumption overhead related to programmatic transceivers reconfiguration assisted with a machine learning (ML) prediction module is marginal with respect to transceivers operation consumption.

The existing research on energy-efficient multilayer networks focuses mainly on hardware-based optimization techniques (e.g., [2]). Traffic-aware approaches to switch the networking devices on and off depending on the demand changes (e.g., [6]) or use traffic grooming to achieve energy

savings by establishing fewer lightpaths (e.g., [7]), are also researched. However, the proposed algorithms rely on analytical models and do not employ traffic forecasting. The recent rapid technological development has enabled, however, new, data-driven optimization methods for the networking community to achieve the discussed energy and cost savings. Automated networks leverage network programmability, monitoring, and data analytics, facilitating dynamic self-optimization to match the current traffic conditions, primarily with the help of ML models [8]–[10]. Gained knowledge helps to discover daily traffic patterns enabling periodic reallocations of time-varying traffic, which enables notable improvements in blocking probability when using shorter reallocation periods [11]. However, an objective measure is required to quantify the potential benefits in realistic settings, i.e., without artificial network oversaturation to achieve bandwidth blocking.

The aim of this work is to explore the energy efficiency of multilayer networks with time-varying traffic. To this end, we analyze various dynamic network optimization variants to match the forecasted upcoming traffic conditions. As a baseline scenario, the network is configured to consider each connection request's predicted peak daily bitrate and avoid any reallocations during the day. Then, we propose three variants with dynamic optimization for different relocation periods, considering the time-varying nature of requests to different extents. Through experiments in a realistic setting, we analyze the number of active transceivers and the amount of provisioned traffic before blocking appears in scenarios with different reallocation frequencies. We quantify the benefits coming from utilizing the dynamic self-optimization. Finally, we examine the benefits of advance resource reservation for further optimization of the dynamic approach.

II. NETWORK MODEL AND ALLOCATION ALGORITHM

In this work, we consider a multilayer network with time-varying traffic of various services and applications. The bottom layer is an elastic optical network (EON), while the top one is a packet (IP) layer. The IP layer is a virtual topology of lightpaths set up in the EON layer. The layers exchange information about the lightpaths, e.g., their current used and free capacity. The cross-layer information exchange enables traffic grooming, where additional connection requests are provisioned in the existing lightpaths according to the remaining free bandwidth. That way, there is more stability in the bottom layer, and bandwidth wastage is minimized. For more details regarding the considered architecture we refer to [12].

The details of our multilayer routing and spectrum allocation (RSA) algorithm are provided in Alg. 1. First, connection requests are sorted by their initial required bandwidth

Algorithm 1 RSA in a multilayer IP-over-EON network

▷ Initial network setup
1: Sort requests by initial bitrate
2: **for** each request **do**
3: **if** a direct lightpath from its source to its destination exists and has enough free space **then**
4: groom the request into this lightpath
5: **end if**
6: **if** request not allocated **then**
7: set up a new lightpath in EON layer using Alg. 2
8: allocate the request into the newly-created lightpath
9: **end if**
10: **end for**

▷ Consecutive iterations in the network lifecycle
11: Sort requests by bitrate predicted for the upcoming period
12: **for** each request with bitrate decrease **do**
13: update free space in each segment of its routing path
14: **end for**
15: **for** each request with bitrate increase **do**
16: check if it still fits in each segment of its routing path
17: **if** request no longer fits in its path **then**
18: process it using Alg. 3
19: **end if**
20: **end for**

(line 1), and then they are processed one by one (line 2). The considered bitrate of each request is based on its prediction (more details in Sec. III). The algorithm first checks if a direct lightpath exists from the request's source to its destination and has enough free space (line 3). In that case, the algorithm performs grooming (line 4), which utilizes the previously existing lightpath topology, saving resources and improving stability, as shown in [7], [12]. If grooming is not possible, a new lightpath is set up in the EON layer (line 7) using Alg. 2 (explained below). In the following iterations, the algorithm proactively reacts to the bandwidth requirement changes of the requests. Each iteration corresponds to a set period (more details in Sec. III). In particular, the algorithm starts each iteration by sorting the requests by their currently considered (forecasted peak for the period) bitrate (line 11). Then, it first processes the ones with a required bandwidth decrease to free the resources (line 13). Next, it checks each increasing request to see if it still fits in its path (line 16). If not, it processes the request using Alg. 3 (explained below).

Algorithm 2 Lightpath allocation in the EON layer

Input: source node, destination node, requested bandwidth

Output: lightpath

1: consider $k = 10$ shortest paths between the requested source and destination node
2: **for** each candidate path **do**
3: using the First Fit heuristic, find a suitable channel for the requested bitrate, choose the most efficient modulation format supporting the required transmission distance
4: **end for**
5: sort the candidate paths by the highest frequency slot index of their found channels, ascending
6: set up a new lightpath on the path with found channel with the lowest channel-ending slot

Alg. 2 details how lightpaths are created in EON layer. A set of $k = 10$ (tuned in the preliminary experiments) shortest candidate paths is considered between node pairs (line 1). To choose the best path, we use a greedy algorithm to minimize spectrum usage (line 3). In particular, the algorithm tries to find a channel on each path with the most spectrally efficient modulation format, supported by the assumed transceiver model, for the path length and requested bitrate. We use the First Fit heuristic to find suitable optical channel on each

considered path. The path candidates with found channels are sorted according to the highest channel-ending slot, ascending (line 5). The lightpath is finally set up on the path with a found channel with the lowest channel-ending slot (line 6).

Algorithm 3 Grooming and routing in the IP layer

Input: source node, destination node, requested bandwidth

Output: allocation information

1: consider $k = 3$ shortest paths according to number of hops in the IP layer between the requested source and destination node
2: sort candidate paths by number of hops, ascending
3: **while** request not allocated and next candidate path exists **do**
4: groom the request to the candidate path if it has enough spare bandwidth
5: **end while**
6: **if** request not allocated **then**
7: set a new lightpath in the EON layer using Alg. 2
8: allocate the request into the newly-created lightpath
9: **end if**

Alg. 3 details the grooming and routing in IP layer. A set of $k = 3$ (tuned in the preliminary experiments) shortest paths in the IP layer is considered between node pairs (line 1). The candidate paths are sorted by the number of hops, ascending (line 2). The algorithm tries to groom the request into the shortest possible path (line 4). If there is no existing path in the IP layer with enough free space, a new lightpath is requested in the EON layer (line 7). The number of candidate paths in both layers was tuned to balance the path length and resource utilization (longer paths in the top layer consume resources in more lightpaths in the bottom layer).

III. TRAFFIC MODEL AND PERIODIC REALLOCATION

The traffic model considered in this work consists of time-varying connection requests (*intents*), i.e., with bitrate changing throughout the day. Each request follows a pattern of a specific network-based service or application (e.g., YouTube, TikTok), as published in the Sandvine report [13]. Because the report only provides the hourly averages of each traffic pattern, we use our approximation and noising algorithm (available as the Traffic Weaver Python package [14]) to create a continuous signal from the available bar plots. In turn, we use *semi-synthetic* data, where real traffic patterns are a base for a request generator. Such an approach allows for a fair simulation with realistic assumptions, considering the lack of detailed real-world data. Because of each service's unique purpose and properties, their traffic patterns, including the number, placement, and range of daily spikes, vary significantly (see Fig. 1). In our simulations, we assume multiple requests of various traffic volumes for each node pair. For more information about the traffic model, refer to [12].




Fig. 1: Illustration of the diversity of time-varying connection requests. Daily traffic patterns of various applications and services as provided in [13].

The proposed algorithm utilizes the dynamic self-optimization to the current traffic conditions. In other words, the network makes the routing and allocation decisions according to the requests' peak traffic within a set period. In

this work, we consider four period lengths (24 hours, 8 hours, 1 hour, 5 minutes), as illustrated in Fig. 2. The baseline algorithm is the traditional one-time (static) request allocation approach to match their forecasted daily peak traffic (purple dashed line on the plot). In this setting, the network does not exploit the daily traffic (dynamic) changes in any way. In turn, Alg. 1 only performs the initial allocation, and the considered bitrate of each connection request is its daily peak. On the contrary, assuming reallocations during the day, the algorithm adapts to the current conditions as explained in Sec. II. As it is easily noticeable from Fig. 2, more extended allocation periods result in vast overprovisioning. On the other hand, the network requires fewer reconfigurations within the day.




Fig. 2: Illustration of allocation for peak traffic for different periods of a day.

Intuitively, provisioning for periodic peak traffic requires traffic forecasting. To this end, as in [12], we employ prediction models for each request based on its one-month history. Following the recommendation from [15], the data-driven models are trained to learn the relationship between the current traffic level and the traffic a day and a week before at the same time, thus having two input features. Such feature engineering is effective in highly seasonal traffic (with daily and weekly patterns), which is present in backbone networks. At the same time, obtaining a forecast for the day ahead (specifically the daily peak) is possible, as the models do not use observations shorter than one day as features. As a ML algorithm of choice, we use Linear Regression. This algorithm proved to be a fast and reliable predictor of network traffic in preliminary experiments. At the same time, as in [16], it was the computationally cheapest method with satisfactory performance to follow the green networking paradigm.

Considering the smallest of the explored reallocation frequencies (5 minutes), the forecasts can be additionally applied to improve network operation further. To avoid unnecessary instabilities and enable more informed planning as the requests fluctuate, in the final algorithm variant, the maximum bitrate from the upcoming 15 minutes is considered for the allocation and reallocation decisions for each request having an increasing trend, thus leaving extra space for its forthcoming peak. In particular, Alg. 2 and 3 make the allocation and reallocation decisions based on the maximum 15-minute-ahead traffic forecasts, while Alg. 1 updates the network state every 5 minutes. That way, additional use of ML is utilized in the network for *advance reservation* (AR). On the plots in the experimental part (Sec. IV), we denote this algorithm modification as "with AR."

IV. EXPERIMENTAL EVALUATION

We run the experiments on the EURO28 and US26 topologies (see Fig. 3). As in [12], in our simulations, we use the Ciena Wavelogic 5 Extreme commercial transceiver model

with specifications as provided in [17]. The experiments were repeated ten times for each topology, with randomly generated sets of requests of bitrate within a 50-150 Gbps range (uniform distribution). The traffic load was increased by increasing the number of requests in the network. In the following part, we discuss the averaged results of our simulations.

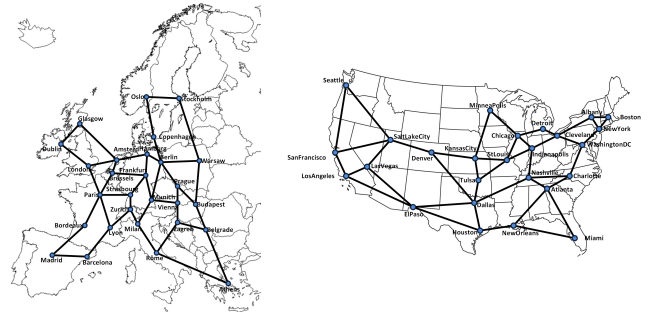


Fig. 3: Considered network topologies: EURO28 with 28 nodes and 82 links (left) and US26 with 26 nodes and 84 links (right).

In Fig. 4, we present the number of active transceivers as a function of traffic load for different analyzed reallocation periods for the considered topologies. Additionally, vertical dashed lines indicate the traffic loads where *bandwidth blocking probability* (BBP) of 1% for subsequent methods appear. The BBP is a measure commonly used to evaluate dynamic RSA algorithms (e.g., [18], [19]), specifying the ratio of provisioned bandwidth to the requested bandwidth; 1% BBP is commonly recognized in the literature as acceptable in operational network scenarios.

Clear benefits of the proposed approach enabling various frequencies of periodic reallocation are visible when compared to the traditional daily allocation. Furthermore, they increase with the increase in traffic load. As an example, consider the traffic load corresponding to 1% BBP for the 24-hour reallocation period (the lightest vertical dashed line on the plots) – the highest load that is possible to be provisioned in the network using the traditional daily allocation (47.5 Tbps for US26 and 40 Tbps for EURO28). For US26, the respective average number of active transceivers is 940.4 for the 24-hour reallocations (purple curve on the plot) and only 791.2 for the 5-minute reallocations (green curve on the plot), which is 16% fewer. In other words, thanks to more frequent reallocations, provisioning the same amount of traffic can be achieved using as much as 16% less energy. For EURO28, the analogous case yields 10% transceiver (power) savings. Moreover, to achieve the network saturation (1% BBP), an additional 7.5 Tbps (US26) or 5 Tbps (EURO28) of traffic can be provisioned in both topologies when using the dynamic 5-minute reallocation period (the darkest vertical dashed line on the plots) when compared to the traditional static case. Comparing the number of transceivers at the network saturation point (1% BBP) for each of the reallocation periods, we still achieve 7% (US26) or 5% (EURO28) transceiver savings. In other words, provisioning a larger amount of traffic can be achieved using as much as 7% less energy.

Overall, 1% BBP appears under heavier traffic load for consecutive reallocation periods, which is another notable benefit of frequent reallocations. In particular, for US26, the network saturation (vertical dashed lines on the plots) was noted for a 10% and 15% higher traffic load for 1-hour- and

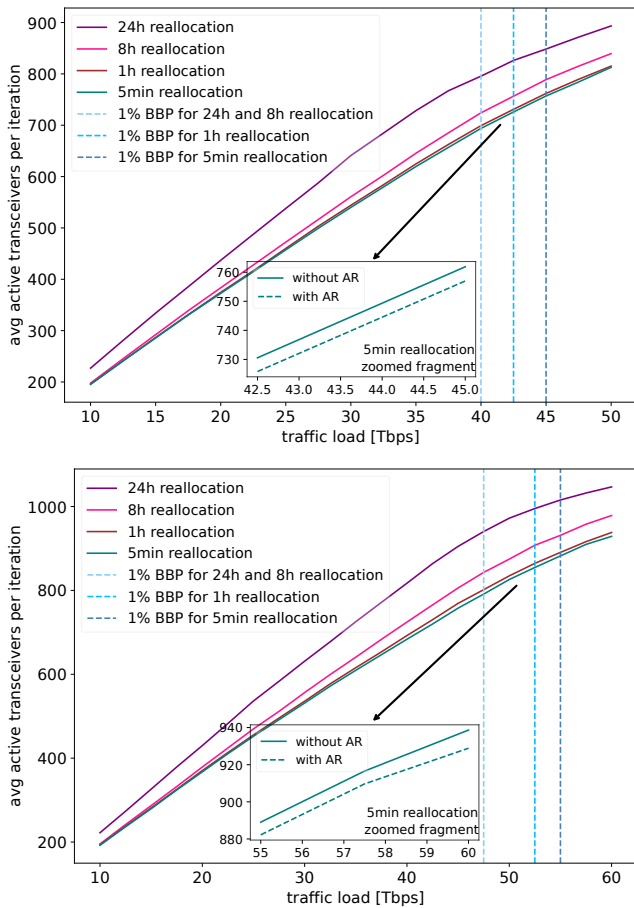


Fig. 4: Number of active transceivers as a function of traffic load for the considered reallocation frequencies, vertical lines indicate the traffic load corresponding to 1% BBP of various reallocation frequencies. US26 (top); EURO28 (bottom).

5-minute reallocation compared to the 24-hour one. The trends for EURO28 are equivalent. Together with the transceiver usage curves, more traffic can be provisioned in the network using the same or less energy. It is evident by comparing the number of active transceivers for the 1-hour and 5-minute reallocation periods. In both topologies, the numbers in terms of the number of active transceivers are very similar; however, the 1% BBP appears later for the more frequent relocations. Finally, even though in our testing, the 24-hour and 8-hour reallocation periods yielded the same amount of accepted traffic for 1% BBP, the number of active transceivers (used power) was 10% (US26) and 9% (EURO28) smaller for the more frequent relocations.

Furthermore, considering the addition of AR to the most efficient 5-minute reallocation period enables even more reduction in transceiver usage and thus power saving. The bottom parts of Fig. 4 illustrate zoomed-in fragments of the plots around each network’s saturation point. It is clearly visible that the number of active transceivers is smaller for the same traffic load in both tested topologies with the additional use of ML.

Finally, we repeated our experiments for other commercial transceiver models explored in [20]: Ciena Wavelogic AI from 2016 and Ciena Wavelogic 3 from 2012. All the trends discussed above also hold for the older equipment, indicating transceiver and, thus, energy saving when using the data-driven adaptive provisioning, therefore confirming the versatility of our study. Considering their often worse spectral efficiency,

the benefits from more frequent reallocations are even more noticeable. In turn, updating the network operation policy is worthwhile despite the used equipment for power saving and increased traffic load provisioning.

V. CONCLUSIONS

In this work, we explored the energy consumption savings that dynamically-optimized multilayer networks bring. We discussed how the number of active transceivers is a good measure of the network’s energy efficiency and operational costs, thus being a versatile network performance estimator. We then provided ways to reduce it, demonstrating the benefits of dynamic network self-optimization to the current traffic conditions, thus decreasing overprovisioning and enabling various resource savings. Through experiments on two benchmark topologies, we showed how a 16% energy saving can be obtained by increasing the relocation frequency in our data-driven approach. Furthermore, we showed how, with the same or less energy spent, more traffic can be provisioned when using more frequent reallocations. Finally, we demonstrated the benefits of employing additional ML-based traffic prediction for further optimization. In the future, we aim to further investigate data-driven intent-based network optimization methods and develop a QoS-aware framework that diversifies request provisioning according to their individual requirements.

REFERENCES

- [1] European Commission, Directorate-General for Communication, “Supporting the green transition: shaping europe’s digital future,” 2020.
- [2] M. Radovic *et al.*, “Experimental optimization of power-aware superchannels in elastic optical networks,” ONDM 2023.
- [3] P. Poggiolini *et al.*, “The GN-model of fiber non-linear propagation and its applications,” J. Light. Technol. 2014.
- [4] J. Perelló *et al.*, “Reducing the number of transceivers with probabilistic constellation shaping in flex-grid over MCF optical backbone networks,” ICC 2023.
- [5] M. M. Hosseini *et al.*, “Multi-period planning in metro-aggregation networks using point-to-multipoint transceivers,” GLOBECOM 2022.
- [6] G. Rizzelli *et al.*, “Energy efficient traffic-aware design of on-off multilayer translucent optical networks,” Comput. Netw. 2012.
- [7] H. Khodakarami *et al.*, “Quality of service provisioning and energy minimized scheduling in software defined flexible optical networks,” J. Opt. Commun. Netw. 2016.
- [8] A. Leivadreas *et al.*, “A survey on intent based networking,” IEEE COMST 2023.
- [9] L. Velasco *et al.*, “Intent-based networking and its application to optical networks,” J. Opt. Commun. Netw. 2022.
- [10] Y. Sharma *et al.*, “SLA management in intent-driven service management systems: A taxonomy and future directions,” ACM Comput. Surv. 2023.
- [11] A. Włodarczyk *et al.*, “Machine learning assisted provisioning of time-varying traffic in translucent optical networks,” IEEE Access 2024.
- [12] A. Knapieńska *et al.*, “On advantages of traffic prediction and grooming for provisioning of time-varying traffic in multilayer networks,” ONDM 2023.
- [13] Sandvine, “The mobile internet phenomena report,” May 2021.
- [14] P. Lechowicz *et al.*, “Traffic weaver: Semi-synthetic time-varying traffic generator based on averaged time series,” SoftwareX 2024.
- [15] A. Knapieńska *et al.*, “Machine-learning based prediction of multiple types of network traffic,” ICCS 2021.
- [16] R. Goścień, “Traffic-aware service relocation in software-defined and intent-based elastic optical networks.”
- [17] Ciena, <https://www.ciena.com/insights/data-sheets/800g-wavelogic-5-extreme-motr-module.html>.
- [18] A. S. Santos *et al.*, “Application-aware service degradation in elastic optical networks,” IEEE TNSM 2022.
- [19] E. Biernacka *et al.*, “Elastic optical bypasses for traffic bursts,” Comput. Commun. 2019.
- [20] A. Knapieńska *et al.*, “Performance analysis of multilayer optical networks with time-varying traffic,” ICTON 2023.