



Provisioning Anycast Connections in Traffic-Prediction-Assisted Multilayer Networks

Downloaded from: <https://research.chalmers.se>, 2025-10-21 02:54 UTC

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

Knapinska, A., Lechowicz, P., Walkowiak, K. (2025). Provisioning Anycast Connections in Traffic-Prediction-Assisted Multilayer Networks. International Conference on Transparent Optical Networks. <http://dx.doi.org/10.1109/ICTON67126.2025.11125440>

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.

Provisioning Anycast Connections in Traffic-Prediction-Assisted Multilayer Networks

Aleksandra Knapieńska^{1,2}, Piotr Lechowicz^{1,2}, Krzysztof Walkowiak¹

¹Department of Systems and Computer Networks, Wrocław University of Science and Technology, Wrocław, Poland

²Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden
{aleksandra.knapinska, piotr.lechowicz, krzysztof.walkowiak}@pwr.edu.pl

Abstract—Anycast, as a one-to-one-of-many communication strategy directing user requests to the nearest or best replica of content stored across multiple data centers (DCs), has gained widespread adoption for improving network performance by reducing latency, and enhancing resilience. However, the growing reliance on anycast in large-scale networks has introduced new challenges, particularly in terms of congestion on the links connecting (DCs), which can degrade network performance and reliability. This paper explores the impact of provisioning anycast connections on multilayer backbone optical networks, particularly under time-varying traffic conditions. We investigate the effects of varying anycast traffic proportions and the number of (DCs) on network performance and resource utilization. Our findings demonstrate the potential of a traffic-prediction-based algorithm to optimize anycast provisioning, contributing to more efficient network management and providing a foundation for future research in dynamic resource allocation and multilayer network design.

Index Terms—anycast, multilayer network, time-varying traffic

I. INTRODUCTION

Anycast, defined as *one to one-of-many* communication, leverages content replicas stored in data centers (DCs) to improve network performance and resilience [1], [2]. This approach enables efficient content distribution by directing user requests to the nearest or otherwise chosen replica, reducing latency and enhancing user experience. Anycast is widely applicable across diverse network types, including terrestrial, optical, and satellite communications [3]–[5]. It is increasingly being adopted as a strategy for improving network robustness, particularly in the face of evolving cybersecurity threats such as AI-driven attacks, which can target centralized infrastructures [6]–[8]. By ensuring that content is replicated across multiple DCs, anycast provides not only faster content delivery but also redundancy, minimizing the risk of service disruption if a single DC is compromised or experiences failure.

While anycast enhances resilience and performance by aggregating spread content into several replicas, its growing adoption has led to increased traffic on DC connections, creating new network management challenges. The increasing reliance on DCs for content storage and distribution has resulted in significant congestion on the links connecting these centers to the broader network [9]. This congestion poses challenges to maintaining high network performance and reliability. As a result, prioritizing DCs in network upgrade strategies has become a crucial consideration for network operators [10]. Another promising idea is to combine anycast and multicast for increased efficiency [11]. Nevertheless, addressing congestion requires advanced routing techniques and resource allocation

strategies that can dynamically adapt to changing network conditions.

To address these challenges, network operators increasingly turn to multilayer network models, which offer a powerful framework for enhancing network efficiency and management, particularly in complex, large-scale networks. These models integrate different network layers (typically optical and packet), to optimize resource utilization and improve resilience [12]–[14]. In backbone optical networks, where high-capacity and low-latency data transmission are essential, multilayer optimization plays a critical role in ensuring seamless service delivery. By leveraging multilayer optimization, network operators can better allocate resources, reduce operational costs, and improve overall performance. However, effective network design must also account for dynamic traffic patterns, as demand fluctuates over time due to user behavior, application requirements, and external factors. Ensuring that optimization algorithms can adapt to these variations is critical for maintaining network stability and efficiency. Incorporating traffic prediction significantly enhances this adaptability [15], [16].

In this paper, we explore how provisioning anycast connections impacts the operation of multilayer backbone optical networks with time-varying traffic. In our experimental evaluation, we examine the impact of the portion of anycast traffic and the number of DCs on network performance and resource utilization. Our study demonstrates the feasibility of utilizing a traffic-prediction-based algorithm for effective anycast provisioning, creating a basis for further research in this area.

In the remainder of the paper, we detail our simulation setup, providing an overview of the network optimization algorithm and traffic model in Section II. Furthermore, we present the results of our experimental evaluation on two benchmark topologies and multiple test scenarios in Section III. Finally, Section IV concludes the paper.

II. EXPERIMENTAL SETUP

We consider a two-layer network model comprising an EON layer at the bottom and an IP layer at the top. The upper layer forms a virtual topology of lightpaths established in the lower layer.

As the baseline routing and spectrum allocation (RSA) policy, we adopt the algorithm proposed in [15]. In brief, time-varying connection requests are processed sequentially in each iteration, sorted by their predicted bitrate over the upcoming five minutes. These predictions are generated by a regression model trained on each request’s month-long history, using traffic patterns from one day and one week prior as




Fig. 1: Considered data center locations as in [1] for EURO28 (left) and US26 (right). The experiments consider three cases with varying numbers of active DCs: (i) three active DCs located at blue nodes, (ii) five active DCs located at blue and green nodes, and (iii) seven active DCs located at blue, green, and pink nodes.

input features. These forecasts guide all (re)allocation and grooming decisions. In the physical layer, the ten shortest paths (based on distance in kilometers) are considered, while in the virtual layer, the three shortest paths (based on hop count) are evaluated. The algorithm first attempts to accommodate each connection using existing lightpaths (virtual links); if this fails, a new lightpath is provisioned in the physical layer.

Anycast connections differ from unicast ones in that a DC serves as their source or destination. As such, they have more potential path candidates, since the destination DC can be selected from a set of options. Furthermore, we allow the DC to change during potential reallocations. This flexibility gives anycast connections a higher likelihood of being provisioned compared to unicast ones, especially as the number of available DCs increases. In our experiments, we rank DCs for each connection request based on geographic distance (in kilometers).

In our experiments, we assume multiple connection requests per node pair, uniformly distributed across the topology. These represent various network services and applications and were generated using *Traffic Weaver* [17], with bitrates scaled to the 50–100 Gbps range. Connections are classified as unicast (city-to-city) or anycast (city-to-DC or DC-to-city). We examine three traffic scenarios, with anycast traffic accounting for 10%, 30%, and 50% of the total connection requests.

The experimental evaluation is performed on two topologies from *SNDlib* [18], namely EURO28 and US26. These topologies, along with the selected DC locations, are shown in Fig. 1. The DC sites were selected as in [1]. We investigate three test cases based on the number of active DCs: (i) three active DCs located at the blue-marked nodes, (ii) five active DCs, including both blue and green nodes, and (iii) seven active DCs, located at the blue, green, and pink nodes.

To achieve stable results and derive meaningful conclusions, we evaluate the network operation of each topology across 35 sets of requests with 36 traffic loads, performing a total of 1260 simulations per test case per topology. In the following section,

we discuss the averaged results. In our simulations, we assume the Ciena Wavelogic 5 Extreme commercial transceiver model, with specifications provided in [19].

III. RESULTS AND DISCUSSION

The primary performance metric we discuss is the bandwidth blocking probability (BBP), describing the percentage of blocked traffic to all traffic for a given load. Fig. 2–3 illustrate the BBP across the evaluated scenarios for US26 and EURO28 topologies, respectively.




Fig. 2: BBP as a function of traffic load for US26.

Starting with US26, the overall trend indicates that more available DCs and lower anycast percentages generally lead to a lower BBP, whether using three or five DCs. Notably, reducing the anycast share from 50% to 10% significantly improves performance in both configurations. However, when the number of DCs increases to seven, BBP drops substantially, regardless of the anycast traffic proportion. Interestingly, the trend reverses in this case: a higher anycast share actually increases the probability of a successful connection.

In the case of EURO28, the overall BBP is generally higher than for US26, particularly under higher loads. The impact




Fig. 3: BBP as a function of traffic load for EURO28.

of the anycast proportion and the number of DCs also differs slightly. For both three and five DCs, the anycast proportion is the primary factor influencing the BBP. Similar to US26, a lower anycast proportion results in a lower BBP. However, increasing the number of DCs to seven dramatically improves network performance and reverses the trend. In this case, a higher anycast proportion actually lowers the BBP, facilitating the provisioning of more connections. This shows, how better scalability is enabled by a greater diversity of candidate routing paths.

To gain further insights, we now focus specifically on the anycast traffic. Figures 4–5 show the BBP for the anycast portion of traffic in US26 and EURO28, respectively. In these cases, the number of available DCs plays a crucial role, especially for EURO28. As expected, increasing the number of DCs allows for the provisioning of more anycast connection requests. For each proportion of anycast traffic, increasing the number of DCs from three to five reduces the BBP for anycast connections. A further increase from five to seven DCs results in a significant performance boost, dramatically improving results. Therefore, even when overall blocking remains similar, expanding the number of DCs is critical for minimizing blocking in the anycast portion of traffic. Finally, the lowest anycast blocking occurs when the proportion of anycast traffic within the network is smallest.




Fig. 4: BBP within anycast connection requests as a function of traffic load for US26.

The next set of results examines the impact of anycast traffic on resource utilization. For space constraints, we will focus




Fig. 5: BBP within anycast connection requests as a function of traffic load for EURO28.

on the example of the EURO28 topology and the 50% ratio of anycast traffic.

Figure 6 shows the number of network links with at least 50% of their slots occupied as a function of traffic load. This metric illustrates how traffic distribution across the topology improves as the number of active DCs increases. Between three and five available DCs, the effect is minimal, as requests are generally served by the nearest DC, and all are located relatively close to one another. However, adding two more DCs, especially as the traffic load increases, enables a much more even distribution of traffic across the network. Specifically, links near the DCs become more congested, prompting the use of alternative paths, which becomes possible as the number of DCs grows. With more available DCs, the algorithm can accommodate more connections while spreading the load more evenly throughout the network.




Fig. 6: Number of network links with at least 50% slots occupied as a function of traffic load, EURO28 topology.

Another measure related to resource allocation is the number of allocated lightpaths, illustrated in Fig. 7. Here, the number of available DCs starts playing a role under heavier traffic loads. More diversity and choice results in more possibilities of setting up new lightpaths.

Further insight can be gained from Fig. 8, which shows the number of activated transceivers. Although the overall trend aligns with the lightpath allocations, the differences are more pronounced. This indicates that the additional lightpaths span longer distances, necessitating the activation of a greater number of transceivers.




Fig. 7: Number of allocated lightpaths as a function of traffic load, EURO28 topology.




Fig. 8: Number of activated transceivers as a function of traffic load, EURO28 topology.

In summary, increasing the number of available DCs enables the successful accommodation of a larger volume of traffic. Anycast connections benefit particularly strongly, with their provisioning improving significantly as more DCs are added. Additionally, network resources are utilized more efficiently, allowing for the establishment of a greater number of lightpaths.

IV. CONCLUSIONS

In this paper, we investigated the provisioning of anycast traffic in traffic-prediction-assisted multilayer networks. We analyzed the impact of the number of DCs and the proportion of anycast traffic on overall network performance. Experimental evaluation on two benchmark topologies revealed several insights, particularly highlighting the benefits of increasing the number of available DCs.

Our results showed that a higher number of DCs consistently improves the BBP, especially under high traffic loads and for anycast connections. The improvements were most significant when scaling from five to seven DCs, underlining the value of infrastructure scalability. Additionally, lower anycast traffic ratios led to better performance when the number of available DCs is smaller. Beyond BBP, we observed that more DCs also enhance resource utilization, supporting more lightpaths and distributing traffic more efficiently throughout the network.

These findings suggest that strategic placement of additional DCs can significantly improve both the efficiency and

robustness of multilayer network architectures aided by traffic prediction.

In the future, we plan to investigate the notion of DC placement and the diversification of their content.

ACKNOWLEDGMENT

This work was supported by the National Science Center, Poland under Grant 2019/35/B/ST7/04272.

REFERENCES

- [1] K. Walkowiak, *Modeling and optimization of cloud-ready and content-oriented networks*. Springer, 2016, vol. 56.
- [2] M. Baentsch, L. Baun, G. Molter, S. Rothkugel, and P. Sturn, "World wide web caching: The application-level view of the internet," *IEEE Communications Magazine*, vol. 35, no. 6, pp. 170–178, 1997.
- [3] L. Zhang, Z. Guo, H. Su, and W. Zhao, "CAAST: Optimizing data communications in satellite networks through cache and anycast," *Computer Networks*, p. 111066, 2025.
- [4] S. Nan, J. Yuan, X. Li, Q. Zhang, and M. Guo, "Routing, modulation and spectrum allocation of dynamic anycast based on ant colony optimization in elastic optical networks," *Optical Fiber Technology*, vol. 74, p. 103049, 2022.
- [5] P. Afsharlar, A. Deylamsalehi, and V. M. Vokkarane, "Delayed spectrum allocation in elastic optical networks with anycast traffic," *OSA Continuum*, vol. 4, no. 8, pp. 2118–2132, 2021.
- [6] C. Develder and B. Jaumard, "Dimensioning backbone networks for multi-site data centers: exploiting anycast routing for resilience," in *7th International Workshop on Reliable Networks Design and Modeling (RNDM)*, 2015, pp. 34–40.
- [7] M. Furdek, C. Natalino, A. Di Giglio, and M. Schiano, "Optical network security management: requirements, architecture, and efficient machine learning models for detection of evolving threats," *Journal of Optical Communications and Networking*, vol. 13, no. 2, pp. A144–A155, 2021.
- [8] S. Das and M. Chatterjee, "Energy efficient resource aware protection with rapid failure recovery in cloud-ready elastic optical networks," *Optical Switching and Networking*, vol. 55, p. 100793, 2025.
- [9] R. Gościński, "Explainable artificial intelligence-based framework for efficient content placement in elastic optical networks," *Expert Systems with Applications*, vol. 262, p. 125541, 2025.
- [10] P. Lechowicz, R. Gościński, R. Rumipamba-Zambrano, J. Perello, S. Spadaro, and K. Walkowiak, "Greenfield gradual migration planning toward spectrally-spatially flexible optical networks," *IEEE Communications Magazine*, vol. 57, no. 10, pp. 14–19, 2019.
- [11] M. Aibin, R. Gościński, and K. Walkowiak, "Multicasting versus anycasting: How to efficiently deliver content in elastic optical networks," in *18th International Conference on Transparent Optical Networks (ICTON)*, 2016, pp. 1–4.
- [12] I. Tomkos, Č. Rožić, M. Savi, P. Sköldström, V. Lopez, M. Chamania, D. Siracusa, C. Matrakidis, D. Klonidis, and O. Gerstel, "Application aware multilayer control and optimization of elastic WDM switched optical networks," in *Optical Fiber Communications Conference (OFC)*, 2018, pp. 1–3.
- [13] V. Lopez, D. Klonidis, D. Siracusa, C. Rozic, I. Tomkos, and J. P. Fernandez-Palacios, "On the benefits of multilayer optimization and application awareness," *Journal of Lightwave Technology*, vol. 35, no. 6, pp. 1274–1279, 2017.
- [14] Č. Rožić, D. Klonidis, and I. Tomkos, "A survey of multi-layer network optimization," in *20th International Conference on Optical Network Design and Modeling (ONDM)*, 2016, pp. 1–6.
- [15] A. Knapińska, P. Lechowicz, S. Spadaro, and K. Walkowiak, "On advantages of traffic prediction and grooming for provisioning of time-varying traffic in multilayer networks," in *27th International Conference on Optical Network Design and Modeling (ONDM)*, 2023, pp. 1–6.
- [16] A. Włodarczyk, A. Knapińska, P. Lechowicz, and K. Walkowiak, "Machine learning assisted provisioning of time-varying traffic in translucent optical networks," *IEEE Access*, vol. 12, pp. 110 193–110 212, 2024.
- [17] P. Lechowicz, A. Knapińska, A. Włodarczyk, and K. Walkowiak, "Traffic weaver: Semi-synthetic time-varying traffic generator based on averaged time series," *SoftwareX*, vol. 28, p. 101946, 2024.
- [18] S. Orlowski, R. Wessälly, M. Pióro, and A. Tomaszewski, "Sndlib 1.0—survivable network design library," *Networks: An International Journal*, vol. 55, no. 3, pp. 276–286, 2010.
- [19] Ciena, <https://www.ciena.com/insights/data-sheets/800g-wavelogic-5-extreme-motr-module.html>.