Hollow-Core Fiber Upgrading for Carbon-Aware Workload Placement in Distributed Data Centers

Yanran Xiao*†, Wei Wang*, Mëmëdhe Ibrahimi†, Qiaolun Zhang†, Jiaheng Xiong†, Jie Zhang*, Massimo Tornatore†
* School of Electronic Engineering; Beijing University of Posts and Telecommunications; Beijing, China
† Department of Electronics, Information and Bioengineering; Politecnico di Milano; Milan, Italy
Corresponding author: jie.zhang@bupt.edu.cn

***Abstract*—The rapid growth of AI applications is increasing the carbon emissions of Data Centers (DCs). To reduce carbon emissions of DCs, carbon-aware workload placement approaches have been developed to deploy computing workloads to remote DCs with cleaner energy. However, such workload shifting is constrained by the latency budget, as the service provider must manage the QoS requirements of applications. That is, reducing the network propagation latency is beneficial for expanding the geographical range that can meet the latency constraint. This paper shows how Hollow-Core Fiber (HCF) can enable more effective carbon-aware workload placement. We study selective HCF upgrading to expand the set of latency-feasible data centers. We jointly design (i) a carbon-driven HCF upgrading heuristic that prioritizes fiber upgrades that lead to higher carbon reductions, and (ii) a carbon-aware workload placement strategy that exploits the expanded feasible region. Simulation results show that combining selective HCF upgrading with carbon-aware placement can reduce carbon emissions by almost 20%.**

***Index Terms*—Hollow-core fiber upgrading, carbon-aware workload placement, geo-distributed data centers**

I. Introduction

According to Independent Commodity Intelligence Services estimates, Data Centers (DCs) in Europe consumed approximately 96 TWh of electricity in 2024, accounting for about 3.1% of total European power demand [1]. In response to this growing footprint, the International Telecommunication Union has defined Paris-aligned emission trajectories for the Information and Communications Technology sector, including DCs, highlighting the urgency of science-based decarbonization strategies [2].

The rapid growth of large-scale AI workloads is increasing the demand for computing capacity and related energy consumption. To mitigate the resulting carbon emissions, existing studies have explored a range of approaches that leverage the spatial and temporal availability of cleaner electricity across geographically-distributed DCs. Among these approaches, carbon-aware workload placement has emerged as an effective mechanism to better utilize low–carbon-intensity energy by shifting computation to cleaner locations [3]. However, such workload placement decisions are inherently constrained by network latency, as workloads can only be migrated to DCs that are reachable within acceptable end-to-end delay bounds over the wide-area network [4] [5]. A substantial body of research has investigated carbon-aware scheduling, migration, and demand shaping at the computing layer [6]–[8].

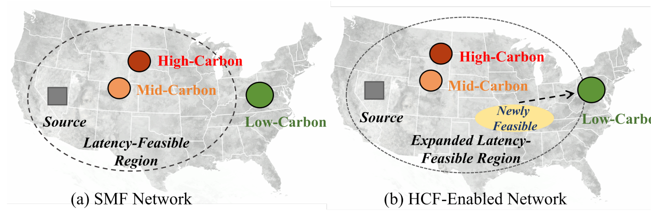

Fig. 1. Latency-feasible data center regions under SMF and HCF-enabled networks.

Recently, Hollow-Core Fiber (HCF) has reached a sufficient level of technological maturity to provide substantially lower propagation latency than conventional solid-core Single-Mode Fiber (SMF) [9]. By guiding light through an air-filled core, HCF enables near-vacuum propagation and reduces latency by roughly one third compared to SMF [10] [11]. This latency reduction has the potential to relax wide-area delay constraints [12], thereby enabling carbon-aware workload placement across a broader geographic region. As illustrated in Fig. 1, conventional SMF-based networks confine execution to nearby sites and often exclude cleaner data centers due to strict latency limits. In contrast, HCF expands the latency-feasible region [13], bringing previously excluded green data centers into the feasible set and thereby unlocking additional opportunities for carbon-aware workload placement. Despite these advantages, the high cost and technological complexity of HCF [14] make full-network upgrades economically impractical, necessitating selective deployment under a constrained HCF upgrading length budget. Recent industrial efforts further indicate that HCF is typically applied only to a subset of links rather than ubiquitously across the entire network [15] [16].

Under practical budget constraints, selectively upgrading a limited number of fiber links provides a potential mechanism for reducing the carbon footprint of geo-distributed DCs. This work examines how selective HCF upgrades can enable broader system-level carbon reductions by expanding the feasible region for carbon-aware workload placement. In this context, we first consider optical-layer latency reduction enabled by HCF upgrading, and then introduce the concept of network-enabled carbon-aware workload placement that

This work is supported by the Fundamental Research Funds for the Central Universities (510224040, ZDYY202102) and the Open Fund of State Key Lab of Information Photonics and Optical Communications (IPOC2025ZZ03).

978-3-903176-78-2

leverages the expanded latency-feasible region, as illustrated in Fig. 2. We develop a carbon-driven HCF upgrading heuristic that prioritizes fiber upgrades according to their estimated carbon-reduction impact, together with a carbon-aware workload placement strategy. Simulation results indicate that limited HCF upgrades can lead to measurable reductions in system-wide carbon emissions, with reductions of up to 19.4%.

II. Network and Workload Model

A. DC and Network Model

We consider a geo-distributed computing infrastructure interconnected by a wide-area optical backbone network. The network is modeled as a directed graph $G = (V, E)$, where nodes represent DC sites and edges represent inter-DC optical links.

Each node $v \in V$ corresponds to a DC cluster characterized by three attributes: (i) Compute capacity C_v, denoting the maximum amount of servers that can be supported; (ii) Grid carbon intensity γ_v, defined as the emission factor (kg CO_2e/kWh) of the electricity supplying node v, which reflects the regional generation mix (e.g., wind, solar, fossil) at that location; and (iii) Power usage effectiveness PUE_v, capturing the efficiency of the facility's cooling and supporting infrastructure.

Following the standard definition, PUE_v is the ratio of the total facility power consumption to the IT equipment power consumption (e.g., CPUs and GPUs) at node v. A lower PUE_v therefore indicates a more energy-efficient DC, as a smaller fraction of power is consumed by non-IT overheads such as cooling and power distribution.

The effective carbon intensity of executing computation at node v is therefore given by

$$CI_v = \gamma_v \times \mathrm{PUE}_v,$$

which jointly reflects the regional energy mix and the DC operational efficiency.

Each directed edge $e_{uv} \in E$ represents a physical inter-DC fiber link with geographical length L_{uv}, deployed using either SMF or HCF. Propagation latency is modeled via round-trip time (RTT), which scales linearly with link length and inversely with propagation speed. Selective HCF upgrading therefore reduces end-to-end latency between distant data centers without altering network topology or bandwidth capacity.

B. Workload Model

We model a set of heterogeneous workloads $\mathcal{T}$ representative of AI-centric inter-DC traffic. Each workload $t \in \mathcal{T}$ is defined by a tuple (s_t, D_t, B_t, τ_t). The source s_t denotes the geographical origin of the request. Each task specifies a compute demand D_t (e.g., number of servers required) and a bandwidth demand B_t for data transfer across the network. The latency τ_t defines a strict upper bound on the allowable round-trip network latency between the source location and the selected execution DC. This constraint captures application-level service requirements and plays a central role in determining the set of feasible execution sites.

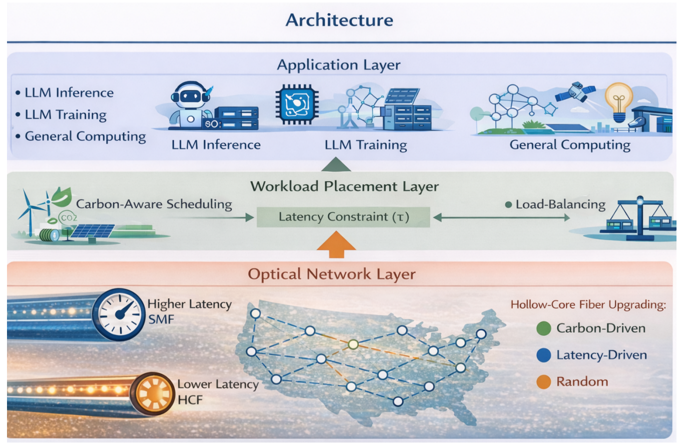

Fig. 2. Layered system architecture of carbon-aware computing.

C. Problem Statement

The problem studied in this work can be stated as follows: **Given** a geo-distributed inter-DC network topology with known link lengths, regional carbon intensity and PUE values, and a set of computing tasks arriving with latency requirements and resource demands, **determine** (i) which optical links should be upgraded from SMF to HCF in the optical network layer of Fig. 2 and (ii) how to assign each task to a DC, **subject to** a limited HCF upgrading budget, DC computing capacity constraints, and end-to-end latency constraints in workload placement layer of Fig. 2. HCF upgrades reduce propagation delay on selected links and thus expand the set of latency-feasible execution sites, while task placement decisions determine the resulting carbon emissions. The overall **objective** is to reduce the carbon emissions of geo-distributed DCs.

Our objective focuses on reducing data-center-side GHG emissions by enlarging the latency-feasible set of destination DCs through HCF upgrading. We do not explicitly account for the additional network energy consumption that may arise from longer multi-hop routes. This modeling choice is motivated by the fact that, in geo-distributed service infrastructures, the carbon impact of data center operation is typically much larger than that of optical transmission, while the main role of the transport network in our framework is to determine latency feasibility. A joint assessment of DC-side and network-side emissions is left for future work.

III. Methodology

Based on the network model defined in Section II, this section presents the algorithmic framework. The key design principle is to first shape end-to-end latency at the network layer, and then perform workload placement at the compute layer under the resulting latency constraints. We first describe the carbon-driven HCF upgrading heuristic, which determines the latency of the network, and then introduce the carbon-aware workload placement strategy that operates on the resulting latency-feasible region.

A. Carbon-Driven HCF Upgrading (CDHU)

The carbon-driven HCF upgrading heuristic aims to selectively reduce network latency along paths that yield the highest potential carbon-emission reduction at the computing layer. Rather than optimizing purely for physical latency savings, this heuristic explicitly couples network upgrades with carbon-aware workload placement decisions.

The heuristic consists of three key steps: (i) estimating potential traffic demand under an ideal low-latency network, (ii) quantifying the carbon attractiveness of destination DCs, and (iii) ranking candidate links according to their joint contribution to latency relief and carbon reduction.

1) Full-HCF Traffic Estimation: To identify links that are critical for reaching low-carbon regions, we estimate the traffic that would be routed toward cleaner DCs if latency were no longer a limiting factor. Specifically, we construct a *full-HCF network* in which all optical links are assumed to be upgraded to HCF, removing propagation latency constraints.

Under this idealized network, tasks are scheduled using the carbon-aware strategy, producing latency-unconstrained routing paths. We denote by F_{uv}^{fullHCF} the aggregate traffic volume traversing link (u, v) across all tasks in the full-HCF scenario. A large F_{uv}^{fullHCF} indicates that the link lies on a preferred path toward carbon-efficient execution sites once latency barriers are relieved.

2) Carbon Attractiveness of DCs: Each DC node v is associated with a carbon intensity

$$CI_v = \text{PUE}_v \cdot \gamma_v,$$

where γ_v denotes the regional carbon intensity and PUE_v denotes the PUE of the DC.

We define the *carbon attractiveness* of node v as

$$G_v = CI_{\max} - CI_v + \varepsilon,$$

where $CI_{\max}$ is the maximum carbon intensity across all nodes and $\varepsilon > 0$ avoids zero-attractiveness cases. Nodes with lower carbon intensity therefore have higher attractiveness.

3) Link Scoring and Selection: After estimating the carbon-driven traffic demand under the full-HCF network and quantifying the carbon attractiveness of destination nodes — two processes that are computed independently — we combine these factors to score each candidate SMF link.

For each link (u, v), we compute a carbon-driven upgrading score

$$Score_{uv}^{\text{carbon}} = F_{uv}^{\text{fullHCF}} \cdot G_v + F_{vu}^{\text{fullHCF}} \cdot G_u.$$

This symmetric formulation accounts for bidirectional traffic demand and jointly captures (i) how much latency-constrained traffic would benefit from upgrading the link, and (ii) how much carbon reduction can be achieved by enabling access to the corresponding low-carbon destinations.

Links are ranked in descending order of $Score_{uv}^{\text{carbon}}$. Given a total HCF upgrading budget B, expressed as the maximum aggregate fiber length that can be replaced by HCF, links are greedily upgraded until the budget is exhausted. This process concentrates upgrades on a small subset of critical links that most effectively unlock low-carbon execution sites.

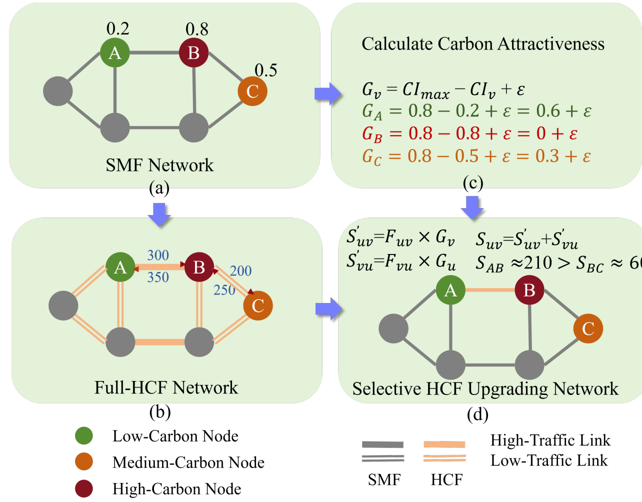

Fig. 3. Illustration of the carbon-driven HCF upgrading heuristic.

Fig. 3 provides an example of the proposed workflow. Fig. 3(a) shows the initial SMF network, where each node is associated with a different regional carbon intensity. In this baseline setting, propagation delay limits how far latency-sensitive traffic can be shifted. To reveal how traffic would redistribute if latency were no longer restrictive, Fig. 3(b) depicts the full-HCF network, where all links are hypothetically upgraded to HCF. Under these relaxed latency constraints, more traffic is shifted toward cleaner regions. The link labels represent the resulting bidirectional traffic volumes, revealing the potential migration patterns. Based on these potential destinations, Fig. 3(c) illustrates the carbon attractiveness of each node, where lower-carbon nodes receive higher attractiveness values, indicating greater carbon-reduction benefit when selected as execution locations. Finally, combining the full-HCF traffic patterns and node attractiveness, Fig. 3(d) shows the selective HCF upgrading result. Links that carry more traffic toward carbon-attractive nodes receive higher priority for upgrade under the HCF budget constraint. Overall, the figure demonstrates how potential traffic tendencies and carbon-aware destination preferences are jointly used to guide carbon-efficient, selective HCF upgrading.

B. Carbon-Aware Workload Placement (CAWP)

Once the HCF upgrading configuration is determined, a second phase of carbon-aware workload placement decides how to place AI workloads over this latency-constrained network.

Let C_v^{remain} denote the remaining compute capacity at DC v when task t arrives. For a task t originating from source s, the latency-feasible candidate set is

$$\mathcal{V}_{\text{feasible}}(t) = \{v \in V \mid RTT_{s,v} \leq \tau_t,\ C_v^{\text{remain}} \geq D_t\}.$$

Only DCs within this set are eligible for execution.

Among all feasible candidates, the strategy selects the destination with the minimum effective carbon intensity, defined as the product of carbon intensity and PUE:

$$v^* = \arg \min_{v \in \mathcal{V}_{\text{feasible}}(t)} \left(\text{PUE}_v \cdot \gamma_v \right).$$

If multiple candidates yield identical carbon cost, the one with lower end-to-end latency is selected.

By coupling a carbon-aware strategy with HCF-enabled latency reduction, the proposed framework enables latency-sensitive workloads to migrate toward geographically distant but cleaner execution sites that are infeasible under conventional fiber infrastructure.

IV. Illustrative Numerical Results

A. Simulation Setup

We evaluate the proposed framework using a discrete-event simulation that models geo-distributed DCs interconnected by optical backbone networks. Two representative backbone topologies are considered to capture different structural characteristics and upgrading scenarios.

Our primary evaluation is conducted on the COST239 topology, a widely used European optical backbone benchmark [17], comprising 11 nodes and 26 bidirectional links, whose link lengths are generally shorter due to its pan-European geographical scale. We additionally evaluate the framework on NSFNet [18], a U.S.-scale backbone topology with 14 nodes and 21 links. To enable a realistic evaluation on these topologies, we configure geographically differentiated data center and energy-related parameters based on real-world infrastructure and regional statistics. DC capacities are configured to reflect heterogeneous upgrading patterns observed in real-world infrastructures, ranging from hyperscale cloud facilities to regional and enterprise data centers [19]. Regional carbon intensity values are derived from the U.S. EPA eGRID 2022 database [20] and the European Environment Agency 2024 [21]. *PUE* values are set to 1.2–1.5, consistent with industry reports on data center energy efficiency [22].

The simulation processes 1,000 computing tasks generated according to a Poisson arrival process. Each task represents a geo-distributed workload that requires execution at a DC and generates inter–DC traffic. Tasks are characterized by compute demand and data-transfer volume, with data rates ranging from Mbps-level (LLM inference) to multi-Gbps (LLM training), while general computing workloads lie in between. Task sources follow a GDP-weighted distribution to approximate metropolitan-level demand heterogeneity. The workload mix includes latency-sensitive LLM inference tasks, general computing workloads, and compute-intensive LLM training tasks.

Driven by the rapid growth of AI services and emerging 6G visions, end-to-end latency requirements for distributed computing are expected to become increasingly stringent [23]. Accordingly, we model heterogeneous latency requirements ranging from 1–500 ms, with LLM inference exhibiting the strictest latency constraints, followed by general computing workloads, and LLM training tasks exhibiting the most relaxed latency tolerance.

The system adopts a two-layer optimization architecture. At the infrastructure layer, three HCF upgrading heuristics are evaluated:

Carbon-Driven HCF Upgrading(CDHU), our proposed approach, which prioritizes links whose latency reduction most effectively enables carbon-aware workload migration toward low-carbon data centers;

Latency-Driven HCF Upgrading(LDHU), which upgrades links that yield the largest absolute propagation latency reduction;

and **Random HCF Upgrading (RDHU)**, which selects links uniformly at random under the same budget constraint, serving as a baseline for comparison.

At the workload layer, two task scheduling strategies are considered: **Carbon-Aware Workload Placement (CAWP)** scheduling, our proposed approach, which assigns tasks to feasible data centers with the lowest effective carbon intensity, and **Load-Balancing Workload Placement(LBWP)** scheduling, which assigns tasks to feasible data centers with the largest remaining compute capacity.

For clarity, we denote a full configuration using the format *[Upgrading Strategy]–[Scheduling Strategy]*. Accordingly, the six evaluated combinations are: **CDHU–CAWP**, **CDHU–LBWP**, **LDHU–CAWP**, **LDHU–LBWP**, **RDHU–CAWP**, and **RDHU–LBWP**.

SMF and HCF are modeled with propagation speeds of 200,000 km/s and 300,000 km/s, respectively. Results are averaged over 10 runs to ensure statistical robustness.

B. Results on the COST239 Topology

1) Carbon Emissions and Reduction Gains: Fig. 4 reports the absolute carbon emissions, while Fig. 5 shows the corresponding relative reduction percentages under different HCF upgrading and scheduling strategies. Across all evaluated configurations, the CDHU–CAWP consistently delivers the strongest decarbonization performance.

A key observation is the strong budget efficiency of CDHU–CAWP on COST239. As shown in Fig. 5, CDHU–CAWP achieves a substantial reduction of **16.7%** with only 4,000 km of HCF upgrades. At the same budget, the LDHU–CAWP yields a negligible reduction of merely **0.4%**, indicating that minimizing physical delay alone is insufficient to unlock low-carbon execution opportunities under tight budget constraints.

The CDHU–CAWP strategy reaches a saturation point of approximately **19.4%** reduction at a budget of 6,000 km and maintains this level as the budget increases further. In contrast, LDHU–CAWP requires nearly the full 14,000 km budget to achieve comparable reduction levels. This gap highlights that carbon-driven upgrading is significantly more effective at translating limited HCF investments into emission reductions on dense backbone networks.

Across all upgrading strategies, LBWP scheduling consistently underperforms carbon-aware scheduling. For instance, CDHU–LBWP plateaus at around **14.3%** reduction, confirming that HCF upgrading must be coupled with carbon-aware workload placement to fully realize its environmental benefits.

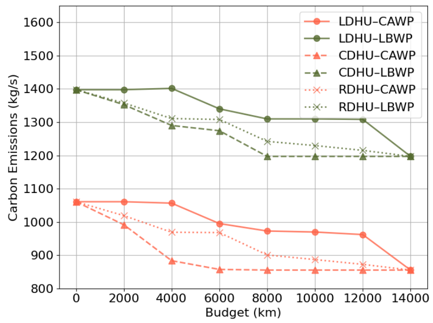

Fig. 4. Total carbon emissions versus HCF upgrade budget under different upgrading and scheduling strategies on COST239.

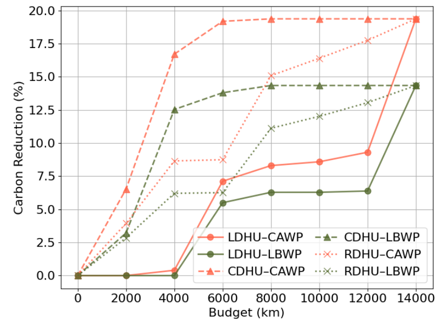

Fig. 5. Carbon reduction relative to the no-HCF baseline as a function of HCF upgrade budget on COST239.

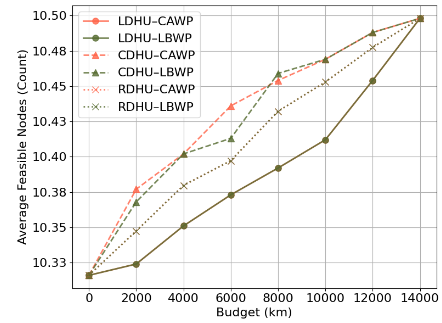

Fig. 6. Average number of latency-feasible nodes versus HCF upgrade budget on COST239.

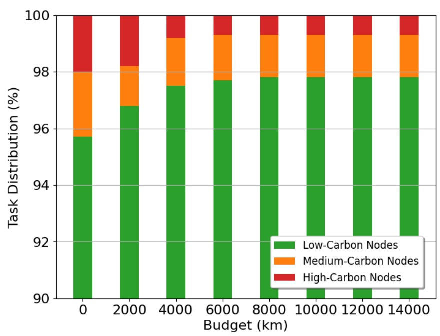

Fig. 7. Workload distribution across low-, medium-, and high-carbon nodes under CDHU–CAWP on COST239.

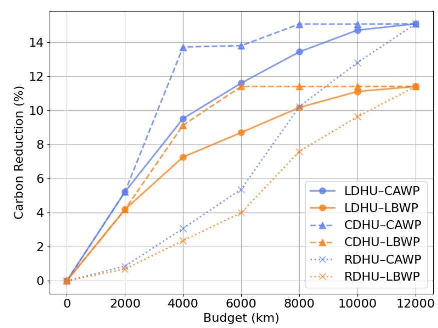

Fig. 8. Carbon reduction versus HCF upgrade budget under different strategies on NSFNet.

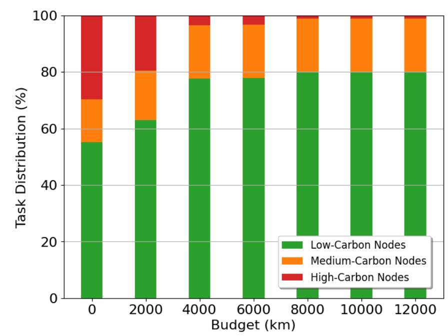

Fig. 9. Workload distribution across low-, medium-, and high-carbon nodes under CDHU–CAWP on NSFNet.

2) Expansion of Latency-Feasible Regions: Figure 6 shows the average number of latency-feasible DCs as the HCF upgrade budget increases. Owing to the dense connectivity of COST239, the baseline feasible region is already large, yet selective HCF upgrading still yields a consistent expansion.

For a given HCF upgrading strategy, the curves corresponding to CAWP and LBWP scheduling overlap. This is expected, as the set of latency-feasible nodes is determined solely by the physical network configuration and HCF upgrading, and is independent of the scheduling. As a result, LDHU–CAWP and LDHU–LBWP (as well as CDHU–CAWP and CDHU–LBWP) exhibit identical feasible-node counts, whereas the RDHU strategy yields non-overlapping data points due to stochastic link selection.

Under carbon-driven upgrading, the average number of feasible nodes increases from **10.32** at zero budget to **10.50** at the maximum budget. Although the absolute increase is modest, it indicates that a small number of geographically distant yet low-carbon nodes become newly reachable, which is sufficient to enable effective carbon-aware workload reallocation.

3) Migration Toward Low-Carbon Regions: Figure 7 illustrates how the expansion of the latency-feasible region translates into systematic workload migration under CDHU–CAWP. Due to space constraints, CDHU–CAWP is used here as a representative case. As HCF upgrading increases, workloads previously constrained by propagation latency are progressively migrated away from carbon-intensive regions.

Specifically, the share of tasks executed at low-carbon nodes increases to **97.8%**, while execution at high-carbon nodes drops from **2.0%** to **0.7%**. Medium-carbon nodes similarly see a decline from **2.3%** to **1.5%**. Although the absolute fraction of migrated tasks appears small, workloads relocated from high-carbon nodes yield a large reduction in total emissions due to their significantly higher carbon intensity. These results show that selective HCF upgrading relaxes residual latency constraints and removes the most carbon-intensive execution paths, achieving substantial emission reductions without compromising latency requirements.

C. Results on the NSFNet Topology

We further evaluate the proposed framework on the NSFNet topology to examine its robustness under a large-scale backbone. Compared to COST239, NSFNet spans a wider geographic area with longer inter-node distances, leading to more stringent end-to-end latency constraints.

Figure 8 shows the carbon reduction achieved under different upgrading and scheduling strategies. Across all budgets, the CDHU–CAWP combination consistently achieves the largest emission reduction, reaching up to approximately **15%**. LDHU and RDHU upgradings yield noticeably smaller gains, especially under limited budgets, underscoring the importance of carbon-aware infrastructure planning.

Figure 9 illustrates the corresponding workload migration behavior. With increasing HCF upgrading, CDHU–CAWP enables workloads to progressively shift from medium- and high-carbon data centers toward lower-carbon regions. However, due to the large geographic scale of NSFNet, some low-carbon data centers remain unreachable within strict latency budgets even after selective HCF upgrades, which limits the overall migration extent and caps the achievable carbon reduction. Other performance metrics, including latency-feasible region expansion and scheduling behavior, exhibit trends consistent with those observed on COST239 and are omitted for brevity.

This paper is organized as follows. Section II presents the network and workload model. Section III describes the proposed methodology. Section IV provides the illustrative numerical results. Finally, Section V concludes the paper.

V. Conclusion

This paper studies how low-latency optical networks using HCF can enable carbon-aware geo-distributed computing. We demonstrated that selective HCF upgrading reshapes end-to-end latency, expanding the set of DCs reachable under strict latency constraints. Rather than directly reducing network energy use, HCF acts as an enabler, allowing latency-sensitive workloads to run at geographically distant, lower-carbon sites. To leverage this effect, we proposed a cross-layer optimization framework that jointly considers selective HCF upgrading at the optical network layer and carbon-aware task scheduling at the compute layer. This framework explicitly couples latency shaping in the optical network with workload placement decisions across geo-distributed data centers, allowing infrastructure upgrades to translate into carbon benefits. Extensive simulations on two backbone topologies showed that the proposed approach can reduce carbon emissions by up to nearly 20% compared with conventional SMF networks, highlighting the importance of co-design between optical networking and geo-distributed computing for sustainable computing. As current deployments of HCFs are happening in urban/regional settings (as opposed to the backbone segment considered in our results) due to several reasons (cost/availability of HCFs, DC availability zones, very strict latency requirements of some rising AI applications), future work will extend our framework to regional-scale topologies to capture the carbon-reduction potential of early HCF upgrading.

References

[1] C. Koronen, M. Åhman, and L. J. Nilsson, "Data centres in future european energy systems—energy efficiency, integration and policy," *Energy efficiency*, vol. 13, no. 1, pp. 129–144, 2020.

[2] "Recommendation itu-t l.1470: Greenhouse gas emissions trajectories for the information and communication technology sector compatible with the unfccc paris agreement," https://www.itu.int/rec/T-REC-L.1470, 2020, establishes ICT GHG emission trajectories aligned with Paris Agreement, including data centers and networks.

[3] J. Yang, Z. Saad, J. Wu, X. Niu, H. Leung, and S. Drew, "A survey on task scheduling in carbon-aware container orchestration," *arXiv preprint arXiv:2508.05949*, 2025.

[4] A. Gupta, U. Mandal, P. Chowdhury, M. Tornatore, and B. Mukherjee, "Cost-efficient live vm migration based on varying electricity cost in optical cloud networks," *Photonic Network Communications*, vol. 30, no. 3, pp. 376–386, 2015.

[5] S. Rahman, A. Gupta, M. Tornatore, and B. Mukherjee, "Dynamic workload migration over backbone network to minimize data center electricity cost," *IEEE Transactions on Green Communications and Networking*, vol. 2, no. 2, pp. 570–579, 2017.

[6] Z. Liu, M. Lin, A. Wierman, S. H. Low, and L. L. Andrew, "Greening geographical load balancing," *ACM SIGMETRICS Performance Evaluation Review*, vol. 39, no. 1, pp. 193–204, 2011.

[7] A. Radovanović, R. Koningstein, I. Schneider, B. Chen, A. Duarte, B. Roy, D. Xiao, M. Haridasan, P. Hung, N. Care *et al.*, "Carbon-aware computing for datacenters," *IEEE Transactions on Power Systems*, vol. 38, no. 2, pp. 1270–1280, 2022.

[8] Y. G. Kim, U. Gupta, A. McCrabb, Y. Son, V. Bertacco, D. Brooks, and C.-J. Wu, "Greenscale: Carbon-aware systems for edge computing," *arXiv preprint arXiv:2304.00404*, 2023.

[9] G. S. Sticca, M. Ibrahimi, F. Musumeci, N. Di Cicco, and M. Tornatore, "Hollow-core fibers for latency-constrained and low-cost edge data center networks," *IEEE Transactions on Network and Service Management*, 2025.

[10] P. Poggiolini and F. Poletti, "Opportunities and challenges for long-distance transmission in hollow-core fibres," *Journal of Lightwave Technology*, vol. 40, no. 6, pp. 1605–1616, 2022.

[11] G. S. Sticca, M. Ibrahimi, N. Di Cicco, F. Musumeci, and M. Tornatore, "Hollow core fiber as a long-term solution for capacity scaling in optical networks," in *2025 Optical Fiber Communications Conference and Exhibition (OFC)*. IEEE, 2025, pp. 1–3.

[12] M. Ibrahimi, G. S. Sticca, F. Musumeci, and M. Tornatore, "On the benefits of hollow-core fiber in next-generation optical networks," in *2025 International Conference on Optical Network Design and Modeling (ONDM)*. IEEE, 2025, pp. 1–5.

[13] K. Borzycki and T. Osuch, "Hollow-core optical fibers for telecommunications and data transmission," *Applied Sciences*, vol. 13, no. 19, p. 10699, 2023.

[14] G. S. Sticca, M. Ibrahimi, N. Di Cicco, F. Musumeci, and M. Tornatore, "Hollow-core-fiber placement in latency-constrained metro networks with edgedcs," in *2024 Optical Fiber Communications Conference and Exhibition (OFC)*. IEEE, 2024, pp. 1–3.

[15] M. Abarinova, "Microsoft preps cloud for ai with corning, heraeus hollow core fiber deal," *Fierce Network*, Sep. 2025, reports on Microsoft's strategic manufacturing partnerships for scaling hollow core fiber production with Corning and Heraeus Covantics globally. [Online]. Available: https://www.fierce-network.com/broadband/microsoft-strikes-new-deal-corning-boost-its-hollow-core-fiber

[16] D. Goovaerts, "Aws is using hollow core fiber to go the distance," *Fierce Network*, Dec. 2025, discusses AWS's development and deployment of hollow core fiber for low-latency, high-bandwidth connectivity across its global cloud infrastructure. [Online]. Available: https://www.fierce-network.com/cloud/aws-using-hollow-core-fiber-go-distance

[17] M. O'Mahony, "Results from the cost 239 project. ultra-high capacity optical transmission networks," in *Proceedings of European Conference on Optical Communication*, vol. 2. IEEE, 1996, pp. 11–18.

[18] K. C. Claffy, G. C. Polyzos, and H.-W. Braun, "Traffic characteristics of the t1 nsfnet backbone," in *IEEE INFOCOM'93 The Conference on Computer Communications, Proceedings*. IEEE, 1993, pp. 885–892.

[19] G. Kamiya and V. C. Coroamă, "Data centre energy use: Critical review of models and results," *IEA 4E TCP Efficient, Demand Flexible Networked Appliances (EDNA)*, 2025.

[20] U.S. Environmental Protection Agency, "Emissions & Generation Resource Integrated Database (eGRID) 2022," https://www.epa.gov/egrid, 2024, includes eGRID2022 data summary tables and technical guide for 2022 emissions and generation statistics.

[21] E. E. Agency, "Greenhouse gas emission intensity of electricity generation in europe," 2021.

[22] A. Shehabi, A. Newkirk, S. J. Smith, A. Hubbard, N. Lei, M. A. B. Siddik, B. Holecek, J. Koomey, E. Masanet, and D. Sartor, "2024 united states data center energy usage report," 2024.

[23] J. Xiong, Q. Zhang, P. Medagliani, M. Ferrero, X. Liu, M. Lian, N. Di Cicco, B. Zhao, M. Ibrahimi, F. Musumeci *et al.*, "Scale-ccl: A scalable collective communication library for wide-area distributed training," in *Proceedings of the 1st Workshop on Inter-networking challenges for AI*, 2025, pp. 21–27.