Hollow-Core-Fiber Placement in Latency-Constrained Metro Networks with edgeDCs

Giovanni Simone Sticca*, Memedhe Ibrahimi, Nicola Di Cicco,
Francesco Musumeci, and Massimo Tornatore
Department of Electronics, Information and Bioengineering, Via Giuseppe Ponzio, 34, 20133 Milano, Italy
**giovannisimone.sticca@polimi.it*

Abstract: We investigate the optimal placement of Hollow-Core Fibers (HCF) in latency-constrained metro networks with edgeDCs, performing physical-layer validation. Upgrading 24% of links to HCF reduces edgeDCs number by 29% compared to a network without HCFs. © 2024 The Author(s)

1. Introduction

Many next-generation applications in fields such as e-health, smart cities, autonomous driving, and extended reality, require optical metro networks supporting extremely high throughput and ultra-low latency. The only viable technical direction to satisfy such hard latency constraints is deploying distributed edge Data Centers (edgeDCs) physically closer to clients [1]. As edgeDCs interwork among them, strict latency requirements not only apply to the communication between edgeDCs and clients, but also extend to the communication across edgeDCs, e.g., for network functions chain deployment, synchronization between different edgeDCs, or distributed Machine Learning training and inference. Due to the inherent limits of the propagation speed in standard single-mode fibers (SSMFs), latency in metro networks can only be substantially reduced by deploying a large number of costly edgeDCs. In this context, Hollow Core Fibers (HCFs) become a powerful instrument to limit the number of edgeDCs. Compared to conventional SSMFs, HCFs provide a lower propagation latency, have four orders of magnitude lower non-linear effects, have lower attenuation, and have a wider usable bandwidth [2]. At the current stage of HCF technology development, the low interconnection losses between SSMFs and HCFs make it feasible to upgrade current optical networks by adding HCFs [3]. As an illustrative example, Fig. 1 shows a 100 km SSMF link that gets upgraded using a parallel HCF link. Upgrade to HCF enables savings of 33% in latency (τ) and an SNR improvement of 2.6 dB, compared to SSMF.

In a similar line of research, past literature investigated augmenting fiber networks with free-space microwave links for low-latency communications [4]. However, microwave links provide significantly lower peak throughput (few Gbps) than fiber links and are sensitive to atmospheric impairments. We argue that HCFs can provide the best of both worlds: high throughput *and* free-space speed-of-light communication.

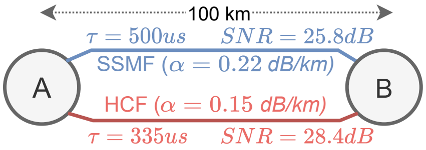

Fig. 1: HCFs vs SSMFs in a 100km link.

In this study, we investigate how to selectively place HCFs in a latency-constrained metro network to minimize the number of deployed edgeDCs. While the problem of edgeDC placement and the advantages of HCFs have been studied in the past, these problems have been investigated separately. For example, Ref. [5] investigates the problem of edgeDC and content placement while minimizing delay and bandwidth consumption, and Ref. [2] investigates the impact of HCFs for long-distance optical systems with emphasis on the network throughput. However, to the best of our knowledge, no prior work has investigated *network planning* with HCFs, specifically optimizing HCF upgrades and edgeDC placement for low-latency communications in metro networks.

Technical contributions: We develop a Mixed Integer Linear Programming (MILP) model for optimal placement of a given number of HCF links to minimize the number of edgeDCs. We test our solution by performing Routing, Modulation format, and Spectrum Assignment (RMSA), and validate it in terms of physical-layer signal quality. We quantify the gains of our solution in terms of network throughput, number of transponders (TXPs), and the lightpaths' Signal-to-Noise Ratio (SNR) achieved after deploying HCF links.

2. *PARALLEL*: Physical-Layer-Aware Hollow-Core-Assisted EdgeDC Placement

To solve this problem, ***P****hysic****a****l-****La****y****er****-****A****ware Ho****ll****ow-Core-Assisted* ***E****dgeDC p****l****acement (PARALLEL)* includes two steps: **1) HCF-DC-MILP**: a MILP model that optimizes the placement of edgeDCs and HCFs; **2) RMSA-PLV**: a post-processing procedure that takes in input the solution of HCF-DC-MILP, performs the RMSA of the traffic, and Physical Layer Validation (PLV) in terms of feasibility of the lightpaths' SNR.

2.1. *Hollow-Core Fiber and EdgeDC Placement Problem (HCF-DC-MILP)*

We formulate the problem of HCF and edgeDC placement as: **Given** a network topology, an HCF budget (i.e., the maximum number of links to be upgraded to HCF), an optical latency threshold T between client nodes and edgeDCs, and a maximum client capacity for each edgeDC, **decide** edgeDCs and HCFs placement and assignment of each client to an edgeDC, **constrained by** an HCF upgrade budget, edgeDC capacity, latency threshold, and client-to-edgeDC assignment, with the **objectives** of minimizing the number of deployed edgeDCs (primary) and the maximum latency between edgeDCs (secondary).

The objective of HCF-DC-MILP is to minimize $z = \alpha\left(\frac{\sum_{i\in N} d_i}{|N|}\right) + \beta\,\frac{t}{M}$, which expresses the weighted sum of a) the total number of edgeDCs (normalized to the number of nodes), and b) the maximum inter-edgeDC latency (normalized to the maximum latency in the topology), subject to the following constraints:

Parameters	Description	Sets and variables	Description
B	HCF upgrade budget	N	set of nodes
τ	propagation latency in SSMF ($\mu s/km$)	N_c	set of client nodes
$\Delta\tau$	latency reduction	L	set of links
G_l	length of link l in km	S_{ij}	set of links in shortest path from node i to node j
T_E	EDFA latency		
ε_l	number of EDFAs in link l	$h_l \in \{0,1\}$	*binary*, 1 if link l is upgraded with HCF
T_N	node latency	$d_i \in \{0,1\}$	*binary*, 1 if an edgeDC is in node i
T	latency threshold	$x_{ij} \in \{0,1\}$	*binary*, 1 if client i is assigned to edgeDC j
K	edgeDC capacity	$t \geq 0$	*continuous*, max. inter-edgeDC latency

Table 1: Parameters, sets, and variables of HCF-DC-MILP

$$\sum_{l\in L} h_l \leq B \qquad (1) \qquad \sum_{j\in N} x_{ij} = \begin{cases} 1 & \text{if } i \in N_c \\ 0 & \text{if } i \in N \setminus N_c \end{cases} \qquad (2) \qquad x_{ij} \leq d_j \quad \forall ij \in N \qquad (3)$$

$$\Big\{\sum_{k\in S_{ij}} [(\tau - \Delta\tau \cdot h_k)\cdot G_k + T_E \cdot \varepsilon_k] + T_N \cdot |S_{ij}|\Big\} \cdot x_{ij} \leq T \quad \forall ij \in N \qquad (4) \qquad \sum_{j\in N} x_{ji} \leq K \quad \forall i \in N \qquad (5)$$

$$t \geq \Big\{\sum_{k\in S_{ij}} [(\tau - \Delta\tau \cdot h_k)\cdot G_k + T_E \cdot \varepsilon_k] + T_N \cdot |S_{ij}|\Big\} \cdot (d_i + d_j - 1) \quad \forall ij \in N \qquad (6)$$

Constraint (1) ensures that the number of links upgraded with HCF does not exceed a maximum budget B. Constraints (2) and (3) assign each client node to one edgeDC. Constraint (4) ensures that the latency for each client-to-edgeDC connection does not exceed the maximum latency threshold. The total latency includes propagation latency and additional latency introduced by optical nodes and EDFAs (the latter two are significantly less than the first). Constraint (5) guarantees that the number of client nodes assigned to an edgeDC does not exceed the maximum edgeDC capacity K. Constraint (6) defines the maximum inter-edgeDC latency.

2.2. *RMSA and Physical Layer Validation (RMSA-PLV)*

Given the solution of HCF-DC-MILP, i.e., the optimal placement of HCFs and edgeDCs, and a traffic matrix, we perform RMSA and PLV. Starting from an initial traffic matrix, we consider an incremental traffic scenario with a 30% per-step traffic increase. For each step, we perform RMSA and we validate it in terms of physical layer impairments, i.e., lightpaths' SNR. The simulation terminates when 1% blocking rate is met.

We consider an elastic optical network with frequency slots of 12.5 GHz and a total capacity of 10.1 THz per fiber (5.9 THz for C-band and 4.2 THz for L-band). We consider three types of TXPs as in [6] and modulation formats ranging from QPSK to 64-QAM. We consider k-shortest-path routing (k=3) with minimal loss [7] and first-fit spectrum allocation. For latency-constrained demands, the link weights are their propagation latency. Otherwise, links are weighted by their length, and ties are broken in favor of SSMFs to avoid overloading HCF links with non-latency-constrained traffic. For each step of traffic increase, we validate the RMSA solution in terms of physical layer impairments. We consider SSMF and HCF operating in (C+L) band. For SSMF, we adopt the physical-layer model in [8], while for HCF we adopt the model in [2]. We assume EDFA amplification and that links operate in ASE loading. A lightpath is considered feasible when its SNR is above a threshold determined by data rate and modulation format, plus an additional 2dB system margin. At the end of the simulation , RMSA-PLV outputs network throughput (i.e., served traffic at the final step), number of TXPs, and lightpaths' SNR.

3. Illustrative numerical results

We consider a realistic metro network of 52 nodes and 144 links (*Metro-52*) [7]. We assume that the latency thresholds of the connection requests are generated from a Gaussian distribution with mean values in the range [100-400] *us* and standard deviation equal to 20 *us*. The latency threshold T that is imposed as a constraint in the MILP is chosen to provide a probability of 99.5% that the latency requirement for the connections is met.

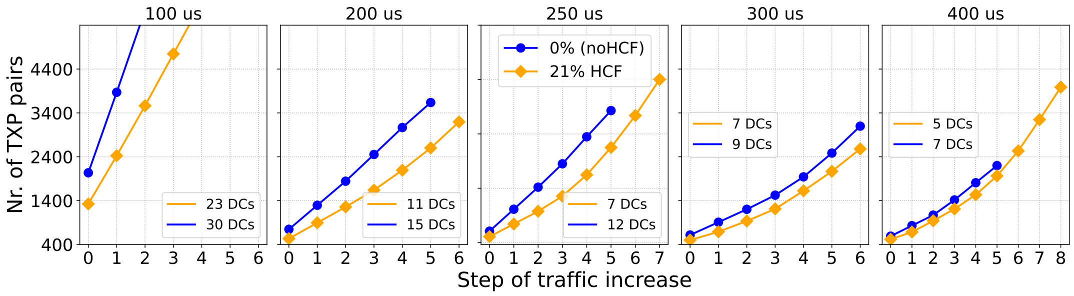

Fig. 3: Number of transponder (TXP) pairs in each step of traffic increase for 0% (noHCF) and 21% HCF budget, for varying latency thresholds between (100-400) *us*.

We consider client-to-edgeDC, client-to-core, and inter-edgeDC traffic with a 40%, 40%, and 20% distribution, respectively. The initial traffic is 200 Tbps, and we assume that client-to-edgeDC traffic is latency-constrained. We consider an edgeDC capacity of 10, i.e., each edgeDC can serve 10 clients at most. The MILP objective function is weighted with $\alpha = 0.9$ and $\beta = 0.1$, which prioritizes minimizing the number of edgeDCs over the maximum inter-edgeDC latency. We compare HCF placement using *PARALLEL* to a baseline *noHCF* scenario where, instead of HCF links, we upgrade the network with additional SSMF links.

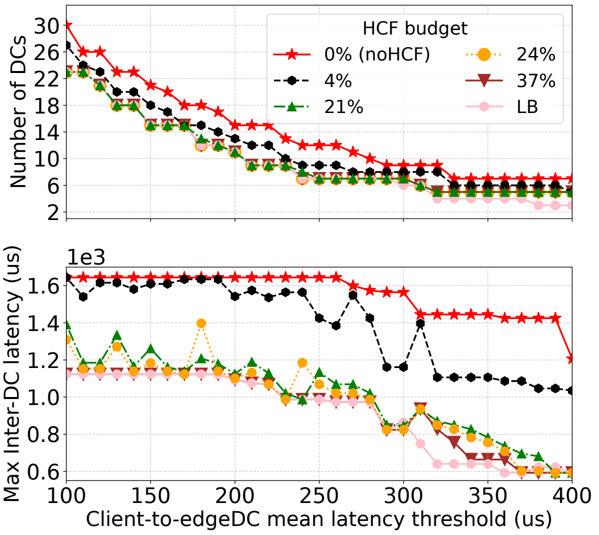

Fig. 2: Number of edgeDCs (top) and max inter-DC latency (bottom), varying the HCF upgrade budget and the client-to-edgeDC mean latency threshold.

Figure 2 shows the number of edgeDCs and the maximum inter-DC latency for varying HCF upgrade budgets and latency thresholds. The HCF budget indicates the number of links upgraded to HCF normalized to the total number of network links. We also report the lower bound (LB), obtained by setting 100% HCF budget and removing the edgeDCs capacity constraints. We observe that, by upgrading only 4% of the links to HCF, we reduce the number of edgeDCs by 16% on average compared to *noHCF*, while increasing the HCF budget to 24% the savings increase up to 29%. HCF upgrade budgets beyond 24% do not bring any further reduction in the number of edgeDCs. In the range [100, 290] *us* we hit the lower bound. This indicates that, in this range, our results are independent of the assumed edgeDC capacity. Finally, increasing the HCF upgrade budget up to 37% reduces the maximum inter-edgeDC latency up to 41% compared to *noHCF*. Further increases do not bring any improvements.

We now illustrate the numerical results for **RMSA-PLV** for an HCF upgrade budget of 21% (similar findings apply to other budget values, not reported due to space constraints). Figure 3 shows the number of TXP pairs required to serve the traffic in the case of *noHCF* and 21% HCF budget, as the network load increases. We report the number of edgeDCs for each scenario of latency threshold, as found in HCF-DC-MILP. By upgrading 21% of the links to HCF we save 24% in the number of TXPs, on average across all latency thresholds (38% for 100 *us*, 29% for 200 *us*, 24% for 250 *us*, 17% for 300 *us*, and 11% for 400 *us*), for the same amount of traffic served. This improvement results from two factors: i) HCFs improve the ligthpaths' SNR, enabling the use of higher-order modulation formats (e.g., average SNR increase is 2.3dB in C-band and 3.3dB in L-band), and ii) the reduction in edgeDCs number leads to a decrease in the number of inter-edgeDC connections which, in turn, reduces the number of TXPs. Moreover, we observe that 21% HCF enables a throughput increase of up to 890 Tbps compared to *noHCF*. As a result, in the case of 250 *us* and 400 *us*, 21% HCF allows more traffic increase steps than *noHCF*. In both *noHCF* and 21% HCF, the number of lightpaths violating the latency threshold is less than 1%.

In conclusion, we have shown that leveraging HCFs allows a 29% average reduction of edgeDCs in a latency-constrained metro network. Moreover, we achieve a remarkable 24% average reduction in the number of TXPs.

Acknowledgement: This work was supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on “Telecommunications of the Future” (PE00000001 - program “RESTART”).

References

1. X. Lin *et al.*, ICOCN, 2023
2. P. Poggiolini *et al.*, JLT, 40(6), pp. 1605-1616, 2022
3. D. Suslov *et al.*, Scientific Reports, 11(1), 2021
4. D. Bhattacherjee *et al.*, NSDI, pp. 1115-1133, 2022
5. Z. Liu *et al.*, ACP, 2022
6. O. Karandin *et al.*, ONDM, 2021
7. M. Ibrahimi *et al.*, JLT, 38(12), pp. 3221-3228, 2020
8. D. Semrau *et al.*, JLT, 37(9), pp. 1924-1936, 2019