

Brief Announcement: Minimizing Congestion in Hybrid Demand-Aware Network Topologies

Wenkai Dai  

Faculty of Computer Science, University of Vienna, Austria

Michael Dinitz  

Computer Science Department, Johns Hopkins University, USA

Klaus-Tycho Foerster  

Computer Science Department, Technical University of Dortmund, Germany

Stefan Schmid  

TU Berlin, Germany

Faculty of Computer Science, University of Vienna, Austria

Abstract

Emerging reconfigurable optical communication technologies enable demand-aware networks: networks whose static topology can be enhanced with demand-aware links optimized towards the traffic pattern the network serves. This paper studies the algorithmic problem of how to jointly optimize the topology and the routing in such demand-aware networks, to minimize congestion. We investigate this problem along two dimensions: (1) whether flows are splittable or unsplittable, and (2) whether routing on the hybrid topology is segregated or not, i.e., whether or not flows either have to use exclusively either the static network or the demand-aware connections. For splittable and segregated routing, we show that the problem is 2-approximable in general, but APX-hard even for uniform demands induced by a bipartite demand graph. For unsplittable and segregated routing, we show an upper bound of $O(\log m / \log \log m)$ and a lower bound of $\Omega(\log m / \log \log m)$ for polynomial-time approximation algorithms, where m is the number of static links. Under splittable (resp., unsplittable) and non-segregated routing, even for demands of a single source (resp., destination), the problem cannot be approximated better than $\Omega(c_{\max}/c_{\min})$ unless $P=NP$, where c_{\max} (resp., c_{\min}) denotes the maximum (resp., minimum) capacity. It is still NP-hard for uniform capacities, but can be solved efficiently for a single commodity and uniform capacities.

2012 ACM Subject Classification Networks → Network architectures; Theory of computation → Design and analysis of algorithms

Keywords and phrases Congestion, Reconfigurable Networks, Algorithms, Complexity

Digital Object Identifier 10.4230/LIPIcs.DISC.2022.13

Category Brief Announcement

Funding Research supported by the European Research Council (ERC), grant agreement No. 864228 (AdjustNet) Horizon 2020, 2020-2025, and NSF Award CCF-1909111.

1 Introduction

Emerging demand-aware networks, whose topologies are typically hybrid, in that a static (and demand-oblivious) network is enhanced with reconfigurable (and demand-aware) links, introduce unprecedented flexibility in adapting the network topology towards current traffic demands. In such hybrid networks, the reconfigurable links are usually enabled by optical circuit switches [1, 8, 13], and particularly, each optical circuit switch provides reconfigurable links by establishing connections between pairs of its ports, i.e., a matching.

Extensive past works studied the question of how to jointly optimize topology and routing of such hybrid (reconfigurable) networks [17] for different networking performance



© W. Dai, M. Dinitz, K.-T. Foerster, and S. Schmid;
licensed under Creative Commons License CC-BY 4.0
36th International Symposium on Distributed Computing (DISC 2022).

Editor: Christian Scheideler; Article No. 13; pp. 13:1–13:3

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Approximation Upper & Lower Bounds (Complexity)	Splittable Flow	Segregated Routing	Restrictions On Demands
2-approximation	yes	yes	
APX-complete	yes	yes	uniform and bipartite demands
$O(\log m / \log \log m)$ -approximation	no	yes	
Lower Bound: $\Omega(\log m / \log \log m)$	no	both	
$(2 \cdot c_{\max} / c_{\min})$ -approximation	yes	no	single source (resp., dest.)
Lower Bound: $\Omega(c_{\max} / c_{\min})$	both	no	single source (resp., dest.)

■ **Table 1** Summary of our approximation upper and lower bounds on the MCHN problem (Def. 1). metrics, e.g., latency [11], throughput [4, 7], routing length [14, 15, 16], flow times [3] etc. Interestingly, *min-congestion*, a most central performance metric in traditional networks, is still not well-understood in hybrid networks. Avin et al. [6] and Pacut et al. [12] study optimal *bounded-degree* topology designs to minimize both the route length and the congestion. Dai et al. [2] worked on the same network model as us, showing that the problem is already NP-hard for *splittable* (resp., *unsplittable*) and *segregated* (resp., *non-segregated*) routing models when the static network is a tree of height at least two, but tractable for static networks of star topologies. Zheng et al. [10] introduced a greedy-based heuristic algorithm for our *segregated* model but on specific topologies of datacenters. However, not much more is known w.r.t. corresponding approximation bounds, which motivates our study, summarized in Table 1.

2 Model

Network Model. We consider a *hybrid* network [5, 9] $N = (V, E, \mathcal{E}, c)$, where a *static* network (V, E) is represented by an bidirected (simple) graph of nodes V , any two distinct nodes $v_i, v_j \in V$ imply a possible *reconfigurable link* denoted by a bidirected edge $\{i, j\}$ in \mathcal{E} , and a function $c : \vec{E} \cup \vec{\mathcal{E}} \mapsto \mathbb{R}_{\geq 0}$ defines *capacities* for both directions of each bidirected link in $E \cup \mathcal{E}$ with the *maximum* (resp., *minimum*) *capacity* denoted by c_{\max} (resp., c_{\min}). The hybrid network N must decide a matching $M \subseteq \mathcal{E}$ to obtain an enhanced graph $N(M) = (V, E \cup M, c)$, which determines the actual topology of the communicating network.

Traffic Demands. A certain communication pattern (*demands*) on nodes V is represented by a matrix $D := (d_{i,j})_{|V| \times |V|}$, where an entry $d_{i,j} \in \mathbb{R}_{\geq 0}$ denotes the traffic load (frequency) or a demand from the node $v_i \in V$ to the node $v_j \in V$.

Routing Models. The *unsplittable* routing requires that flows of each demand must be sent along a single (directed) path, otherwise the routing is called *splittable*. For a hybrid network, *segregated* routing requires that each demand $d_{i,j}$ is either sent on the reconfigurable link $\{i, j\}$, if it exists, or purely on the static network, otherwise it is *unsegregated* routing. Hence, we consider *four different routing models*: *Unsplittable & Segregated (US)*, *Unsplittable & Non-segregated (UN)*, *Splittable & Segregated (SS)*, and *Splittable & Non-segregated (SN)*.

► **Definition 1** (Min-Congestion Hybrid Network Problem (MCHN)). *Given a hybrid network $N = (V, E, \mathcal{E}, c)$, a routing model $\tau \in \{US, UN, SS, SN\}$, and a demand matrix D , find a matching $M \subseteq \mathcal{E}$, s.t., the congestion λ , i.e., the maximum load on $\vec{E} \cup \vec{M}$, to serve D in $N(M)$ is minimized.*

3 Our Contributions

We initiate the study of approximation algorithms for minimizing congestion in hybrid demand-aware networks (for a given matrix of demands). Our results include an overview of approximation results and complexity characterizations in general settings. We also provide a fine-grained algorithmic analysis for restricted cases:

Segregated Routing: We can give a mixed-integer programming formulation for segregated and un-/splittable flow models whose LP relaxation can be solved efficiently. For splittable flows, we provide a 2-approximation algorithm by a novel deterministic rounding approach, and also prove APX-hardness even if demands are uniform and bipartite. However, we also show that the problem becomes tractable for demands with a single source (resp., dest.). For unsplittable flows, we show that the hybrid network problem cannot be approximated better than the min-congestion multi-commodity unsplittable flow problem (MCMF) [18], but any ρ -approximation algorithm based on rounding techniques for the MCMF problem can be utilized to give a 2ρ -approximation. This implies an approximability of $\Theta(\log m / \log \log m)$ for segregated and unsplittable routing, where $m = |E|$.

Non-Segregated Routing: Under the splittable (resp., unsplittable) flow model, even for demands of a single source (resp., destination), the problem cannot be approximated better than $\Omega(c_{\max}/c_{\min})$ unless $P=NP$, but still $(2 \cdot c_{\max}/c_{\min})$ -approximable for the splittable flow, where c_{\max} (resp., c_{\min}) denotes the maximum (resp., minimum) capacity on all links, and it still remains NP-hard for *uniform capacities*, i.e., $c : \vec{E} \cup \vec{E} \mapsto \{a\}$ for $a \in \mathbb{R}_{>0}$. However, the problem with uniform capacities becomes efficiently solvable for demands of a single commodity under un-/splittable flow.

References

- 1 S. Aleksic. The future of optical interconnects for data centers: A review of technology trends. In *2017 14th International Conference on Telecommunications (ConTEL)*, June 2017.
- 2 W. Dai et al. Load-optimization in reconfigurable networks: Algorithms and complexity of flow routing. *SIGMETRICS Perform. Evaluation Rev.*, 48(3), 2020.
- 3 M. Dinitz and B. Moseley. Scheduling for weighted flow and completion times in reconfigurable networks. In *INFOCOM*. IEEE, 2020.
- 4 A. Singla et al. High throughput data center topology design. In *NSDI*. USENIX, 2014.
- 5 B. Venkatakrisnan et al. Costly circuits, submodular schedules and approximate carathéodory theorems. In *SIGMETRICS*, 2016.
- 6 C. Avin et al. Demand-aware network design with minimal congestion and route lengths. In *INFOCOM*. IEEE, 2019.
- 7 D. Nikhil et al. Stable matching algorithm for an agile reconfigurable data center interconnect (MSR-TR-2016-1140). Technical report, Microsoft Research, June 2016.
- 8 G. Wang et al. c-through: part-time optics in data centers. In *SIGCOMM*. ACM, 2010.
- 9 H. Liu et al. Scheduling techniques for hybrid circuit/packet networks. In *CoNEXT*. ACM, 2015.
- 10 J. Zheng et al. Dynamic load balancing in hybrid switching data center networks with converters. In *ICPP*. ACM, 2019.
- 11 M. Ghobadi et al. Projector: Agile reconfigurable data center interconnect. In *SIGCOMM*. ACM, 2016.
- 12 M. Pacut et al. Improved scalability of demand-aware datacenter topologies with minimal route lengths and congestion. *Perform. Evaluation*, 152, 2021.
- 13 N. Farrington et al. Helios: a hybrid electrical/optical switch architecture for modular data centers. In *SIGCOMM*. ACM, 2010.
- 14 T. Fenz et al. Efficient non-segregated routing for reconfigurable demand-aware networks. *Comput. Commun.*, 164, 2020.
- 15 K.-T. Foerster et al. On the complexity of non-segregated routing in reconfigurable data center architectures. *Comput. Commun. Rev.*, 49(2):2–8, 2019.
- 16 K.-T. Foerster and S. Schmid. Survey of reconfigurable data center networks: Enablers, algorithms, complexity. *SIGACT News*, 50(2), 2019.
- 17 M. N. Hall et al. A survey of reconfigurable optical networks. *Opt. Switch. Netw.*, 41, 2021.
- 18 Vijay V. Vazirani. *Approximation algorithms*. Springer, 2001.