

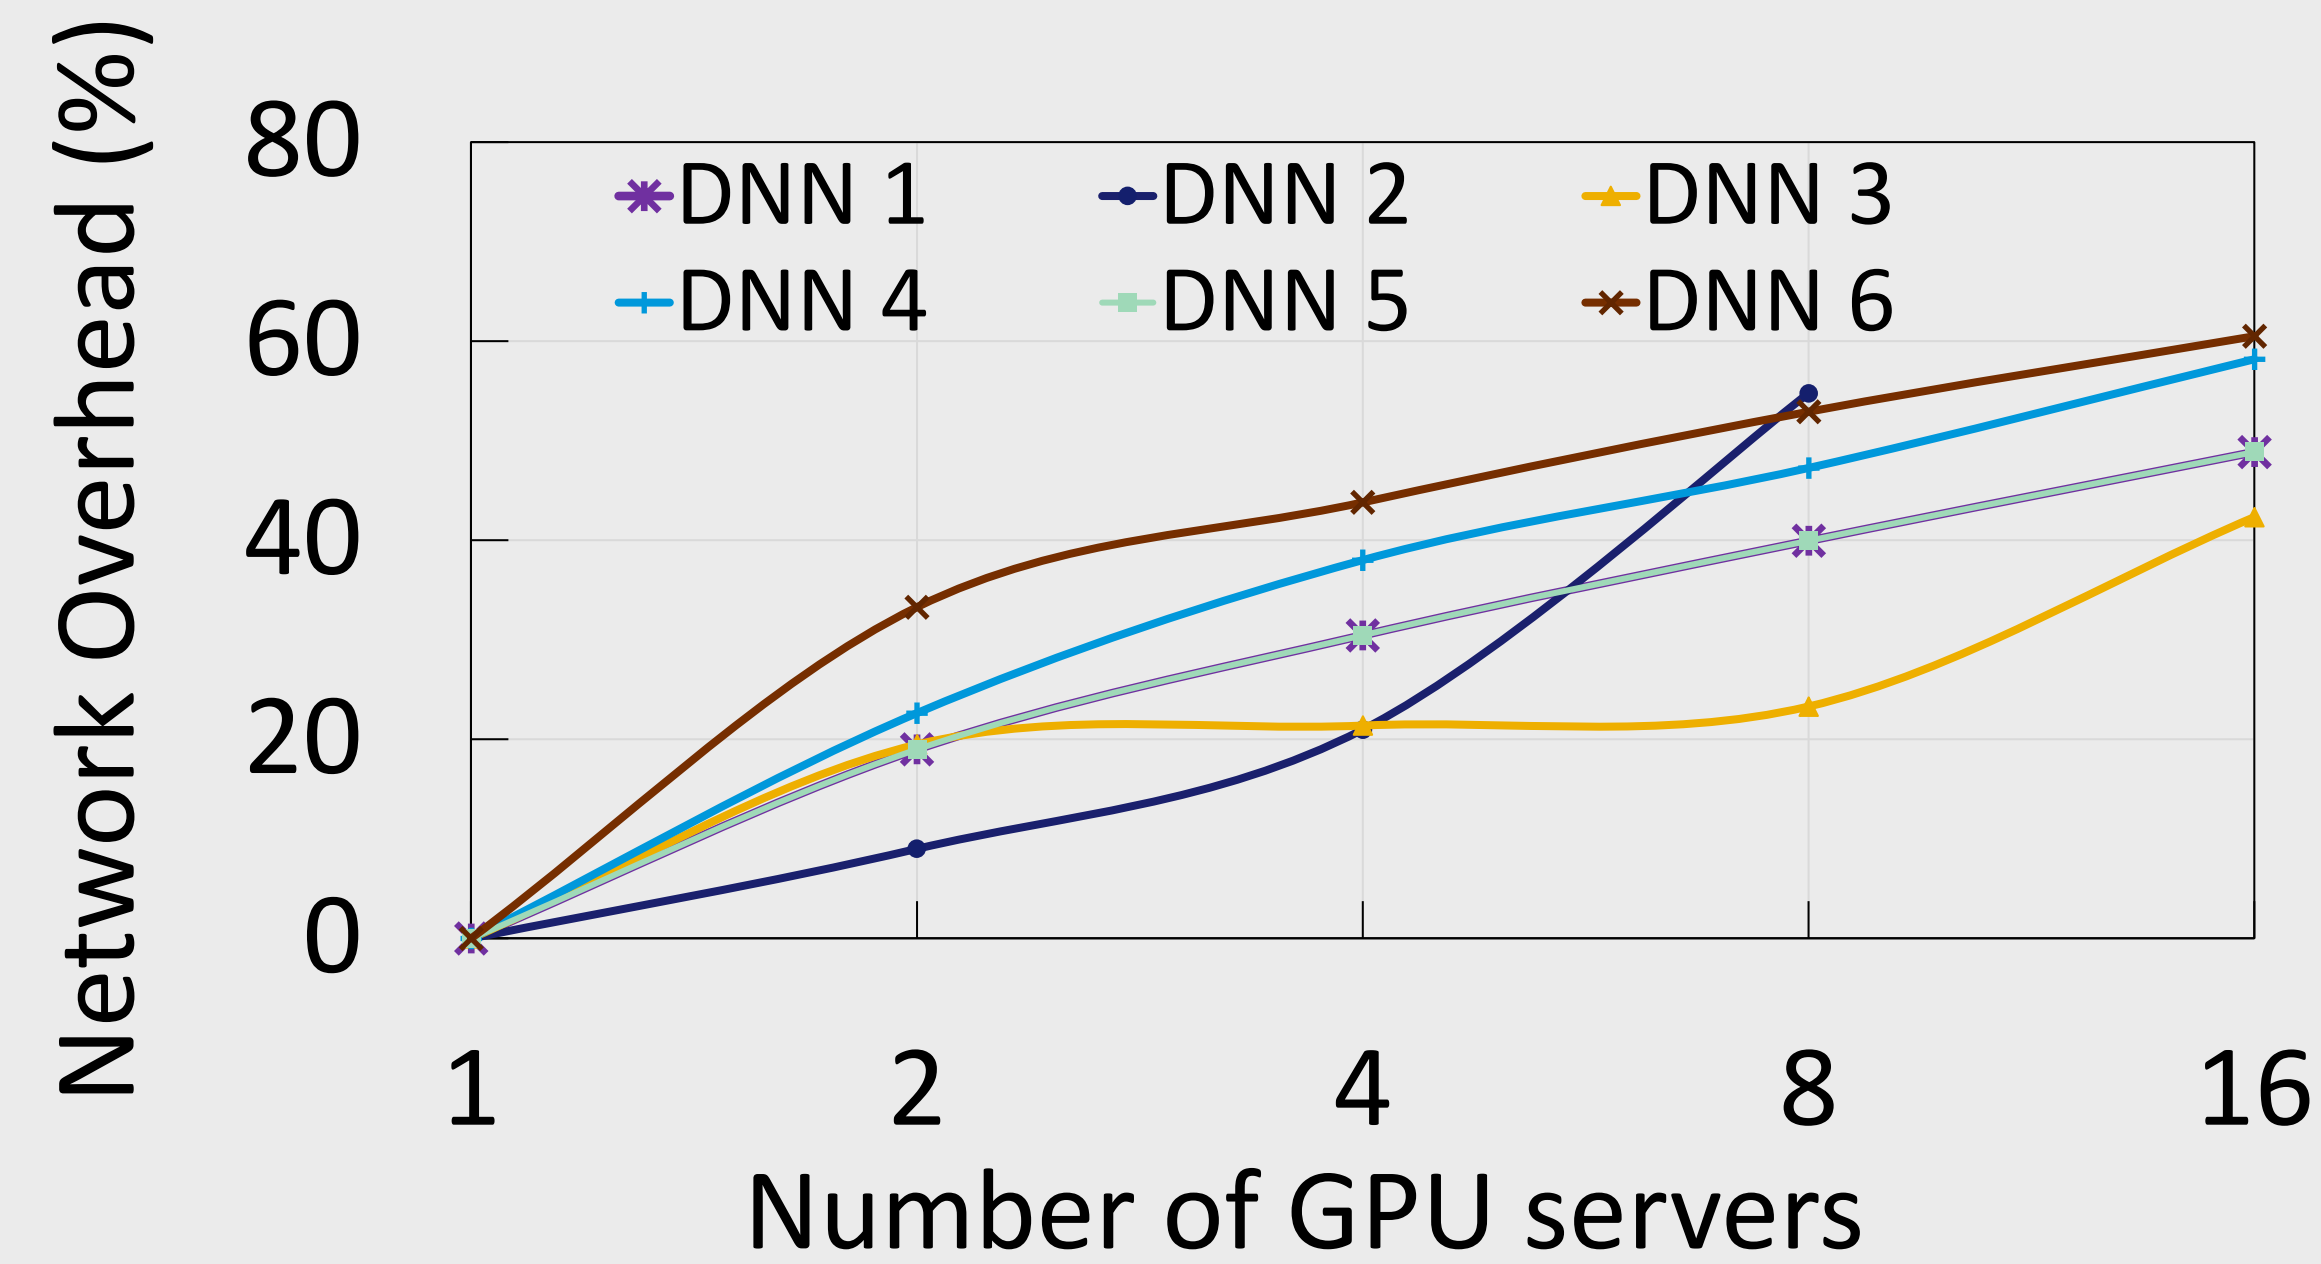
TopoOpt: Co-optimizing Network Topology and Parallelization Strategy for Distributed Training Jobs

nsdi'23



Weiyang (Frank) Wang, Moein Khazraee, Zhizhen Zhong, Manya Ghobadi, Zhihao Jia, Dheevatsa Mudigere, Ying Zhang, Anthony Kewitsch

Network is becoming a bottleneck for DNN training



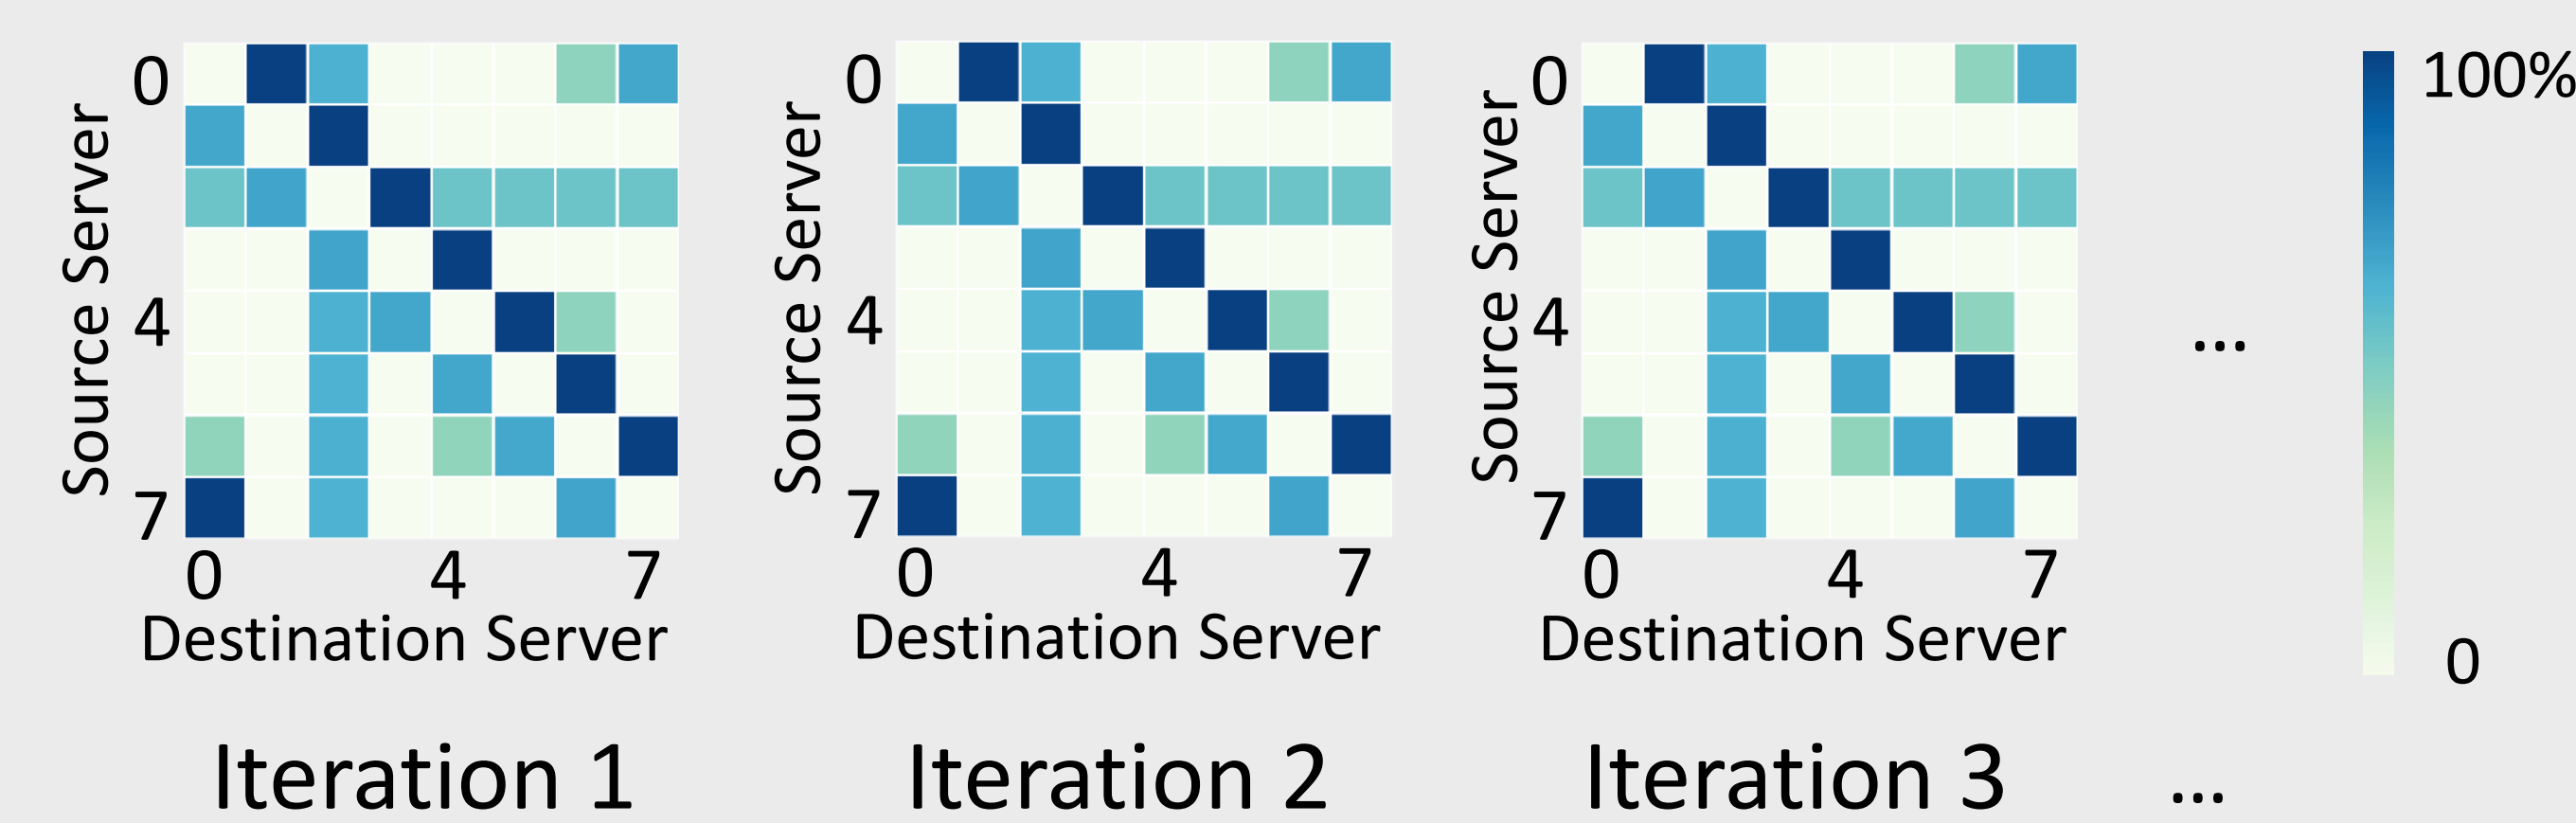
Network Overhead: the amount of time spent on communication only

DNN training traffic pattern

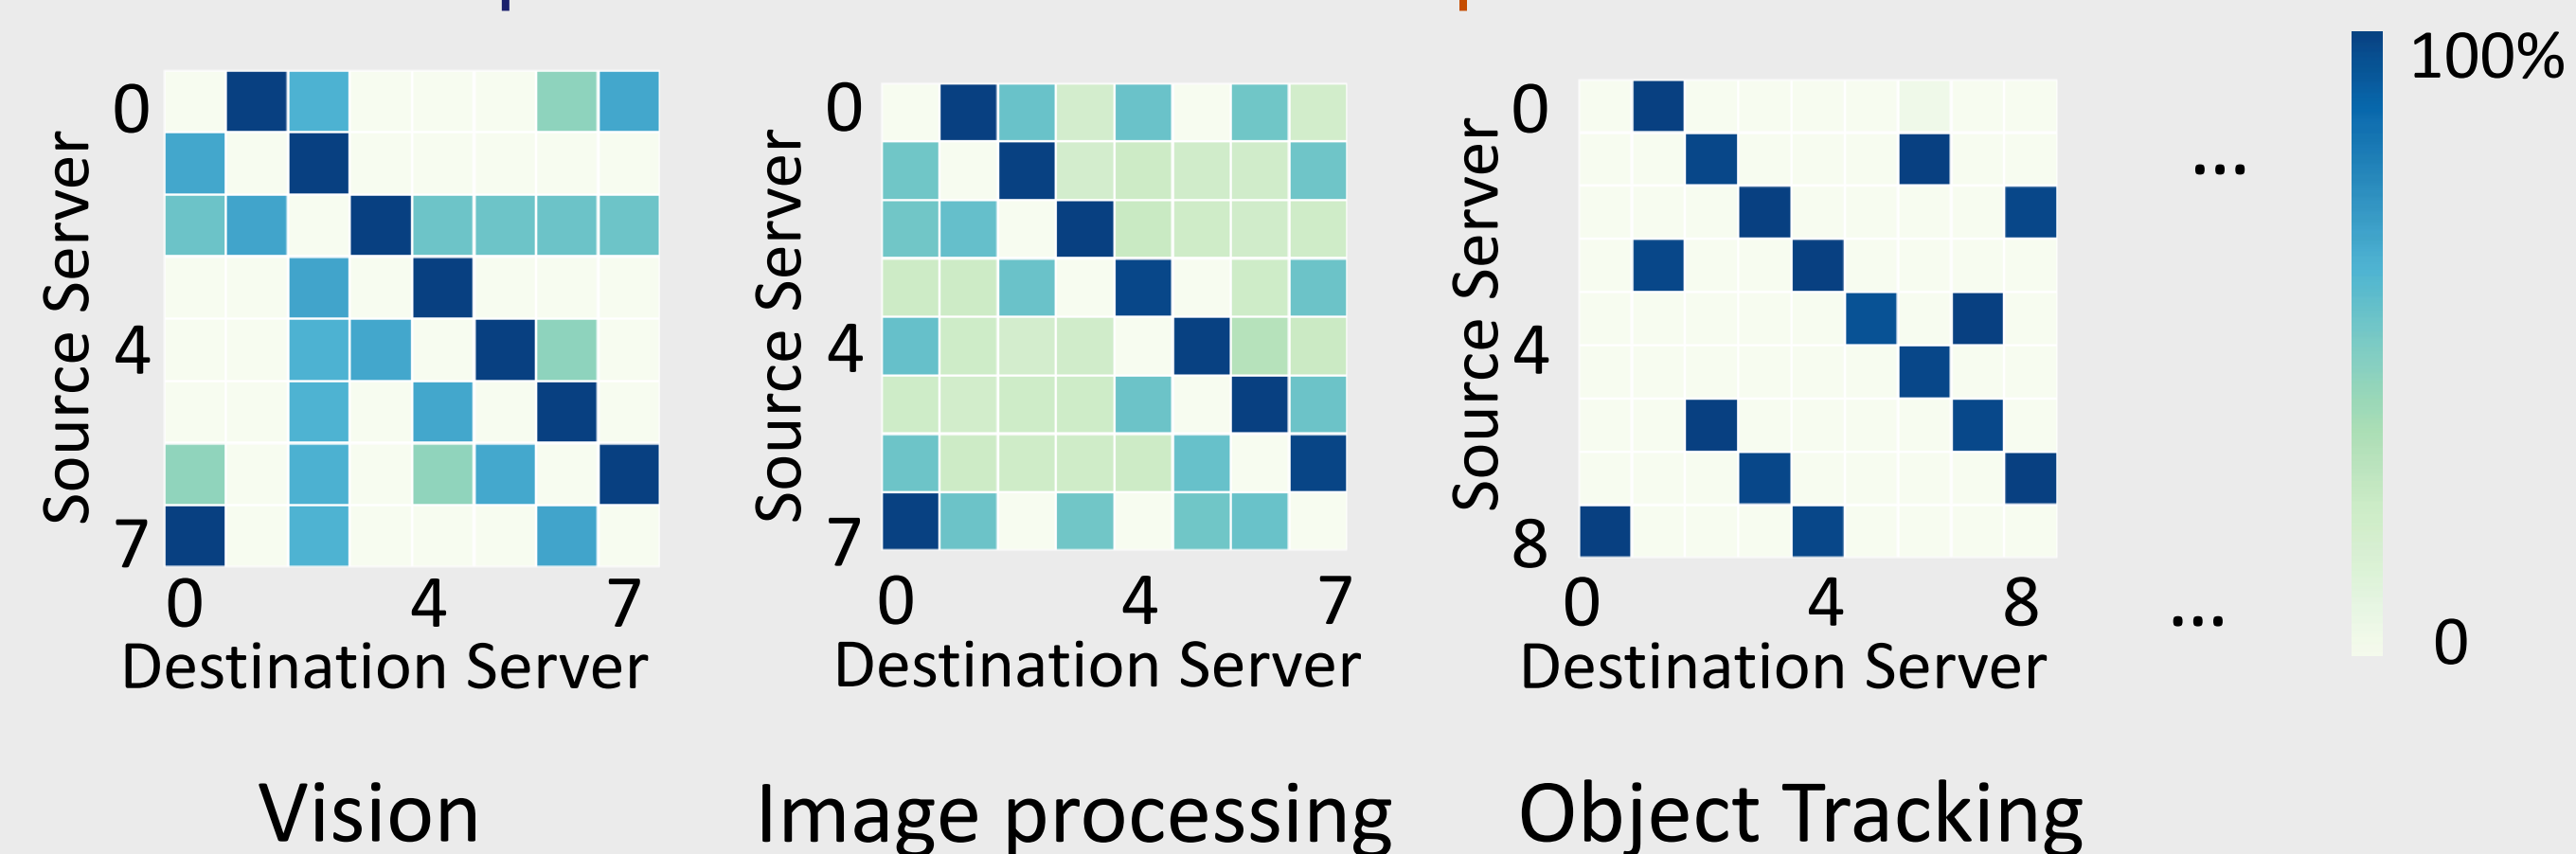
- Existing datacenter networks adapt to **unpredictable** workloads consisting **mainly of short transfers**
- These characteristics are no longer valid for DNN training

- For DNN training:

- Traffic patterns **do not change** across training iterations



- Traffic patterns are **model-dependent**



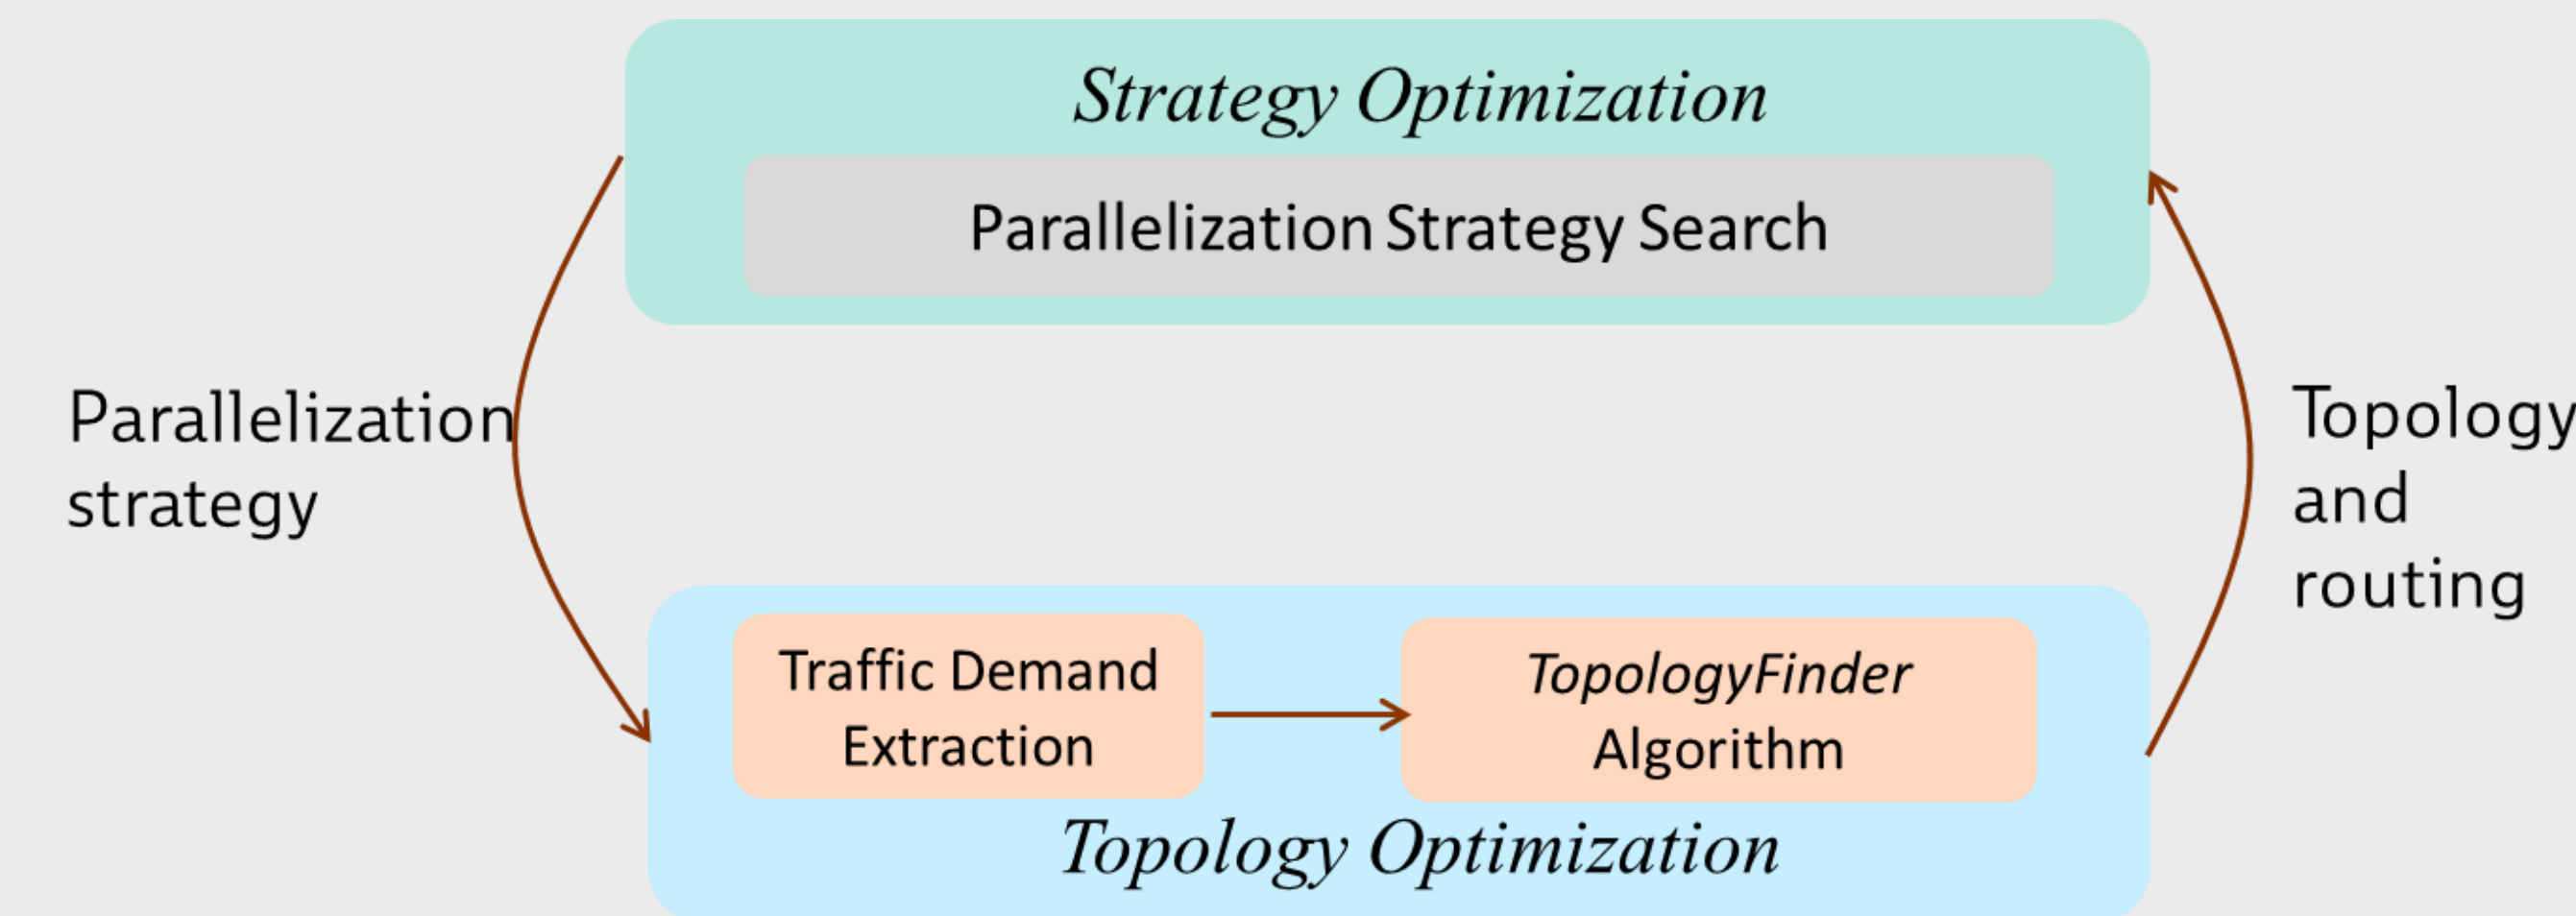
Our Solution: TopoOpt

The **first system** to leverage **reconfigurable network**, to co-optimize network topology and parallelization strategy for distributed training.


3.4x faster DNN training, compared to similar cost Fat-Trees

TopoOpt System Design

Alternating optimization framework



Direct-connect network using optical switches

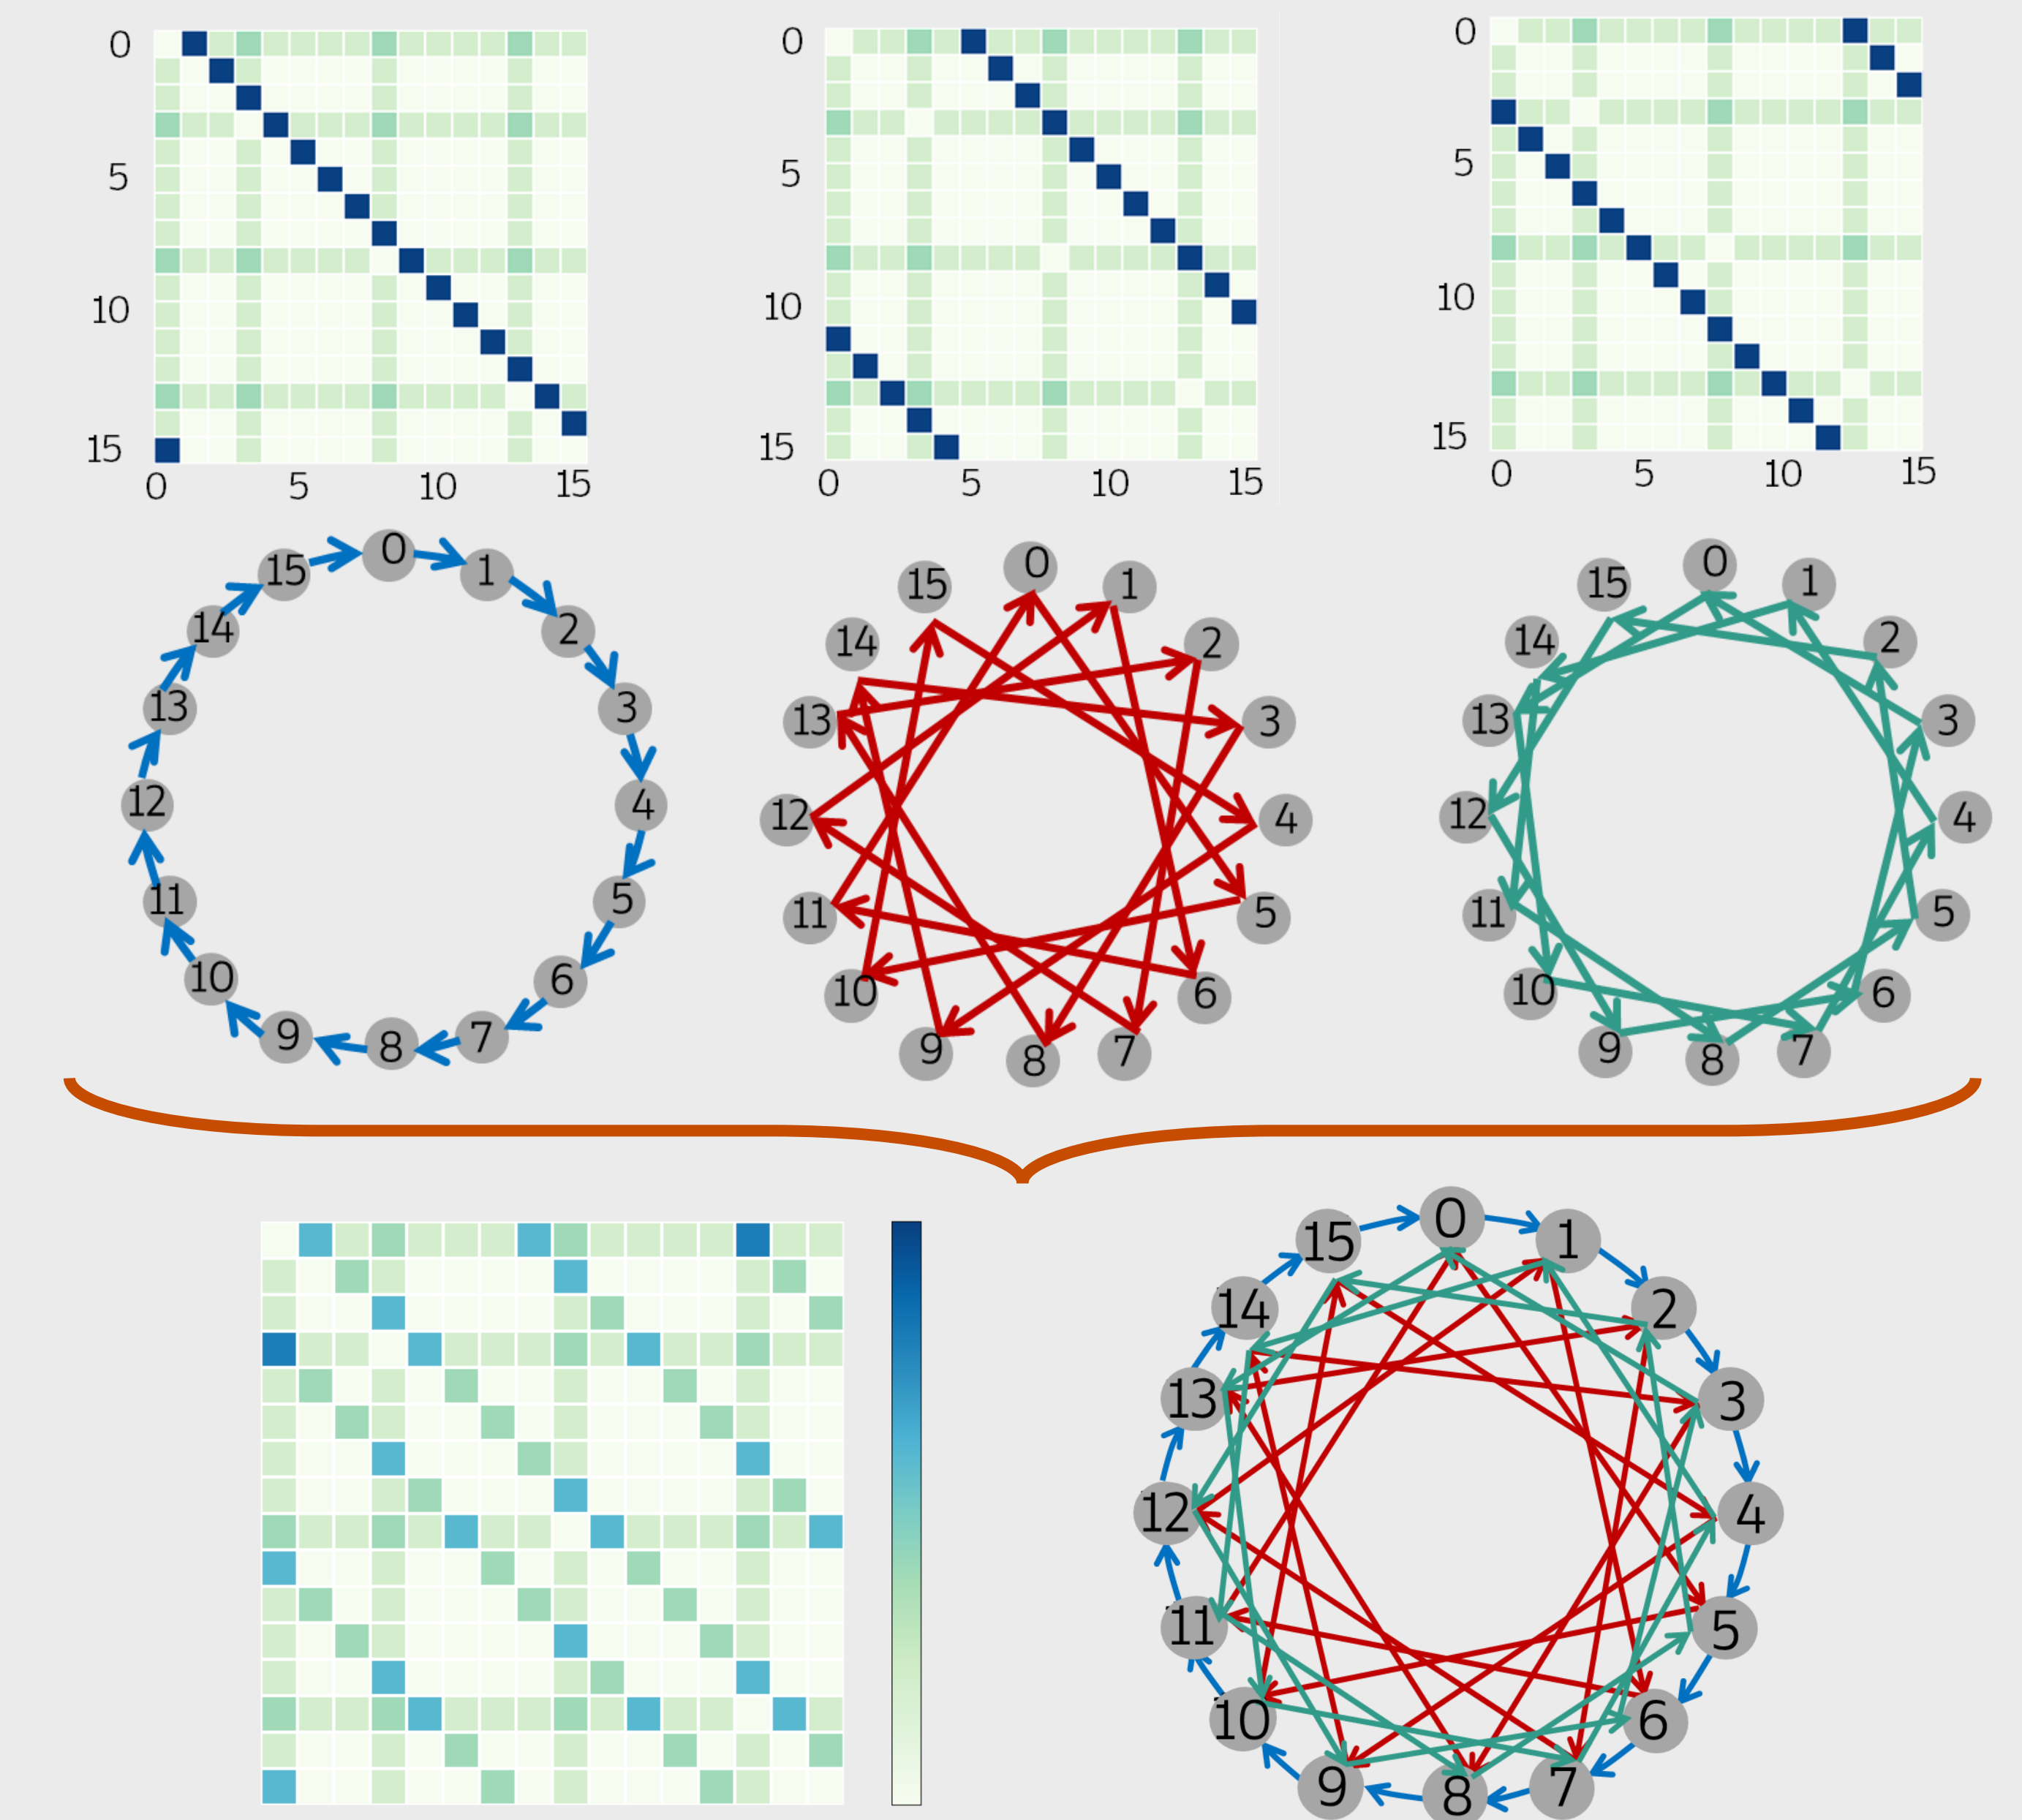


Challenge: Different characteristics for different transfer types

Transfer Type	Characteristics	Network Requirement
Collective Communication	Large, Sparse	Ample Bandwidth
Point-to-Point Communication	Small, Dense	Low hop-count

Key idea: Mutate the traffic matrix


- Collective Communications (CC) are **mutable**
- Point-to-Point transfers (P2P) are not mutable



- Leverage the mutability of CC transfers to achieve high bandwidth for CC & low hop-count for P2P!

Result: Faster DNN training

Simulation - 432 servers shared among DNN jobs



Testbed - Training a VGG19 model

