CS59200: AI/DC Networking

Instructor: Vamsi Addanki
Computer Science Department
Purdue University

Fall, 2025

Administrative Details

Course Code: CS59200
Semester: Fall 2025
Credit hours: 3
Prerequisites: CS536 or instructor consent. Students are expected to have prior background in computer networks, as well as familiarity with basic algorithms, linear optimization, and introductory graph theory.
Location & timings: TBD
Instructor: Vamsi Addanki
Email: vaddank@purdue.edu

Target Audience

This course is particularly valuable for early-stage PhD students in computer science who are interested in pursuing research at the intersection of AI, networking, and computing systems.

Course Overview

Have you ever wondered what makes massive AI models possible? How do thousands of GPUs communicate to train today's largest models? Can AI itself help configure and optimize the networks that connect them? From cutting-edge datacenter architectures to adaptive photonic interconnects that bend light, this course dives into the critical role networking plays in enabling AI at scale — and how AI, in turn, is transforming the way we build and manage networks. In this course, we will explore the algorithms, architectures, and open research challenges at the frontier where AI and networking meet.

This seminar-style course will cover a range of topics, including: datacenter network topologies, collective communication algorithms (GPU-to-GPU communication), photonic interconnects, network congestion control and load balancing, AI-assisted algorithm design, and the use of AI in network management and optimization. The tentative weekly schedule is as follows. Optional reading material is just for your reference to explore the literature (strongly recommended), but not explicitly required for the course.

Week	Tentative Date	Paper Title	Presenter	Optional Reading
Warmup!				
1	Day 1: 27 August 2025	Introduction	All	-
	Day 2: 29 August 2025	How to read a paper [28]	[Student Name]	[47, 48]
Large-Scale LLM Training Architectures – Experience Shared by HyperScalers				
2	Day 1: 3 September 2025	RDMA over Ethernet for Distributed AI Training at Meta Scale [16]	[Student Name]	[22, 46, 59]
	Day 2: 5 September 2025	Alibaba HPN: A Data Center Network for Large Language Model Training [45]	[Student Name]	[56, 17]
Collective Communication Algorithms I: Primitives and Allreduce				
3	Day 1: 10 September 2025	Optimization of Collective Communication Operations in MPICH [51]	[Student Name]	[9]
	Day 2: 12 September 2025	Swing: Short-cutting Rings for Higher Bandwidth Allreduce [49]	[Student Name]	[30]
Collective Communication Algorithms II: Synthesis				
4	Day 1: 17 September 2025	Synthesizing optimal collective algorithms [8]	[Student Name]	[54]
	Day 2: 19 September 2025	Rethinking Machine Learning Collective Communication as a Multi-Commodity Flow Problem [36]	[Student Name]	[50, 60]
Collective Communication Algorithms III: Straggler GPUs				
5	Day 1: 24 September 2025	Accelerating AllReduce with a Persistent Straggler [11]	[Student Name]	[23]
	Day 2: 26 September 2025	OptiReduce: Resilient and Tail-Optimal AllReduce for Distributed Deep Learning in the Cloud [57]	[Student Name]	[27, 10]
		Assignment-1 Due	All	
Photonic Interconnects I: Oblivious & Traffic-Aware				
6	Day 1: 1 October 2025	RotorNet: A Scalable, Low-complexity, Optical Datacenter Network [39]	[Student Name]	[38, 6, 1, 5]
	Day 2: 3 October 2025	Scheduling techniques for hybrid circuit/packet networks [34]	[Student Name]	[42, 15, 18, 33]
Photonic Interconnects II: TPU Clusters				

Week	Tentative Date	Paper Title	Presenter	Optional Reading
7	Day 1: 8 October 2025	TPU v4: An Optically Reconfigurable Supercomputer for Machine Learning with Hardware Support for Embeddings [25]	[Student Name]	[26]
	Day 2: 10 October 2025	Resiliency at Scale: Managing Google's TPUv4 Machine Learning Supercomputer [62]	[Student Name]	[43]
Photonic Interconnects III: Topologies for Collective Communication				
8	Day 1: 15 October 2025	SiP-ML: high-bandwidth optical network interconnects for machine learning training [29]	[Student Name]	[61]
	Day 2: 17 October 2025	TopoOpt: Co-optimizing Network Topology and Parallelization Strategy for Distributed Training Jobs [55]	[Student Name]	[32]
		Assignment-2 Due	All	
Photonic Interconnects IV: Chip-to-Chip Networks				
9	Day 1: 22 October 2025	A case for server-scale photonic connectivity [31]	[Student Name]	[12, 52]
	Day 2: 24 October 2025	*Midterm examination*	All	-
AI for Networks I: The fun part of applying LLMs				
10	Day 1: 29 October 2025	Enhancing Network Management Using Code Generated by Large Language Models [37]	[Student Name]	[7]
	Day 2: 31 October 2025	What do LLMs need to Synthesize Correct Router Configurations? [40]	[Student Name]	[24, 53]
AI for Networks II: The not so fun but beautiful part of performance guarantees				
11	Day 1: 5 November 2025	Credence: Augmenting Datacenter Switch Buffer Sharing with ML Predictions [2]	[Student Name]	[44]
	Day 2: 7 November 2025	Towards Integrating Formal Methods into ML-Based Systems for Networking [19]	[Student Name]	[20]
AI for Networks III: It gets serious with wide-area networks				
12	Day 1: 12 November 2025	DOTE: Rethinking (Predictive) WAN Traffic Engineering [41]	[Student Name]	[21]

Week	Tentative Date	Paper Title	Presenter	Optional Reading
	Day 2: 14 November 2025	Transferable Neural WAN TE for Changing Topologies [4]	[Student Name]	[35]
AI for Networks IV: It's fun, but congestion is a never-ending problem				
13	Day 1: 19 November 2025	TCP ex machina: computer-generated congestion control [58]	[Student Name]	[3]
	Day 2: 21 November 2025	PCC: Re-architecting congestion control for consistent high performance [13]	[Student Name]	[14]
		Assignment-3 Due	All	
Thanksgiving Break				
14	Day 1: 26 November 2025	-	-	-
	Day 2: 28 November 2025	-	-	-
Discussions and Feedback on Project/Paper Progress				
15	Day 1: 3 December 2025	Submissions due by next week	All	-
	Day 2: 5 December 2025	Submissions due by next week	-	-
Week 16				
16	Day 1: 10 December 2025	Final Presentations	All	-
	Day 2: 12 December 2025	Final Presentations	All	-

Assignments, Midterm and Final Project

The course is structured around student-led presentations and discussions held during weekly sessions, with the instructor providing guidance and facilitating exploration of the material. Each student will present assigned research papers to the class and participate in discussions to enhance collective understanding. Course evaluation is based on three assignments, one midterm exam, and a final research project.

Assignment 1. Each student will be assigned an AllReduce algorithm (or a synthesized variant) to implement in the Astra-Sim simulator. The simulation should use a ring topology of 16 nodes, each connected by 400 Gbps links with a 500 ns propagation delay. The goal is to evaluate the algorithm's completion time and compare it against the baseline Ring AllReduce across a range of message sizes.

Assignment 2. Building on the first assignment, extend the implementation to a reconfigurable ring topology where nodes are connected via a photonic switch. The objective is to optimize the circuit-switching schedule to minimize the AllReduce completion time. Students may submit either: (i) A proof showing the minimized completion time based on an optimized schedule, or (ii) Simulation results using Astra-Sim, along with a clear description of the optimization method used.

Assignment-3 (Option 1): Implement the assigned AllReduce algorithm using NVIDIA Collective Communication Library (NCCL) by writing a CUDA code[1] or using torch in python, and evaluate the performance in an 8 GPU cluster (access will be provided).

[1]nccl-tests repository provides a good starting point for implementation.

Assignment-3 (Option 2): Extend HTTP/3 QUIC transport protocol with a learning-augmented congestion control algorithm (e.g., having Cubic or Reno as base algorithms and leveraging ML predictions about network conditions) and implement in aioquic. Test the final implementation by sending iperf traffic to different remote servers and compare the throughput-latency-loss performance against the baseline QUIC implementation. Students may explore any learning algorithm of their choice, with emphasis on the techniques and methods discussed in the course schedule.

Midterm: The format of the midterm will be announced during the semester and will focus on the core concepts underlying the algorithms and protocol designs covered in the weekly readings.

Final project: The course concludes with a final research project, to be submitted via an internal HotCRP portal. Evaluation of the project will primarily consider the originality of the algorithmic or systems design proposal, the depth of related work understanding, and the quality of the presentation.

Learning Objectives

By the end of this seminar, students will be able to:

- Develop critical thinking in networking problems.
- Lead and participate in academic discussions.
- Analyze and present research papers effectively.
- Explore and propose innovative solutions to current challenges in AI/DC networking.
- Implement and test GPU communications algorithms with production-grade tools, such as NCCL, and with simulators, such as Astra-Sim.
- Write papers on the topics discussed in the course.

Evaluation

- Assignments: 25% (3 assignments, 8.33% each)
- Midterm examination: 25%
- Final project/paper: 50%

References

[1] Vamsi Addanki, Chen Avin, and Stefan Schmid. Mars: Near-optimal throughput with shallow buffers in reconfigurable datacenter networks. *Proc. ACM Meas. Anal. Comput. Syst.*, 7(1), March 2023. `doi:10.1145/3579312`.

[2] Vamsi Addanki, Maciej Pacut, and Stefan Schmid. Credence: Augmenting datacenter switch buffer sharing with ML predictions. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*, pages 613–634, Santa Clara, CA, April 2024. USENIX Association. URL: `https://www.usenix.org/conference/nsdi24/presentation/addanki-credence`.

[3] Neil Agarwal, Rui Pan, Francis Y. Yan, and Ravi Netravali. Mowgli: Passively learned rate control for Real-Time video. In *22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI 25)*, pages 579–594, Philadelphia, PA, April 2025. USENIX Association. URL: `https://www.usenix.org/conference/nsdi25/presentation/agarwal`.

[4] Abd AlRhman AlQiam, Yuanjun Yao, Zhaodong Wang, Satyajeet Singh Ahuja, Ying Zhang, Sanjay G. Rao, Bruno Ribeiro, and Mohit Tawarmalani. Transferable neural wan te for changing topologies. In *Proceedings of the ACM SIGCOMM 2024 Conference*, ACM SIGCOMM '24, page 86–102, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672237.

[5] Daniel Amir, Nitika Saran, Tegan Wilson, Robert Kleinberg, Vishal Shrivastav, and Hakim Weatherspoon. Shale: A practical, scalable oblivious reconfigurable network. In *Proceedings of the ACM SIGCOMM 2024 Conference*, ACM SIGCOMM '24, page 449–464, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672248.

[6] Hitesh Ballani, Paolo Costa, Raphael Behrendt, Daniel Cletheroe, Istvan Haller, Krzysztof Jozwik, Fotini Karinou, Sophie Lange, Kai Shi, Benn Thomsen, and Hugh Williams. Sirius: A flat datacenter network with nanosecond optical switching. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM '20, page 782–797, New York, NY, USA, 2020. Association for Computing Machinery. doi:10.1145/3387514.3406221.

[7] Luca Beurer-Kellner, Martin Vechev, Laurent Vanbever, and Petar Veličković. Learning to configure computer networks with neural algorithmic reasoning. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 730–742. Curran Associates, Inc., 2022. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/04cc90ec6868b97b7423dc38ced1e35c-Paper-Conference.pdf.

[8] Zixian Cai, Zhengyang Liu, Saeed Maleki, Madanlal Musuvathi, Todd Mytkowicz, Jacob Nelson, and Olli Saarikivi. Synthesizing optimal collective algorithms. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP '21, page 62–75, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3437801.3441620.

[9] Ernie Chan, Marcel Heimlich, Avi Purkayastha, and Robert van de Geijn. Collective communication: theory, practice, and experience. *Concurrency and Computation: Practice and Experience*, 19(13):1749–1783, 2007. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.1206, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/cpe.1206, doi:10.1002/cpe.1206.

[10] James Cipar, Qirong Ho, Jin Kyu Kim, Seunghak Lee, Gregory R Ganger, Garth Gibson, Kimberly Keeton, and Eric Xing. Solving the straggler problem with bounded staleness. In *14th Workshop on Hot Topics in Operating Systems (HotOS XIV)*, 2013.

[11] Arjun Devraj, Eric Ding, Abhishek Vijaya Kumar, Robert Kleinberg, and Rachee Singh. Accelerating allreduce with a persistent straggler. *arXiv preprint arXiv:2505.23523*, 2025. URL: https://arxiv.org/abs/2505.23523.

[12] Eric Ding and Rachee Singh. Pipswitch: A circuit switch using programmable integrated photonics. *arXiv preprint arXiv:2501.18136*, 2025. URL: https://arxiv.org/abs/2501.18136.

[13] Mo Dong, Qingxi Li, Doron Zarchy, P Brighten Godfrey, and Michael Schapira. {PCC}: Re-architecting congestion control for consistent high performance. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, pages 395–408, 2015. URL: https://www.usenix.org/system/files/conference/nsdi15/nsdi15-paper-dong.pdf.

[14] Mo Dong, Tong Meng, Doron Zarchy, Engin Arslan, Yossi Gilad, Brighten Godfrey, and Michael Schapira. PCC vivace: Online-Learning congestion control. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*, pages 343–356, Renton, WA, April 2018. USENIX Association. URL: https://www.usenix.org/conference/nsdi18/presentation/dong.

[15] Nathan Farrington, George Porter, Sivasankar Radhakrishnan, Hamid Hajabdolali Bazzaz, Vikram Subramanya, Yeshaiahu Fainman, George Papen, and Amin Vahdat. Helios: a hybrid electrical/optical switch architecture for modular data centers. In *Proceedings of the ACM SIGCOMM 2010 Conference*, SIGCOMM '10, page 339–350, New York, NY, USA, 2010. Association for Computing Machinery. doi:10.1145/1851182.1851223.

[16] Adithya Gangidi, Rui Miao, Shengbao Zheng, Sai Jayesh Bondu, Guilherme Goes, Hany Morsy, Rohit Puri, Mohammad Riftadi, Ashmitha Jeevaraj Shetty, Jingyi Yang, Shuqiang Zhang, Mikel Jimenez Fernandez, Shashidhar Gandham, and Hongyi Zeng. Rdma over ethernet for distributed training at meta scale. In *Proceedings of the ACM SIGCOMM 2024 Conference*, ACM SIGCOMM '24, page 57–70, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672233.

[17] Alexandru M. Gherghescu, Vlad-Andrei Bădoiu, Alexandru Agache, Mihai-Valentin Dumitru, Iuliu Vasilescu, Radu Mantu, and Costin Raiciu. I've got 99 problems but flops ain't one. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*, HotNets '24, page 195–204, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3696348.3696893.

[18] Monia Ghobadi, Ratul Mahajan, Amar Phanishayee, Nikhil Devanur, Janardhan Kulkarni, Gireeja Ranade, Pierre-Alexandre Blanche, Houman Rastegarfar, Madeleine Glick, and Daniel Kilper. Projector: Agile reconfigurable data center interconnect. In *Proceedings of the 2016 ACM SIGCOMM Conference*, SIGCOMM '16, page 216–229, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2934872.2934911.

[19] Fengchen Gong, Divya Raghunathan, Aarti Gupta, and Maria Apostolaki. Towards integrating formal methods into ml-based systems for networking. In *Proceedings of the 22nd ACM Workshop on Hot Topics in Networks*, HotNets '23, page 48–55, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3626111.3628188.

[20] Fengchen Gong, Divya Raghunathan, Aarti Gupta, and Maria Apostolaki. Zoom2net: Constrained network telemetry imputation. In *Proceedings of the ACM SIGCOMM 2024 Conference*, ACM SIGCOMM '24, page 764–777, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672225.

[21] Fei Gui, Songtao Wang, Dan Li, Li Chen, Kaihui Gao, Congcong Min, and Yi Wang. Redte: Mitigating subsecond traffic bursts with real-time and distributed traffic engineering. In *Proceedings of the ACM SIGCOMM 2024 Conference*, ACM SIGCOMM '24, page 71–85, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672231.

[22] Chuanxiong Guo, Haitao Wu, Zhong Deng, Gaurav Soni, Jianxi Ye, Jitu Padhye, and Marina Lipshteyn. Rdma over commodity ethernet at scale. In *Proceedings of the 2016 ACM SIGCOMM Conference*, SIGCOMM '16, page 202–215, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2934872.2934908.

[23] Aaron Harlap, Henggang Cui, Wei Dai, Jinliang Wei, Gregory R. Ganger, Phillip B. Gibbons, Garth A. Gibson, and Eric P. Xing. Addressing the straggler problem for iterative convergent parallel ml. In *Proceedings of the Seventh ACM Symposium on Cloud Computing*, SoCC '16, page 98–111, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2987550.2987554.

[24] Zhiyuan He, Aashish Gottipati, Lili Qiu, Xufang Luo, Kenuo Xu, Yuqing Yang, and Francis Y. Yan. Designing network algorithms via large language models. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*, HotNets '24, page 205–212, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3696348.3696868.

[25] Norm Jouppi, George Kurian, Sheng Li, Peter Ma, Rahul Nagarajan, Lifeng Nai, Nishant Patil, Suvinay Subramanian, Andy Swing, Brian Towles, Clifford Young, Xiang Zhou, Zongwei Zhou, and David A Patterson. Tpu v4: An optically reconfigurable supercomputer for machine learning with hardware support for embeddings. In *Proceedings of the 50th Annual International Symposium on Computer Architecture*, ISCA '23, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3579371.3589350.

[26] Norman P. Jouppi and Andy Swing. A machine learning supercomputer with an optically reconfigurable interconnect and embeddings support. In *2023 IEEE Hot Chips 35 Symposium (HCS)*, pages 1–24, 2023. doi:10.1109/HCS59251.2023.10254691.

[27] Can Karakus, Yifan Sun, Suhas Diggavi, and Wotao Yin. Straggler mitigation in distributed optimization through data encoding. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL: https://proceedings.neurips.cc/paper_files/paper/2017/file/663772ea088360f95bac3dc7ffb841be-Paper.pdf.

[28] S. Keshav. How to read a paper. *SIGCOMM Comput. Commun. Rev.*, 37(3):83–84, July 2007. doi:10.1145/1273445.1273458.

[29] Mehrdad Khani, Manya Ghobadi, Mohammad Alizadeh, Ziyi Zhu, Madeleine Glick, Keren Bergman, Amin Vahdat, Benjamin Klenk, and Eiman Ebrahimi. Sip-ml: high-bandwidth optical network interconnects for machine learning training. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, SIGCOMM '21, page 657–675, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3452296.3472900.

[30] Dmitry Kolmakov and Xuecang Zhang. A generalization of the allreduce operation. *arXiv preprint arXiv:2004.09362*, 2020. URL: https://arxiv.org/abs/2004.09362.

[31] Abhishek Vijaya Kumar, Arjun Devraj, Darius Bunandar, and Rachee Singh. A case for server-scale photonic connectivity. In *Proceedings of the 23rd ACM Workshop on Hot Topics in Networks*, HotNets '24, page 290–299, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3696348.3696856.

[32] Abhishek Vijaya Kumar, Eric Ding, Arjun Devraj, Darius Bunandar, and Rachee Singh. Lumion: Fast fault recovery for ml jobs using programmable optical fabrics. *arXiv preprint arXiv:2505.23105*, 2025. URL: https://arxiv.org/abs/2505.23105.

[33] Cong Liang, Xiangli Song, Jing Cheng, Mowei Wang, Yashe Liu, Zhenhua Liu, Shizhen Zhao, and Yong Cui. Negotiator: Towards a simple yet effective on-demand reconfigurable datacenter network. In *Proceedings of the ACM SIGCOMM 2024 Conference*, ACM SIGCOMM '24, page 415–432, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672222.

[34] He Liu, Matthew K. Mukerjee, Conglong Li, Nicolas Feltman, George Papen, Stefan Savage, Srinivasan Seshan, Geoffrey M. Voelker, David G. Andersen, Michael Kaminsky, George Porter, and Alex C. Snoeren. Scheduling techniques for hybrid circuit/packet networks. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, CoNEXT '15, New York, NY, USA, 2015. Association for Computing Machinery. doi:10.1145/2716281.2836126.

[35] Ximeng Liu, Shizhen Zhao, Yong Cui, and Xinbing Wang. Figret: Fine-grained robustness-enhanced traffic engineering. In *Proceedings of the ACM SIGCOMM 2024 Conference*, ACM SIGCOMM '24, page 117–135, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672258.

[36] Xuting Liu, Behnaz Arzani, Siva Kesava Reddy Kakarla, Liangyu Zhao, Vincent Liu, Miguel Castro, Srikanth Kandula, and Luke Marshall. Rethinking machine learning collective communication as a multi-commodity flow problem. In *Proceedings of the ACM SIGCOMM 2024 Conference*, ACM SIGCOMM '24, page 16–37, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672249.

[37] Sathiya Kumaran Mani, Yajie Zhou, Kevin Hsieh, Santiago Segarra, Trevor Eberl, Eliran Azulai, Ido Frizler, Ranveer Chandra, and Srikanth Kandula. Enhancing network management using code generated by large language models. In *Proceedings of the 22nd ACM Workshop on Hot Topics in Networks*, HotNets '23, page 196–204, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3626111.3628183.

[38] William M. Mellette, Alex Forencich, Rukshani Athapathu, Alex C. Snoeren, George Papen, and George Porter. Realizing rotornet: Toward practical microsecond scale optical networking. In *Proceedings of the ACM SIGCOMM 2024 Conference*, ACM SIGCOMM '24, page 392–414, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672273.

[39] William M. Mellette, Rob McGuinness, Arjun Roy, Alex Forencich, George Papen, Alex C. Snoeren, and George Porter. Rotornet: A scalable, low-complexity, optical datacenter network. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, SIGCOMM '17, page 267–280, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3098822.3098838.

[40] Rajdeep Mondal, Alan Tang, Ryan Beckett, Todd Millstein, and George Varghese. What do llms need to synthesize correct router configurations? In *Proceedings of the 22nd ACM Workshop on Hot Topics in Networks*, HotNets '23, page 189–195, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3626111.3628194.

[41] Yarin Perry, Felipe Vieira Frujeri, Chaim Hoch, Srikanth Kandula, Ishai Menache, Michael Schapira, and Aviv Tamar. DOTE: Rethinking (predictive) WAN traffic engineering. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 1557–1581, Boston, MA, April 2023. USENIX Association. URL: https://www.usenix.org/conference/nsdi23/presentation/perry.

[42] George Porter, Richard Strong, Nathan Farrington, Alex Forencich, Pang Chen-Sun, Tajana Rosing, Yeshaiahu Fainman, George Papen, and Amin Vahdat. Integrating microsecond circuit switching into the data center. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, SIGCOMM '13, page 447–458, New York, NY, USA, 2013. Association for Computing Machinery. doi:10.1145/2486001.2486007.

[43] Leon Poutievski, Omid Mashayekhi, Joon Ong, Arjun Singh, Mukarram Tariq, Rui Wang, Jianan Zhang, Virginia Beauregard, Patrick Conner, Steve Gribble, Rishi Kapoor, Stephen Kratzer, Nanfang Li, Hong Liu, Karthik Nagaraj, Jason Ornstein, Samir Sawhney, Ryohei Urata, Lorenzo Vicisano, Kevin Yasumura, Shidong Zhang, Junlan Zhou, and Amin Vahdat. Jupiter evolving: transforming google's datacenter network via optical circuit switches and software-defined networking. In *Proceedings of the ACM SIGCOMM 2022 Conference*, SIGCOMM '22, page 66–85, New York, NY, USA, 2022. Association for Computing Machinery. doi:10.1145/3544216.3544265.

[44] Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving online algorithms via ml predictions. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL: https://proceedings.neurips.cc/paper_files/paper/2018/file/73a427badebe0e32caa2e1fc7530b7f3-Paper.pdf.

[45] Kun Qian, Yongqing Xi, Jiamin Cao, Jiaqi Gao, Yichi Xu, Yu Guan, Binzhang Fu, Xuemei Shi, Fangbo Zhu, Rui Miao, Chao Wang, Peng Wang, Pengcheng Zhang, Xianlong Zeng, Eddie Ruan, Zhiping Yao, Ennan Zhai, and Dennis Cai. Alibaba hpn: A data center network for large language model training. In *Proceedings of the ACM SIGCOMM 2024 Conference*, ACM SIGCOMM '24, page 691–706, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3651890.3672265.

[46] Saeed Rashidi, Srinivas Sridharan, Sudarshan Srinivasan, and Tushar Krishna. Astra-sim: Enabling sw/hw co-design exploration for distributed dl training platforms. In *2020 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pages 81–92, 2020. doi:10.1109/ISPASS48437.2020.00018.

[47] Timothy Roscoe. Writing reviews for systems conferences, 2007. URL: https://people.inf.ethz.ch/troscoe/pubs/review-writing.pdf.

[48] Henning Schulzrinne. Writing technical articles. URL: https://www.cs.columbia.edu/~hgs/etc/writing-style.html.

[49] Daniele De Sensi, Tommaso Bonato, David Saam, and Torsten Hoefler. Swing: Short-cutting rings for higher bandwidth allreduce. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*, pages 1445–1462, Santa Clara, CA, April 2024. USENIX Association. URL: `https://www.usenix.org/conference/nsdi24/presentation/de-sensi`.

[50] Aashaka Shah, Vijay Chidambaram, Meghan Cowan, Saeed Maleki, Madan Musuvathi, Todd Mytkowicz, Jacob Nelson, Olli Saarikivi, and Rachee Singh. TACCL: Guiding collective algorithm synthesis using communication sketches. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 593–612, Boston, MA, April 2023. USENIX Association. URL: `https://www.usenix.org/conference/nsdi23/presentation/shah`.

[51] Rajeev Thakur, Rolf Rabenseifner, and William Gropp. Optimization of collective communication operations in mpich. *The International Journal of High Performance Computing Applications*, 19(1):49–66, 2005. `arXiv:https://doi.org/10.1177/1094342005051521`, `doi:10.1177/1094342005051521`.

[52] Mark Wade, Erik Anderson, Shahab Ardalan, Pavan Bhargava, Sidney Buchbinder, Michael L. Davenport, John Fini, Haiwei Lu, Chen Li, Roy Meade, Chandru Ramamurthy, Michael Rust, Forrest Sedgwick, Vladimir Stojanovic, Derek Van Orden, Chong Zhang, Chen Sun, Sergey Y. Shumarayev, Conor O'Keeffe, Tim T. Hoang, David Kehlet, Ravi V. Mahajan, Matthew T. Guzy, Allen Chan, and Tina Tran. Teraphy: A chiplet technology for low-power, high-bandwidth in-package optical i/o. *IEEE Micro*, 40(2):63–71, 2020. `doi:10.1109/MM.2020.2976067`.

[53] Changjie Wang, Mariano Scazzariello, Alireza Farshin, Simone Ferlin, Dejan Kostić, and Marco Chiesa. Netconfeval: Can llms facilitate network configuration? *Proc. ACM Netw.*, 2(CoNEXT2), June 2024. `doi:10.1145/3656296`.

[54] Guanhua Wang, Shivaram Venkataraman, Amar Phanishayee, Nikhil Devanur, Jorgen Thelin, and Ion Stoica. Blink: Fast and generic collectives for distributed ml. In I. Dhillon, D. Papailiopoulos, and V. Sze, editors, *Proceedings of Machine Learning and Systems*, volume 2, pages 172–186, 2020. URL: `https://proceedings.mlsys.org/paper_files/paper/2020/file/cd3a9a55f7f3723133fa4a13628cdf03-Paper.pdf`.

[55] Weiyang Wang, Moein Khazraee, Zhizhen Zhong, Manya Ghobadi, Zhihao Jia, Dheevatsa Mudigere, Ying Zhang, and Anthony Kewitsch. TopoOpt: Co-optimizing network topology and parallelization strategy for distributed training jobs. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 739–767, Boston, MA, April 2023. USENIX Association. URL: `https://www.usenix.org/conference/nsdi23/presentation/wang-weiyang`.

[56] Xizheng Wang, Qingxu Li, Yichi Xu, Gang Lu, Dan Li, Li Chen, Heyang Zhou, Linkang Zheng, Sen Zhang, Yikai Zhu, Yang Liu, Pengcheng Zhang, Kun Qian, Kunling He, Jiaqi Gao, Ennan Zhai, Dennis Cai, and Binzhang Fu. SimAI: Unifying architecture design and performance tuning for Large-Scale large language model training with scalability and precision. In *22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI 25)*, pages 541–558, Philadelphia, PA, April 2025. USENIX Association. URL: `https://www.usenix.org/conference/nsdi25/presentation/wang-xizheng-simai`.

[57] Ertza Warraich, Omer Shabtai, Khalid Manaa, Shay Vargaftik, Yonatan Piasetzky, Matty Kadosh, Lalith Suresh, and Muhammad Shahbaz. OptiReduce: Resilient and Tail-Optimal AllReduce for distributed deep learning in the cloud. In *22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI 25)*, pages 685–703, Philadelphia, PA, April 2025. USENIX Association. URL: `https://www.usenix.org/conference/nsdi25/presentation/warraich`.

[58] Keith Winstein and Hari Balakrishnan. Tcp ex machina: computer-generated congestion control. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, SIGCOMM '13, page 123–134, New York, NY, USA, 2013. Association for Computing Machinery. `doi:10.1145/2486001.2486020`.

[59] William Won, Taekyung Heo, Saeed Rashidi, Srinivas Sridharan, Sudarshan Srinivasan, and Tushar Krishna. Astra-sim2.0: Modeling hierarchical networks and disaggregated systems for large-model training at scale. In *2023 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pages 283–294, 2023. doi:10.1109/ISPASS57527.2023.00035.

[60] Guanbin Xu, Zhihao Le, Yinhe Chen, Zhiqi Lin, Zewen Jin, Youshan Miao, and Cheng Li. AutoCCL: Automated collective communication tuning for accelerating distributed and parallel DNN training. In *22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI 25)*, pages 667–683, Philadelphia, PA, April 2025. USENIX Association. URL: https://www.usenix.org/conference/nsdi25/presentation/xu-guanbin.

[61] Liangyu Zhao, Siddharth Pal, Tapan Chugh, Weiyang Wang, Jason Fantl, Prithwish Basu, Joud Khoury, and Arvind Krishnamurthy. Efficient Direct-Connect topologies for collective communications. In *22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI 25)*, pages 705–737, Philadelphia, PA, April 2025. USENIX Association. URL: https://www.usenix.org/conference/nsdi25/presentation/zhao-liangyu.

[62] Yazhou Zu, Alireza Ghaffarkhah, Hoang-Vu Dang, Brian Towles, Steven Hand, Safeen Huda, Adekunle Bello, Alexander Kolbasov, Arash Rezaei, Dayou Du, Steve Lacy, Hang Wang, Aaron Wisner, Chris Lewis, and Henri Bahini. Resiliency at scale: Managing Google's TPUv4 machine learning supercomputer. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*, pages 761–774, Santa Clara, CA, April 2024. USENIX Association. URL: https://www.usenix.org/conference/nsdi24/presentation/zu.

Instructor Contact: vaddank@purdue.edu
Office Hours: TBA