Provisioning Networks for AI Supercomputers: A Trace-Driven Study of Performance Sensitivity at Unprecedented Scale

Tommaso Bonato[1]
ETH Zurich; Microsoft Corporation
Zurich, Switzerland; Redmond, USA
tommaso.bonato@inf.ethz.ch

Amirreza Rastegari
Microsoft Corporation
Redmond, USA
arastegari@microsoft.com

Daniele De Sensi
Sapienza University of Rome
Rome, Italy
desensi@di.uniroma1.it

Timo Schneider
ETH Zurich
Zurich, Switzerland
timo.schneider@inf.ethz.ch

Siyuan Shen
ETH Zurich
Zurich, Switzerland
siyuan.shen@inf.ethz.ch

Ales Kubicek
ETH Zurich
Zurich, Switzerland
akubicek@student.ethz.ch

Prabhat Ram[1]
Microsoft Corporation
Redmond, USA
prabhat.ram@microsoft.com

Torsten Hoefler
ETH Zurich
Zurich, Switzerland
torsten.hoefler@inf.ethz.ch

***Abstract*— AI workloads on supercomputers make interconnect performance a primary design constraint. Hardware specifications and microbenchmarks alone do not show how real workloads respond to latency, bandwidth, and jitter. We present a trace-driven study of network sensitivity for training and inference workloads at scales up to 4,096 GPUs. Our methodology combines signature-based trace decomposition with linear programming to make broad network-parameter sweeps feasible. It preserves the recorded NCCL schedules and program dependencies, and estimates a lower bound on runtime under the LogGP model. We validate the compositional analysis against monolithic linear programs, LogGOPSim, and hardware measurements at feasible scales. The resulting latency-bandwidth envelopes identify low-sensitivity regions and performance cliffs that vary substantially across workloads. They also show when bandwidth upgrades provide meaningful returns and when additional network performance has little effect. These results translate execution traces into practical guidance for evaluating network configurations while making the model assumptions explicit.**

***Index Terms*—artificial intelligence, high performance computing, interconnection networks, performance analysis, linear programming**

I. Introduction

The scale of AI training infrastructure has grown by orders of magnitude in less than five years [1]–[3]. Clusters built to train frontier models now contain tens of thousands of accelerators interconnected by networks whose aggregate bandwidth exceeds petabits per second [4], [5]. The capital expenditure for a single such facility can easily exceed one billion dollars, and roughly 20% of the cost of training frontier models is attributable to the network fabric [2]. At this investment scale, the question of how much network performance AI workloads actually need is not merely academic. Overprovisioning wastes capital. Underprovisioning degrades training throughput, increasing time to convergence and inflating operational costs.

[1]Tommaso Bonato and Prabhat Ram have since left Microsoft Corporation.

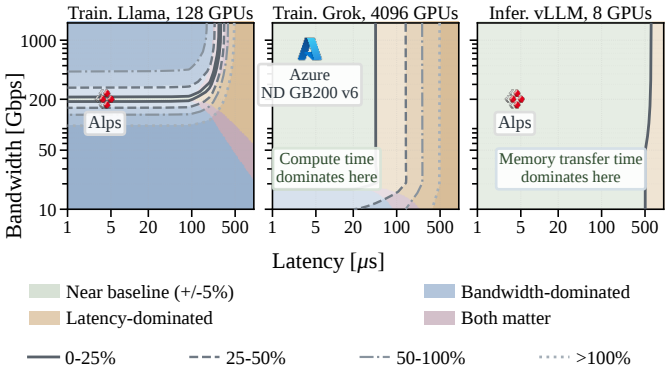

Figure 1. Latency/bandwidth sensitivity maps constructed from separate model-derived sweeps for each AI workload. Colors indicate the dominant runtime sensitivity regime; contours show increasing runtime change relative to each system's logo-marked baseline operating point.

Despite the stakes, the trade-off between network and model performance is not well studied. The dominant approach to AI cluster design treats network performance as a constraint to be maximized within a power and cost envelope: more bandwidth is better, less latency is better. This approach is reasonable when the relationship between network parameters and application performance is unknown, but it can lead to overprovisioning for some models and parallelism strategies. As shown in Figure 1, obtained with the toolchain presented in Section III, workloads such as xAI's Grok-1 [6] and vLLM [7] with Llama 8B [8] exhibit large regions of the latency-bandwidth parameter space that produce nearly identical training and inference runtimes under our model. Llama 70B, by contrast, is predominantly bandwidth-bound and becomes latency-sensitive only at very high latency values.

Historically, studying this trade-off has been difficult for multiple reasons: First, AI training workloads exhibit com-

SC26, November 15-20, 2026, Chicago, Illinois, USA
979-8-3195-4789-7/26/$31.00 © 2026 IEEE

munication patterns that are far more complex than those of traditional HPC applications. A single training iteration of a large language model (LLM) interleaves computation with collective communication across multiple CUDA streams, using a mix of tensor, pipeline, data, and expert parallelism [9]–[12]. The resulting execution graph is a dense directed acyclic graph (DAG) with millions of vertices whose critical path depends on the network parameters in non-obvious ways. Second, the scale of production training (over tens of thousands of GPUs) makes direct experimentation prohibitively expensive. Running a single training job for long enough to obtain stable throughput measurements at one network configuration consumes thousands of GPU-hours; sweeping over a meaningful parameter grid is infeasible. Third, packet-level network simulators such as NS-3 [13] or htsim [14], while accurate, are too slow to serve as the basis for systematic parameter studies at scale [15]–[17].

This question also cannot be answered from microbenchmarks alone. Real training executions combine communication and computation in dependency-constrained phases whose performance depends on temporal ordering, overlap, burstiness, and synchronization structure [16]. Two workloads with similar aggregate communication volume may expose different critical paths once these effects are accounted for. Even within one workload, the relevant network bottleneck can shift across phases and scale points. Synthetic kernels remain valuable controls, but they do not capture the complete behavior of end-to-end AI workloads.

On the other side, small-scale hardware experiments are expensive and do not always tell the full large-scale story. As jobs scale, the balance between exposed communication and useful computation changes, collective algorithms switch regimes (from being communication-bound to computation-bound and vice versa), and synchronization amplifies small service differences into iteration-level slowdowns [18], [19]. The resulting sensitivity landscape is inherently multidimensional. Latency affects dependency chains and exposed message startup costs. Bandwidth determines transfer time once communication is on the critical path. Jitter perturbs both average behavior and tails, which is particularly relevant in synchronous training where the slowest participant drives global progress.

Existing methods each capture only part of this landscape. Direct hardware experiments provide indispensable baselines but are limited to available fabrics and costly for broad what-if exploration. Packet-level simulation offers richer controls but becomes infeasible when sweeping multiple network dimensions across production-scale traces. Analytical or reduced-order models can be efficient, yet they often rely on abstractions that are difficult to validate on realistic AI executions [20], [21]. What is missing is a validated approach that makes large sensitivity studies feasible while preserving the workload semantics that determine the answer.

This paper addresses this gap by using realistic traces at scales of up to 4,096 GPUs to analyze how latency, bandwidth, and jitter influence end-to-end performance. The main result is a set of latency-bandwidth curves for representative AI workloads, plotted alongside current fabric technologies. These curves identify low-sensitivity regions, transition points, and configurations where runtime increases sharply under the recorded schedule and LogGP model.

To make this study feasible at production scale, we develop a compositional trace analysis methodology grounded in linear programming. Rather than building a single monolithic model for an entire execution, we decompose the workload into its constituent collective operations, analyze each unique communication structure independently, and compose the results at the program level. The method exploits a structural property of AI training: collectives within an iteration are drawn from a small set of distinct signatures, and operations on each CUDA stream execute sequentially. Composition preserves the traced stream order and measured inter-collective compute gaps. This decomposition reduces analysis costs by up to two orders of magnitude. We validate the method against runs on real hardware, monolithic linear programming (LP) solutions, and LogGOPSim [22] simulations on the studied configurations.

The primary contributions of this work are as follows:

1) We provide a large-scale, trace-grounded characterization of AI network sensitivity across realistic workloads, fabrics, and scales up to 4,096 GPUs. The complete dataset contains more than 2 TB of trace data, and the trace inputs used in this study are publicly available. We separate network latency, bandwidth, and jitter, and identify the regions where each dominates as the bottleneck.
2) We develop and validate a methodology that makes these studies feasible at production size. It combines compositional, linear programming based, trace analysis with signature-based caching and is designed to preserve the regime distinctions and sensitivity classifications that determine provisioning decisions.
3) We derive latency-bandwidth performance envelopes that expose low-sensitivity regions and performance cliffs. These envelopes provide model-based provisioning guidance and enable direct comparison with modern fabrics.

II. Background

This section introduces the concepts and tools that underpin the methodology presented in Section III.

A. NCCL and Collective Communication in AI Training

Modern distributed AI workloads rely on collective communication to synchronize gradients, broadcast parameters, and redistribute activations across GPUs. On NVIDIA hardware, these collectives are typically implemented by the NVIDIA Collective Communications Library (NCCL) [23], [24], which provides optimized routines for collectives and point-to-point operations.

NCCL divides the execution of collectives into multiple *channels*, which execute the collective's algorithm (typically Ring or Tree) in parallel. Within a single GPU kernel launch, channels map to distinct thread blocks, and, depending on the hardware topology, they can either share a single physical

NIC or be routed across multiple NICs to maximize inter-node bandwidth. Within a channel, communication proceeds in a pipelined fashion: data is divided into *slices*, and consecutive slices overlap their send, receive, and optional reduce operations across pipeline stages. The slice size, number of pipeline steps, and protocol (Simple, LL, or LL128) are determined at runtime based on the message size and hardware topology [24].

A training iteration typically issues hundreds of collective calls. For instance, in data-parallel training with the distributed optimizer [11], each backward-pass layer produces a ReduceScatter of gradients followed by an AllGather of updated parameters, yielding $O(n_{\text{layers}})$ collective invocations per iteration. These collectives are launched on CUDA streams, and operations within a single stream execute sequentially. Different streams may execute concurrently, enabling overlap between communication and computation.

B. Execution Traces and the GOAL Format

Analyzing the network sensitivity of a full training iteration requires a representation that captures not only the communication operations but also their dependencies, timing, and concurrency structure. We build on the tracing and GOAL [25] generation pipeline introduced in ATLAHS [16].

Starting from NVIDIA Nsight Systems [26] profiling traces of a real training run, the generator reconstructs each NCCL collective by parsing NVTX markers, kernel metadata, and NCCL runtime logs. For each collective instance, the generator identifies the algorithm (Ring [27] or Tree [28]), protocol (Simple, LL, LL128), number of channels, data size, slice configuration, and the actual ring or tree topology for each channel. It then emits a GOAL [25] schedule, represented as a directed acyclic graph (DAG) of *send*, *recv*, and *calc* operations connected by dependency edges. Operations are assigned to the recorded streams and NIC paths.

The GOAL format was originally designed for MPI programs [22], [29]. ATLAHS applied it to NCCL collectives by mapping operations on each CUDA stream to a distinct GOAL compute stream. This preserves the parallelism and inter-channel NIC sharing that characterize multi-channel collectives. Sends from multiple streams are serialized when they share the same modeled scale-out injection path. Computation costs for reduction and copy primitives are derived from per-operation hardware measurements collected with NPKit [30] at the target cluster and from event timestamps in the traces.

C. The LogGP Model

The LogGP model [31], [32] characterizes point-to-point communication using a small set of parameters: the *latency L* (time for the first byte to traverse the network), the *gap g* (minimum interval between consecutive message injections at the NIC), the *gap per byte G* (the inverse of per-link bandwidth, determining transfer time for long messages), the *overhead o* (fixed CPU cost per send or receive operation), and the number of processors P. The time to deliver a message of m bytes from send initiation to receive completion is:

$$T_{\text{msg}} = L + (m-1)\cdot G + 2o \tag{1}$$

LogGOPSim (LGS) [22] is a message-level simulator that replays GOAL schedules under the LogGP model, and reports the runtime of the full schedule. LGS is one of the fastest available trace simulators (minutes for traces with millions of operations) and has been validated for MPI workloads on HPC clusters [17], [18], [29]. We use LGS as a validation baseline for the LP-based predictions.

D. LP-Based Sensitivity Analysis

LLAMP [17], developed for MPI HPC applications, introduced the idea of exposing LogGP model parameters as LP variables rather than fixed constants. Given a GOAL dependency graph $G = (V, E)$, the LP assigns a start-time variable s_v to every vertex $v \in V$ and encodes the LogGP timing constraints as inequalities. For a send-receive pair with message size m bytes between ranks on different nodes:

$$s_{\text{recv}} \;\geq\; s_{\text{send}} + L + m\cdot G + o \tag{2}$$

When L or G is a decision variable, minimizing the makespan $T = \max_v s_v$ yields a one-dimensional optimal value function, such as $T^*(L)$, that is piecewise affine:

$$T^*(L) = \max_{k=1,\ldots,K} \big(C_k + \lambda_k \cdot L\big) \tag{3}$$

Each piece (C_k, λ_k) corresponds to a distinct critical path. The slope λ_k counts the number of inter-node messages on that path and represents the *latency sensitivity*: the additional runtime per unit of added latency. Analogous formulations hold when the bandwidth parameter G is the decision variable, yielding bandwidth sensitivity.

The piecewise-affine structure follows from LP duality: the optimal value of a parametric LP is a convex piecewise-affine function of its right-hand-side parameters. For a one-dimensional sweep, the active slopes can be obtained from the LP dual variables. To visualize the combined parameter space, we evaluate the independently derived latency and bandwidth envelopes on a common finite grid and classify the dominant effect relative to the baseline. These maps summarize the separate one-dimensional sweeps; they are not solutions of a joint parametric LP.

a) Limitations of the existing approach: LLAMP was effective in performing sensitivity analysis. However, it was designed for MPI workloads on traditional HPC workloads, where each rank typically corresponds to a single CPU process on a single node. Applying LP-based analysis to NCCL workloads on multi-GPU clusters raises three challenges that the LLAMP toolchain does not address:

1) **Scale.** A full training iteration produces dependency graphs with billions of vertices which would imply billions of constraints. The resulting monolithic LP exceeds the practical limits of current solvers for the kind of parametric

sweeps needed for sensitivity studies (hundreds of parameter configurations across multiple network dimensions).

2) **NIC injection serialization.** NCCL multi-channel collectives simultaneously inject messages from multiple channels using a shared NIC. This results in contention, which is not present in single-process-per-node MPI traces, where each rank typically has dedicated NIC access. Without explicit modeling of NIC serialization, the LP underestimates the bandwidth-bound component of runtime.
3) **Scale-up vs. scale-out communication.** A node typically hosts multiple GPUs (e.g., four or eight) connected by high-bandwidth scale-up interconnects such as NVLink or InfinityFabric, with latencies of a few hundred nanoseconds and bandwidths exceeding 600 GB/s. Inter-node transfers traverse the scale-out network (e.g., InfiniBand [33], Slingshot [34], or RoCE) with latencies instead in the microsecond range and per-link bandwidths of 25–200 GB/s [35]. The network parameters under study (L, G) only model the scale-out network. An LP formulation that treats all communication uniformly will conflate the two regimes, producing sensitivity estimates that do not reflect the actual network dependence. In this work, we extend the approach proposed by LLAMP and differentiate between scale-up and scale-out networks (both during tracing and in the LP formulation).

Section III describes how we address all three challenges.

III. Methodology

We present a compositional approach, called *Composite LP*, to network sensitivity analysis of distributed AI workloads. Given an execution trace of NCCL collective operations, our method produces one-dimensional piecewise-affine functions $T(L)$ and $T(G)$ that characterize how the program's end-to-end runtime responds to changes in network latency L or bandwidth parameter G. We combine these one-dimensional envelopes on a finite grid to visualize the latency-bandwidth parameter space.

A. From Traces to Per-Collective LPs

For each unique collective signature σ (described in Sec. III-B), we construct and solve an LP. The construction proceeds in three steps.

a) Step 1: GOAL generation: From the signature parameters, we generate a standalone GOAL file that encodes the collective's communication pattern. For instance, for a Ring AllReduce on n ranks with c channels, this includes the reduce-scatter phase ($n-1$ steps of receive-reduce-send per channel) followed by the allgather phase ($n-1$ steps of receive-copy-send per channel), with explicit per-slice pipelining and inter-channel parallelism on distinct compute streams.

b) Fixed NCCL schedule: The analysis replays the NCCL algorithm, protocol, channel count, chunking, and channel topology recorded in the trace. These choices remain fixed while L or G is varied. The method therefore isolates network-parameter sensitivity for an observed schedule; it does not model NCCL selecting a different schedule for another fabric.

c) Step 2: Dependency graph and LP: The GOAL file is then parsed into a dependency graph, and the graph is converted to an LP following the formulation of Equation (2). We extend the LP formulation from Section II-D with two refinements required for multi-GPU clusters:

- **Scale-up/Scale-out distinction.** For rank pairs running on GPUs on the same node (connected via a scale-up network), the latency and bandwidth parameters are fixed constants ($L_{\text{scale-up}}$, $G_{\text{scale-up}}$) rather than LP variables. This ensures that the sensitivity analysis reflects changes to the scale-out fabric only, without conflating scale-up transfers that are unaffected by network reconfiguration.
- **NIC injection serialization.** In multi-channel collectives, sends from the same source rank to the same destination may share a physical NIC injection path. We model this contention by introducing ordering constraints between consecutive sends in each per-destination chain:

$$s_{\text{send}_{j+1}}^{(r,d)} \;\geq\; s_{\text{send}_j}^{(r,d)} + m_j \cdot G + o \tag{4}$$

  where $s_{\text{send}_j}^{(r,d)}$ is the j-th send from rank r to destination d, and m_j is the message size. Without these constraints, the LP would allow simultaneous injection of all channels' messages and underestimate the bandwidth-bound runtime component.

d) Step 3: Parametric sweep and piece extraction: We sweep the decision variables (e.g., L from 0 to $L_{\max}$) in discrete steps. At each point, the LP is re-optimized using dual simplex warm-starting: after tightening the lower bound on L, the previous basis remains dual feasible, and re-optimization requires a much smaller number of pivots. From each solve, we extract the reduced cost of the selected decision variable, which gives the slope of the currently active piece. The full set of pieces $\{(\lambda_k, C_k)\}$ is collected across the sweep in a single pass, without requiring separate sensitivity analysis.

When bandwidth sensitivity is analyzed, we leave G as an LP decision variable instead of L and sweep G from a high-bandwidth to a low-bandwidth regime. For the latency-bandwidth maps, we evaluate the separate $T(L)$ and $T(G)$ envelopes on a common two-dimensional grid. At each point, we compare the changes relative to the baseline and classify whether latency, bandwidth, or both have a substantial effect. The maps are derived summaries of the one-dimensional sweeps, rather than a joint two-parameter LP or an exact enumeration of parametric regions.

B. Decomposition via Collective Signatures

The per-collective LP described above is exact for a single collective but would need to be repeated for each of the hundreds of collectives in a training iteration. We observe that these collectives are drawn from a small set of *collective signatures*. A signature σ is defined by the tuple:

$$\sigma = \big(\mathit{type},\ \mathit{algo},\ \mathit{proto},\ d,\ n,\ c,\ b,\ \tau\big) \tag{5}$$

where $\mathit{type} \in \{\text{AllReduce}, \text{AllGather}, \text{ReduceScatter}, \ldots\}$ is the collective type, $\mathit{algo} \in \{\text{Ring}, \text{Tree}\}$ is the algorithm, $\mathit{proto} \in$

{Simple, LL, LL128} is the NCCL protocol, d is the data size in bytes, n is the number of ranks, c is the number of channels, b is the slice size in bytes, and τ is the ring or tree topology (the permutation defining each rank's predecessor and successor in each channel).

Two collective instances with the same signature produce identical dependency graphs and therefore identical piecewise-affine functions $T_\sigma(L)$. This property follows from the deterministic mapping between signature parameters and GOAL structure: the number of pipeline stages, message sizes, and inter-rank dependencies are fully determined by σ.

The number of unique signatures in a training iteration is bounded by the model architecture, not by the number of collective invocations. In data-parallel training, each layer produces collectives whose data size is determined by the parameter tensor shape. The algorithm, protocol, and topology are selected identically for tensors of the same size. In a Llama 8B training iteration with 202 collectives per GPU, we find only 15 unique signatures (12 Ring-Simple collectives of varying data size plus 3 negligibly small Tree-LL scalar reductions).

C. Program-Level Composition

Given the per-collective results $T_{\sigma_i}(L)$ for each operation $i = 1, \ldots, N$ in the program's execution order, plus the inter-collective compute gaps g_i measured from the trace, we formulate a lightweight *program-level LP* that composes individual collective sensitivities into a program-level $T_{\text{prog}}(L)$.

The decomposition uses two structural properties of NCCL workloads. First, collectives on each CUDA stream execute sequentially: the NCCL runtime issues collectives in program order, and each collective completes before the next one on the same stream begins. This sequential ordering means that the finish time of one collective and the start time of the next are linked by a precedence constraint. Second, the compute gaps between collectives are measured from the trace and remain fixed during a network-parameter sweep. The trace conversion obtains cross-stream constraints in a best-effort manner by merging non-overlapping regions, following ATLAHS [16]. Within each collective, the GOAL graph retains NCCL's channel-level send, receive, and compute operations, per-slice pipeline dependencies, scale-up and scale-out distinction, NIC assignment, and injection serialization.

The program LP introduces, for each stream, variables s_i (start time) and f_i (finish time) for each operation, a free variable L, and a makespan variable T:

$$f_i \geq s_i + C_k + \lambda_k \cdot L \quad \forall\, k \in \text{pieces}(\sigma_i),\ \forall\, i \in \mathcal{C} \tag{6}$$

$$f_i \geq s_i + g_i \quad \forall\, i \in \mathcal{G} \tag{7}$$

$$s_{i+1} \geq f_i \quad \forall\, i = 1, \ldots, N-1 \tag{8}$$

$$T \geq f_N \tag{9}$$

where $\mathcal{C}$ is the set of collective operations, $\mathcal{G}$ is the set of compute gaps, and $\text{pieces}(\sigma_i)$ denotes the piecewise-affine pieces extracted for collective i's signature.

The number of variables in this LP is $2N + 2$ (two per operation plus L and T), and the number of constraints is $N + \sum_{i \in \mathcal{C}} |\text{pieces}(\sigma_i)|$, which in practice amounts to a few hundred for a full training iteration.

D. Signature-Based Caching

Because the per-collective LP depends only on the signature σ, results can be cached and reused for later identical collectives and across runs. The cache key is a deterministic hash of all signature fields, so reuse is valid only when the collective type, algorithm, protocol, message size, rank count, channel count, slice size, and ring or tree topology match.

Caching provides two levels of amortization:

1) **Within a trace:** The LP solve cost is proportional to the number of unique signatures U, not the total number of collective invocations N. For Llama 8B, $U = 15$ while $N = 202$, so the per-collective solve phase touches each signature once and reuses cached results for the remaining 187 invocations.
2) **Across traces:** Different training configurations can share identical signatures for corresponding collectives. A cache populated from one trace can be reused when analyzing another trace only when all signature fields and fixed model parameters match.

Moreover, the U per-signature solves can be executed in parallel since each one operates on an independent LP instance with no shared state.

E. Complexity and Scalability

Table I
SCALING COMPARISON FOR A TRACE WITH N COLLECTIVE INSTANCES, U UNIQUE SIGNATURES, DEPENDENCY GRAPH $G = (V, E)$ FOR THE MONOLITHIC TRACE, AND PER-SIGNATURE SUBGRAPHS $G_\sigma = (V_\sigma, E_\sigma)$. COMPOSITE ENTRIES GIVE THE LARGEST INDEPENDENT SIGNATURE; ITS PROGRAM LP HAS $O(N)$ VARIABLES.

	Monolithic LP	**Composite LP**
LP variables	$O(\lvert V\rvert)$	$O(\max_\sigma \lvert V_\sigma\rvert)$
LP constraints	$O(\lvert V\rvert + \lvert E\rvert)$	$O(\max_\sigma(\lvert V_\sigma\rvert + \lvert E_\sigma\rvert))$
Memory	$O(\lvert V\rvert + \lvert E\rvert)$	$O(\max_\sigma(\lvert V_\sigma\rvert + \lvert E_\sigma\rvert))$

Table I summarizes the scaling of the monolithic and compositional approaches. The critical difference is that the monolithic approach builds a single LP over the full dependency graph $G = (V, E)$, whose size grows with the total number of collective instances, while the compositional approach solves U independent LPs, one per unique signature, each over a subgraph G_σ, and composes the results in a program LP with $O(N)$ variables. Since $U \ll N$ in the reported Llama 8B case (15 vs. 202; the ratio grows further with more training iterations), the compositional approach is faster and less memory intensive by several factors as shown in Section VI, an improvement that enables the analysis of much larger traces.

Fig. 2 illustrates the end-to-end workflow. Starting from a single GPU's perspective, the ATLAHS generator decomposes each CUDA stream into a sequence of collective operations and profiles their signatures (step ①). Unique signatures are solved

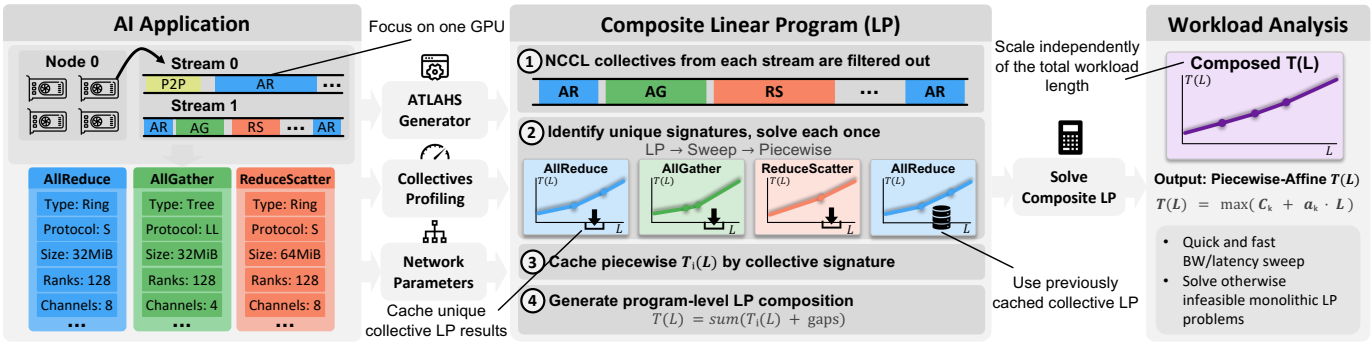

Figure 2. Compositional sensitivity analysis pipeline.

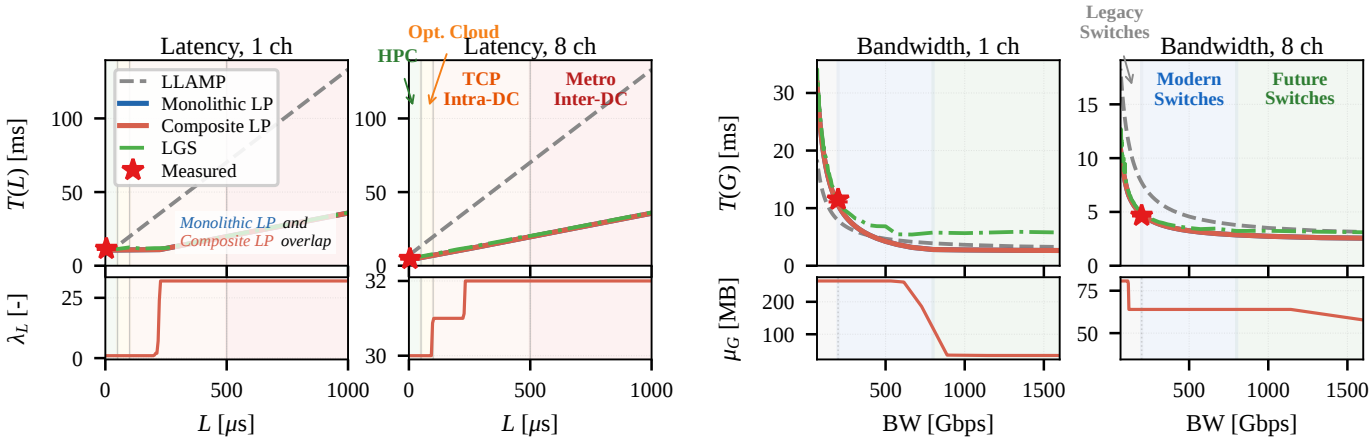

Figure 3. Single AllReduce collective of 128MiB running on 16 nodes (64 GPUs) varying the number of NCCL channels.

independently via LP sweeps (step ②), and results are cached by signature hash (step ③). The program-level Composite LP (step ④) then assembles the cached piecewise-affine functions with the traced compute gaps to produce the final $T(L)$ curve. Because the Composite LP has only $O(N)$ variables, the sweep over hundreds of latency or bandwidth configurations completes in under a second, enabling the multidimensional parameter studies presented in the following sections.

IV. Trace Dataset

A central goal of this paper is to ground the sensitivity analysis in realistic workloads. To this end, we collect a set of execution traces spanning different model architectures, parallelism strategies, hardware platforms, and scales. The trace inputs and their checksums are available at http://storage2.spcl.ethz.ch/traces/ai/.

The traces are collected on two supercomputer platforms: the Alps [3] system at CSCS (NVIDIA GH200, Slingshot interconnect [34]), which provides up to 200 Gbps of scale-out bandwidth, and a large-scale Microsoft Azure GPU cluster, named ND GB200 v6, which provides up to 800 Gbps. The Azure cluster comprises GB200 nodes, each featuring a dual-socket Grace CPU with 144 cores (72 per socket) and 1.7 TB of RAM. Each node contains four NVIDIA B200 GPUs with 185 GB of HBM3e each. GPUs are interconnected via NVLink both within a node and within racks, where each rack houses 18 nodes (72 GPUs). Each GPU is equipped with a dedicated NIC enabling GPUDirect RDMA [36] over the backend network, which is organized as a rail-optimized fat tree. The dataset comprises the following workloads:

a) Llama 3.3: Meta's Llama 3 family [8] represents the current generation of open-weight large language models. We trace Llama 3.3 training at multiple scale points using Megatron-LM [9] with data parallelism (DP), tensor parallelism (TP), and pipeline parallelism (PP).

b) Grok: The Grok workload uses xAI's 314-billion-parameter Mixture-of-Experts (MoE) architecture [6], [37], which adds expert-parallel All-to-All communication to the standard DP/TP/PP pattern. We trace Grok at the largest scale in our dataset (4,096 GPUs), producing over 2 TB of trace data containing more than 16 billion GOAL vertices and 20 billion dependencies.

c) vLLM inference (DeepSeek, Llama): To extend the study beyond training, we include inference traces collected with the vLLM serving engine [7] running DeepSeek [12] and Llama models. Inference exhibits different communication patterns: shorter bursts, smaller message sizes. It tends to be more memory-bound.

We ran the Grok traces on Azure ND GB200 v6, while all the other workloads on Alps.

V. Validation

Before using the Composite LP method for large-scale case studies, we validate it against three baselines: (1) the monolithic LP formulation applied to the full GOAL trace, (2) LogGOPSim (LGS) simulation, and (3) hardware measurements. We proceed from controlled single-collective experiments to full training iterations. For the validation runs, we use the Alps supercomputer.

Our LP-based approach produces a *lower bound* on runtime. Real networks introduce effects such as congestion, routing variability, transport state, and protocol overheads that are not captured by the LogGP model [22]. The LP prediction should therefore be interpreted as the best runtime allowed by the recorded dependency structure and modeled parameters, not as a guarantee for a deployed fabric. A gap from hardware or a more detailed simulator may include these omitted effects as well as measurement and modeling error. The LP identifies performance cliffs forced by the recorded dependency graph under the LogGP abstraction. Modeling all omitted effects at every point would make broad parameter sweeps at thousands of GPUs impractical. Our method trades this detail for the ability to evaluate many candidate network configurations while preserving the recorded program dependencies.

A. Single Collective

We begin with a single 128 MiB AllReduce collective on 64 GPUs across 16 nodes (4 GPUs per node), using the Ring algorithm with NCCL's Simple protocol. We run a forced 1-channel configuration and then also use automatic channel selection, which converges to 8 channels. We do this to demonstrate the feasibility of our toolchain to adapt to different configurations. Fig. 3 shows the latency and bandwidth sensitivity curve $T(L)$ predicted by four methods: LLAMP (the original LP formulation created for HPC apps), Monolithic LP (our extended formulation), our Composite LP approach, and LGS simulation. The hardware measurement, averaged over 30 runs at the cluster's native network parameters, is shown as a reference point (star marker). In each panel, the upper plot reports the predicted runtime as a function of either latency or bandwidth, while the lower plot reports the corresponding sensitivity coefficient. $\lambda_L = \partial T/\partial L$ is dimensionless and counts the inter-node messages on the active critical path. The bandwidth sensitivity $\mu_G = \partial T/\partial G$ has units of bytes because G is the gap per byte. For readability, the bandwidth panels place μ_G against the equivalent bandwidth $B = 1/G$ in consistent units. A derivative taken directly with respect to B would instead have units of time squared per byte.

To aid interpretation, we annotate the plots with shaded regions corresponding to representative network operating regimes. For latency, we distinguish four regimes: specialized HPC and modern AI fabrics such as Slingshot and InfiniBand HDR/NDR, which typically operate below 50 μs [34], [38]; commodity cloud interconnects using RDMA, which range from roughly 50 to 100 μs [39]; TCP-based protocols, which introduce software stack overhead and fall in the 100 to 500 μs range [40]; and metro or inter-datacenter links, which exceed 500 μs due to propagation delay [41], [42]. For bandwidth, we mark three regimes: current-generation switch bandwidth (200 to 800 Gbps), legacy bandwidth (below 200 Gbps), and next-generation or future bandwidth (above 800 Gbps) [38]. These boundaries are approximate and serve as visual anchors rather than strict thresholds; we use a consistent color scheme across all figures for ease of cross-referencing.

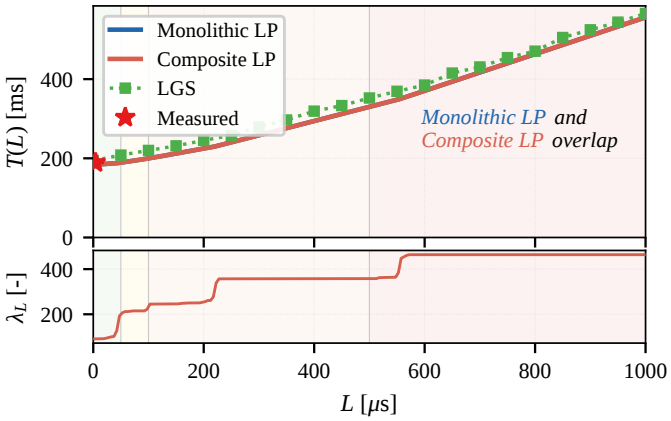

Figure 4. Predicted total runtime (top) and latency sensitivity (bottom) of a workload of 20 randomly selected NCCL collectives on 16 nodes (64 GPUs); the Composite LP matches the Monolithic LP result.

Three observations stand out. First, the Monolithic LP and Composite LP curves are indistinguishable, confirming that the piecewise extraction faithfully represents the full LP solution in this experiment. Second, LGS and the Monolithic LP closely match throughout the latency sweep, providing a check against an independent simulator. Third, LLAMP underestimates runtime (15% error at $T(0)$) and overestimates latency sensitivity ($\lambda = 15$ vs. $\lambda = 4$), because it treats all communication uniformly without distinguishing scale-up and scale-out networks or modeling NIC serialization and NCCL streams. This result shows why the extensions described in Section III-A are needed for the studied multi-GPU NCCL collective.

B. Mixed Collectives

To further validate the composition methodology, we construct a workload of 20 randomly selected NCCL collectives with data sizes drawn uniformly from 16 to 64 MiB separated by random compute gaps, executed on 64 GPUs across 16 nodes. Fig. 4 compares the Monolithic LP, the Composite LP, LGS, and the hardware wall-clock time. We note that the collectives are separated by a random delay that we intentionally keep short (below 10% of the total runtime) to clearly analyze the communication patterns.

The Composite LP matches the Monolithic within 0.8% at $T(0)$ and reproduces the same latency sensitivity ($\lambda = 104$). This confirms that the per-collective decomposition and program-level recomposition preserve the critical-path structure of the full trace in this experiment. LGS reports a higher $T(0)$ (5.3% above the Monolithic LP), consistent with its stricter NIC serialization model: LGS enforces a single NIC queue per rank across all

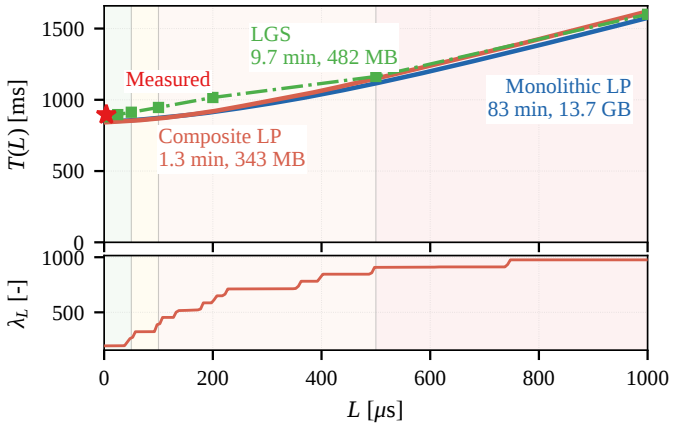

Figure 5. Llama 8B training iteration (DP16, 4 nodes, 16 GPUs). Top: runtime $T(L)$; bottom: latency sensitivity λ_L. The Composite LP matches the Monolithic LP while reducing solve time from 83 minutes to under 2 minutes.

destinations, while the LP serializes sends only within each per-destination channel group. At the measured hardware point, the LP result is closer to the observed runtime.

C. End-to-End Training: Llama 8B

We now scale the validation to a full Llama 8B training iteration with 16-way data parallelism (4 nodes, 16 GPUs in total), comprising 202 collective operations per GPU. Fig. 5 compares the Composite LP approach against the Monolithic LP.

The results demonstrate both accuracy and scalability. At each value where the Monolithic LP is feasible, the Composite prediction matches within 0.7% at $T(0)$. The full Monolithic LP (2.3 million variables, 83 minutes solve time, 14 GB of memory) confirms the Composite prediction of $T(0) = 841.8$ ms vs. the Monolithic's 847.6 ms.

The scalability advantage is already significant at this small scale. Composite LP solves all 202 collectives in 45 seconds (77 s including cold-start LP solves for all 12 unique Ring-Simple signatures) using at most 343 MB of memory. The monolithic approach requires 83 minutes and 14 GB. If the 12 per-signature LP solves are run in parallel, with 8 cores, the solve phase completes even faster in just 22 seconds.

D. Validation Across Scale

We also evaluate Composite LP and LGS across nine Grok scale points, from 16 to 4,096 GPUs. At 4 μs, in line with the inter-node latency on Alps, Fig. 8 lets us validate the model at different GPU counts against both LGS and hardware. Across the eight scale points whose exports contain all ranks, Composite LP and LGS have mean errors of about 9% and 3%, respectively. At 4,096 GPUs, both remain within 10% of the available hardware reference.

The remaining panels compare Composite LP and LGS for higher hypothetical latencies. Raising inter-node latency from 4 μs to 1 ms changes the predicted runtime by less than 0.3% at 16 GPUs. At 4,096 GPUs, the Composite LP prediction rises from about 11 to 31 s, while LGS rises from about 12 to 33 s. The two methods therefore show the same increase in latency sensitivity as the job grows.

VI. Evaluation

Having validated the compositional methodology on the configurations in Section V, we now apply it to the trace dataset introduced in Section IV. At production scale, the Monolithic LP is infeasible, and repeated LGS runs across broad parameter sweeps are costly.

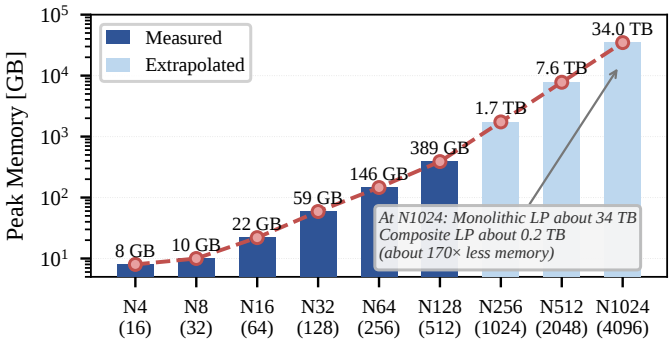

Figure 6. Monolithic LP's peak memory footprint for Grok as a function of GPU count. Dark blue bars are measured; light blue bars are extrapolated. At 4,096 GPUs the monolithic formulation would require over 34 TB.

Fig. 6 quantifies this limitation for Grok. We measure the peak memory footprint of the Monolithic LP at scale points from 4 to 128 nodes (16 to 512 GPUs), and extrapolate the larger configurations starting at 1,024 GPUs through 4,096 GPUs. Memory footprint grows superlinearly with GPU count: 8 GB at 16 GPUs, 59 GB at 128, 146 GB at 256, and 389 GB at 512 GPUs. Extrapolating to the full 4,096-GPU configuration yields an estimated 34 TB. The Composite LP, by contrast, solves each collective signature independently; its memory footprint is bounded by the largest individual collective and peaks at around 200 GB for 4,096 GPUs. Composite LP takes around 4 hours for this case, while we estimate that Monolithic LP would take several weeks. For this reason, all results in this section use Composite LP, with LGS simulation and hardware measurements as cross-checks where feasible.

A. Latency and Bandwidth Sensitivity

Fig. 7 presents the latency and bandwidth sensitivity profiles for three representative workloads, spanning AI training and inference: Llama 70B training on 128 GPUs, Grok training on 4,096 GPUs, and vLLM Llama 8B inference on 8 GPUs. The top row sweeps inter-node latency L from 0 to 1 ms; the bottom row sweeps bandwidth from 10 to 1,600 Gbps.

a) Latency: The two training workloads respond to latency inflation in qualitatively different ways. Llama 70B exhibits a piecewise structure with two key jump points near $L = 100\,\mu$s and $L = 200\,\mu$s: below the smaller threshold, the runtime is nearly flat, while above it, latency begins to dominate. The λ_L subplot confirms this, rising from near zero to approximately 2,500 and then 5,000 as latency grows. This transition reflects the

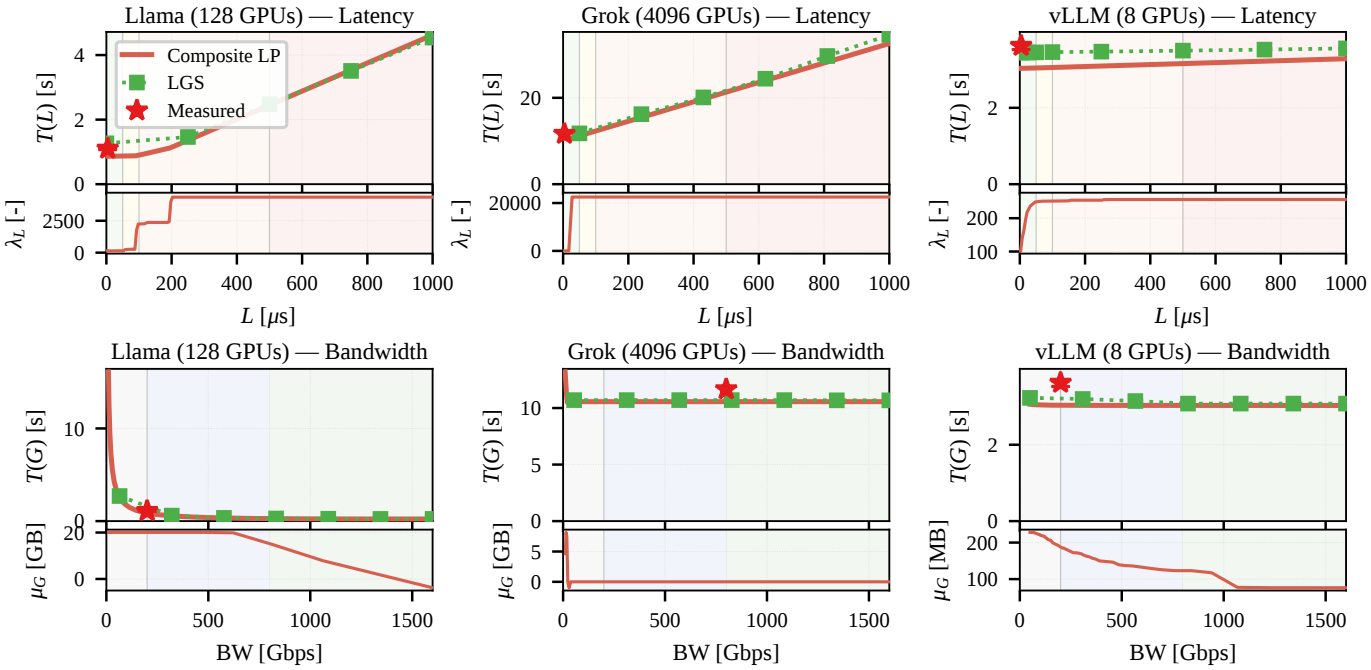

Figure 7. Model-derived latency (top) and bandwidth (bottom) sensitivity for Llama 70B (128 GPUs), Grok (4,096 GPUs), and vLLM Llama-3.1-8B (8 GPUs). Each panel shows the predicted runtime and corresponding sensitivity coefficient (λ_L for latency, μ_G for bandwidth). Available hardware measurements and LGS predictions are overlaid. The bandwidth hardware points use each system's native scale-out rate: 200 Gbps on Alps for Llama and vLLM, and 800 Gbps on Azure ND GB200 v6 for Grok.

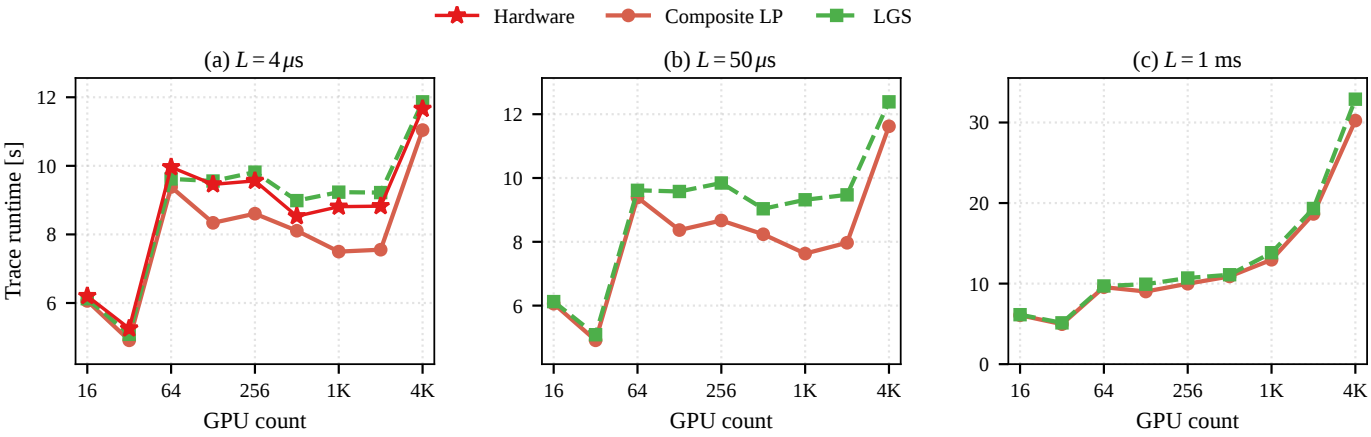

Figure 8. Grok trace runtime from 16 to 4,096 GPUs, with four GPUs per node, at inter-node latencies of (a) 4 μs, (b) 50 μs, and (c) 1 ms. The measured runtime at the native hardware settings is shown in (a).

point at which inter-node communication displaces computation or bandwidth from the critical path.

Grok at 4,096 GPUs shows only a small low-sensitivity region below 50 μs. From that point, its λ_L is approximately 22,500 across the rest of the sweep, meaning that latency is on the critical path for most evaluated points. At 1,024 nodes, ring traversals cross enough inter-node hops that communication becomes exposed even at relatively low latency values.

For vLLM, latency sensitivity similarly rises to $\lambda_L \sim 60$ for $L \lesssim 50\,\mu$s, beyond which it stays relatively flat. However, the overall effect of latency variation on end-to-end runtime is limited compared with the training workloads investigated here.

b) Bandwidth: Bandwidth sensitivity tells a different story. The Llama 70B bandwidth curve reveals a sharp transition below 200 Gbps, where reducing bandwidth causes runtime to increase rapidly. Above 800 Gbps, additional bandwidth yields diminishing returns, as computation and latency-bound paths dominate the critical path. The bandwidth sensitivity, μ_G, confirms this structure. This study keeps compute performance fixed while varying network latency and bandwidth. Faster compute subsystems could expose more communication and therefore change the sensitivity envelope; these sweeps should not be read as predictions for future accelerators.

Grok, on the other hand, exhibits near-zero bandwidth

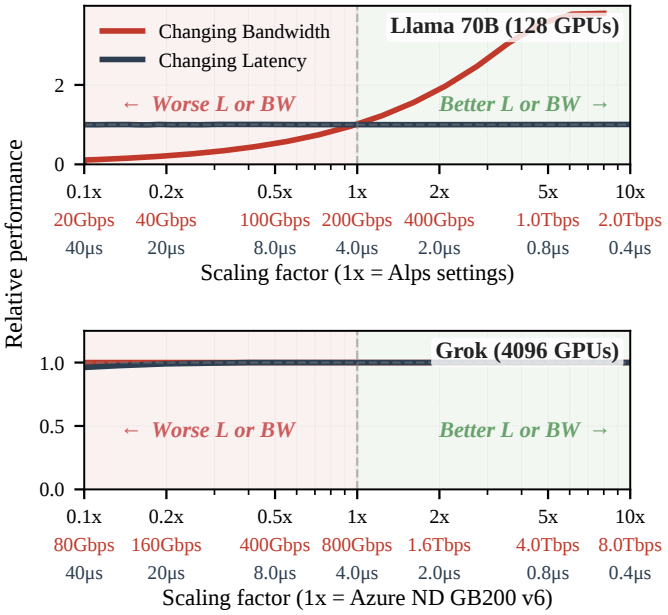

Figure 9. Latency and bandwidth sensitivity for Llama 70B and Grok focusing on AI/HPC specialized cluster network parameters.

sensitivity, as its runtime is largely driven by compute costs. For inference and vLLM 8B, we observe a similar trend: bandwidth matters little, and the critical path carries only a very small byte volume.

c) Current Data Center Networks: However, the latency range swept in Fig. 7 extends well beyond what modern datacenter fabrics deliver. Production Slingshot and InfiniBand networks typically operate in the 1–50 μs latency range at 200–800 Gbps bandwidth. When the analysis is restricted to this window, the picture changes considerably, as shown in Fig. 9. For both Llama and Grok, latency variations within this range have negligible effect on runtime. Bandwidth sensitivity, however, diverges: Grok remains largely insensitive, whereas Llama exhibits high sensitivity, particularly to bandwidth reductions. At 20 Gbps, Llama's iteration time increases by nearly 30× relative to the baseline.

In Figs. 9 and 10, relative performance is $P_{\text{rel}} = T_{\text{base}}/T$, where T_{base} is the modeled runtime at the corresponding native bandwidth or latency baseline. Thus, $P_{\text{rel}} = 1$ matches the baseline and $P_{\text{rel}} = 0.01$ denotes a 100× slowdown.

Within the latency and bandwidth ranges characteristic of current data center fabrics, bandwidth is the dominant network parameter for the studied Llama configuration. In the model, latency variation within the low-microsecond range has little effect for these workloads. The wider sweep also shows that latency becomes more important for configurations that cross metropolitan or inter-data-center distances. These observations are workload-specific and should be evaluated together with the assumptions in Section VI-D.

B. Cost-Performance Tradeoff

The bandwidth sensitivity results do not show how the performance gains compare with the additional network cost. To provide this comparison, Fig. 10 maps bandwidth operating points onto estimated switch and network-adapter costs for a fixed 4,096-GPU fat-tree deployment. The topology, radix, NIC count, and GPU configuration remain fixed while the per-link bandwidth changes.

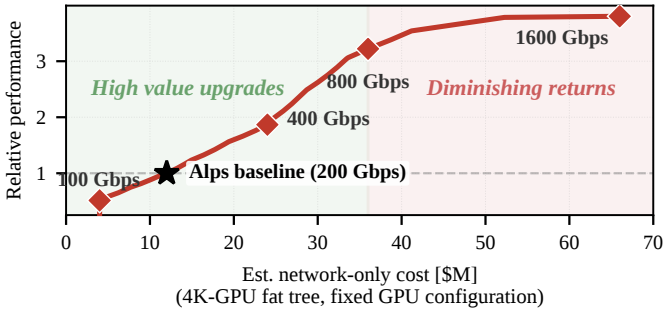

Figure 10. Estimated network-only cost versus relative performance for Llama 70B on 128 GPUs. Each point corresponds to a different bandwidth (100–1,600 Gbps) under fixed topology, radix, NIC count, and GPU configuration. The baseline marks 200 Gbps, as used by Alps in our experiments.

The plot shows a hyperbolic shape with two distinct regions. On the left, upgrading from 100 to 800 Gbps delivers large iteration-time reductions per dollar spent. The transition from 100 to 200 Gbps alone roughly halves the iteration time. On the right, upgrading from 800 to 1,600 Gbps nearly doubles the network cost while reducing iteration time by a small fraction.

For the same modeled GPU fleet, the performance gained per unit of network cost depends on the operating point relative to the workload's bandwidth breakpoint. Below that point, bandwidth upgrades produce larger modeled gains. Above it, the gains diminish.

Detailed network cost modeling is beyond the scope of this paper. Our estimates use prior work and public list prices for switches and network adapters [5], [43]–[45]. They provide a relative network-only comparison under fixed design assumptions, not a procurement quote or an estimate of total system cost.

C. Jitter

The results so far treat network parameters as deterministic: every inter-node message sees the same latency L. In production networks, however, some messages may experience elevated delays due to congestion, routing variability, or transient failures [18], [35], [46]. We use a controlled perturbation to study how delayed messages affect the modeled runtime. Table II reports minimum and maximum latency and bandwidth values from previous network characterization studies.

We use a controlled two-state model with a nominal latency of 5 μs and an elevated latency of 250 μs. For an elevated-latency fraction p, we interpolate the Composite LP predictions as $T(p) = (1-p)T(5\,\mu\text{s}) + pT(250\,\mu\text{s})$. We sweep p from 0 to 50%.

Fig. 11 shows the results for Grok. Under this two-state model, small elevated-latency fractions have a limited effect: at $p =$ 0.1% the slowdown is 0.1%, and at $p = 1\%$ it remains within

Table II
REFERENCE NETWORK LATENCY AND BANDWIDTH VALUES FROM PREVIOUS CHARACTERIZATION STUDIES.

System	Min Lat.	Max Lat. (Increase)	Max Bw.	Min Bw. (Decrease)
Alps (Slingshot 11) [35]	3.64 μs	6.48 μs (1.78×)	763 Gb/s	763 Gb/s (1×)
LUMI (Slingshot 11) [35]	2.31 μs	4.62 μs (2×)	766 Gb/s	766 Gb/s (1×)
Leonardo (IB HDR) [35]	1.9 μs	131.86 μs (69.4×)	395 Gb/s	210.11 Gb/s (1.88×)
Daint (Aries) [18]	1.192 μs	610.30 μs (512×)	75.24 Gb/s	37.62 Gb/s (2×)
Oracle (RoCEv2) [18]	1.66 μs	49.80 μs (30×)	97.53 Gb/s	2.24 Gb/s (43.5×)
AWS (SRD) [18]	16.79 μs	537.28 μs (32×)	78.74 Gb/s	39.37 Gb/s (2×)

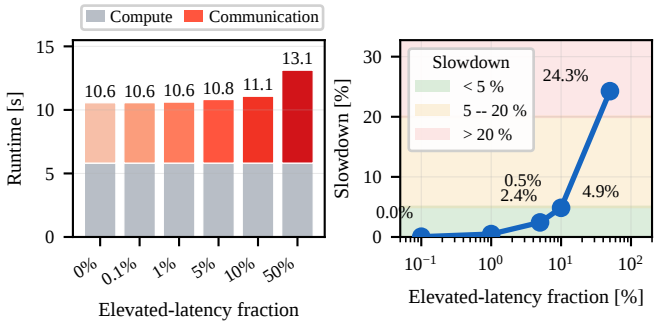

Figure 11. Impact of the two-state latency perturbation model on Grok (4,096 GPUs). (a) Iteration time as a function of the elevated-latency fraction, decomposed into compute and communication. (b) Slowdown as a function of the elevated-latency fraction.

1%. At $p = 10\%$ the slowdown reaches 5%, and at $p = 50\%$ it exceeds 24%.

These results complement the deterministic latency and bandwidth analysis with a controlled trace-based perturbation. They do not represent correlated congestion or a measured tail-latency distribution.

D. Scope and Limitations

The reported envelopes are structural lower bounds under the recorded dependency graph and the LogGP abstraction. They do not model topology-specific routing, adaptive routing, background traffic, congestion, transport state, retransmissions, or other fabric effects. These effects can increase runtime and can shift the observed breakpoints. A low-sensitivity region in our model is therefore not a guarantee of operational performance on a loaded system.

The analysis also keeps the traced NCCL schedule and compute behavior fixed. It does not retune the collective algorithm, protocol, channel count, chunking, or topology when network parameters change. Within collectives, the GOAL graph retains channel-level operations, pipeline dependencies, scale-up and scale-out transfers, and modeled NIC serialization. At program level, it preserves CUDA stream order and measured compute gaps, while cross-stream constraints are inferred in a best-effort manner. Changes in accelerator performance, overlap, or schedule selection can therefore produce different envelopes. Future work could dynamically select NCCL algorithms and protocols or add more detailed routing and transport effects. Both extensions would increase analysis time, which is the main tradeoff against the broad sweeps studied here.

ATLAHS currently relies on kernel-level activity and NVTX metadata in Nsight Systems exports. Recent Nsight Systems versions trace CUDA Graphs at graph granularity by default, where the underlying node activities are not collected [26]. This can hide kernels in inference workloads that rely heavily on CUDA Graph replay. Node-level graph tracing is available, but it can add significant profiling overhead. Supporting graph-level inference traces is an important direction for future work.

Finally, the jitter study uses a two-state interpolation and is not a model of bursty or correlated network events. Hardware validation is performed where native measurements are feasible, while the broad and large-scale parameter sweeps remain model-based. The quantitative conclusions apply to the recorded workloads, software choices, and scale points; the toolchain can be rerun for other configurations.

VII. RELATED WORK

a) Network performance modeling: The LogP [31] and LogGP [32] families capture point-to-point communication using a small set of abstract cost parameters. Building on these models, LogGOPSim [22] replays GOAL schedules and has been validated on large-scale MPI workloads [18], [29]. LLAMP [17] reformulates the LogGP constraints as a linear program, enabling analytical sensitivity and tolerance analysis for MPI applications without repeated simulation. Our work builds on the LP formulation but targets NCCL-based AI workloads and introduces a compositional decomposition that reduces analysis cost by up to two orders of magnitude.

b) AI training simulation: Simulators such as ASTRA-sim [20] and SimAI [47] model distributed training by coupling compute, memory, and network simulators, enabling what-if studies on parallelism strategies and hardware configurations. Chakra [21] proposes an open format to store execution traces of ML workloads and benchmarks and supports trace-driven simulation. Both frameworks target design-space exploration at the system level; they do not expose one-dimensional parametric network sensitivity as a piecewise-affine function. Packet-level simulators such as NS-3 [13] and htsim [14] provide high-fidelity network modeling but are too costly for the multidimensional sweeps required by our study [15]. ATLAHS [16] bridges these approaches by converting real application traces to GOAL and feeding them to multiple simulator backends; GOAL traces are 2×–11× smaller than corresponding Chakra traces [16]. We adopt the ATLAHS tracing infrastructure and extend the analysis layer with LP-based composition.

c) Network sensitivity and provisioning: Prior work on network latency sensitivity has focused on MPI workloads. Underwood et al. [48] and Patke et al. [49] inject latency on real hardware to measure application response, an approach that is effective but expensive and limited to available fabrics. Hoefler et al. [29] study the influence of system noise on HPC applications via simulation. De Sensi et al. [18] characterize network variability in cloud environments. These efforts examine individual dimensions (latency or noise) on HPC codes.

Our study examines latency, bandwidth, and jitter on large AI workload traces, derives piecewise-affine one-dimensional envelopes, and uses the latency and bandwidth envelopes to construct finite-grid sensitivity maps.

VIII. Conclusion

We presented a trace-grounded study of network sensitivity for AI workloads on supercomputers. By decomposing execution traces into per-collective LP problems and composing them at the program level, we reduced the cost of parametric sensitivity analysis by up to two orders of magnitude compared to monolithic formulations, making multidimensional sweeps over latency, bandwidth, and jitter feasible at large scale, including the 4,096 GPUs test case presented earlier.

The resulting performance envelopes show that, depending on the workload, large regions of the latency-bandwidth space produce near-identical runtime under the model and that transitions to degraded operation can be sharp. These findings provide provisioning guidance by identifying low-sensitivity regions, performance cliffs, and points of diminishing returns. They should be interpreted together with the lower-bound model assumptions and workload scope discussed in Section VI-D.

The execution traces used in this study are available at http://storage2.spcl.ethz.ch/traces/ai/. They enable the analysis to be repeated for the workloads and scale points studied in this work.

Acknowledgments

This work is supported by the European Union's Horizon Europe under grant 101175702 (NET4EXA), by Sapienza University Grants ADAGIO and D2QNeT (*Bando per la ricerca di Ateneo* 2023 and 2024), and by a research grant from Microsoft Azure. The research was also conducted as part of the FastTrackAI project at the Singapore-ETH Centre, which was established collaboratively between ETH Zurich and the National Research Foundation, Singapore. Additionally, this research is supported by the National Research Foundation, Singapore (NRF), and the Ministry of Digital Development and Information (MDDI) under the AI Visiting Professorship (Award No. AIVP-2025-005). We also thank the Swiss National Supercomputing Center (CSCS) for providing the computational resources used in this work.

ChatGPT 5.4 and Opus 4.6 were used for language editing.

References

[1] K. Wagner, "Meta is building new $800 million ai-focused data center in indiana," *Bloomberg*, 2024. [Online]. Available: https://www.bloomberg.com/news/articles/2024-01-25/meta-building-new-800-million-ai-focused-data-center-in-indiana

[2] B. Cottier, R. Rahman, L. Fattorini, N. Maslej, T. Besiroglu, and D. Owen, "The rising costs of training frontier ai models," 2024. [Online]. Available: https://doi.org/10.48550/arXiv.2405.21015

[3] CSCS, "New research infrastructure: 'alps' supercomputer inaugurated," *Swiss National Supercomputing Center*, 2024. [Online]. Available: https://www.cscs.ch/publications/news/2024/new-research-infrastructure-alps-supercomputer-inaugurated

[4] K. Qian, Y. Xi, J. Cao, J. Gao, Y. Xu, Y. Guan, B. Fu, X. Shi, F. Zhu, R. Miao, C. Wang, P. Wang, P. Zhang, X. Zeng, E. Ruan, Z. Yao, E. Zhai, and D. Cai, "Alibaba hpn: A data center network for large language model training," in *Proceedings of the ACM SIGCOMM 2024 Conference*, 2024, pp. 691–706. [Online]. Available: https://doi.org/10.1145/3651890.3672265

[5] T. Hoefler, T. Bonato, D. De Sensi, S. Di Girolamo, S. Li, M. Heddes, D. Goel, M. Castro, and S. Scott, "Hammingmesh: A network topology for large-scale deep learning," *Commun. ACM*, vol. 67, no. 12, p. 97–105, Nov. 2024. [Online]. Available: https://doi.org/10.1145/3623490

[6] xAI, "Open release of Grok-1," https://x.ai/news/grok-os, 2024, accessed: 2026-08-29.

[7] W. Kwon, Z. Li, S. Zhuang, Y. Sheng, L. Zheng, C. H. Yu, J. E. Gonzalez, H. Zhang, and I. Stoica, "Efficient memory management for large language model serving with pagedattention," in *Proceedings of the 29th Symposium on Operating Systems Principles*, 2023, pp. 611–626. [Online]. Available: https://doi.org/10.1145/3600006.3613165

[8] A. Grattafiori, A. Dubey, A. Jauhri *et al.*, "The llama 3 herd of models," 2024. [Online]. Available: https://doi.org/10.48550/arXiv.2407.21783

[9] M. Shoeybi, M. Patwary, R. Puri, P. LeGresley, J. Casper, and B. Catanzaro, "Megatron-lm: Training multi-billion parameter language models using model parallelism," *arXiv preprint arXiv:1909.08053*, 2019. [Online]. Available: https://doi.org/10.48550/arXiv.1909.08053

[10] Y. Huang, Y. Cheng, A. Bapna, O. Firat, D. Chen, M. Chen, H. Lee, J. Ngiam, Q. V. Le, Y. Wu, and Z. Chen, "Gpipe: Efficient training of giant neural networks using pipeline parallelism," in *Advances in Neural Information Processing Systems*, vol. 32, 2019. [Online]. Available: https://proceedings.neurips.cc/paper/2019/hash/093f65e080a295f8076b1c5722a46aa2-Abstract.html

[11] S. Rajbhandari, J. Rasley, O. Ruwase, and Y. He, "Zero: Memory optimizations toward training trillion parameter models," in *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, 2020, pp. 1–16. [Online]. Available: https://doi.org/10.1109/SC41405.2020.00024

[12] DeepSeek-AI, A. Liu, B. Feng *et al.*, "Deepseek-v3 technical report," 2024. [Online]. Available: https://doi.org/10.48550/arXiv.2412.19437

[13] G. F. Riley and T. R. Henderson, "The ns-3 network simulator," in *Modeling and tools for network simulation*. Springer, 2010, pp. 15–34. [Online]. Available: https://doi.org/10.1007/978-3-642-12331-3_2

[14] C. Raiciu, C. Pluntke, S. Barre, A. Greenhalgh, D. Wischik, and M. Handley, "Data center networking with multipath tcp," in *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, 2010, pp. 1–6. [Online]. Available: https://doi.org/10.1145/1868447.1868457

[15] M. Besta, M. Schneider, S. D. Girolamo, A. Singla, and T. Hoefler, "Towards million-server network simulations on just a laptop," 2021. [Online]. Available: https://doi.org/10.3929/ethz-b-000526502

[16] S. Shen, T. Bonato, Z. Hu, P. Jordan, T. Chen, and T. Hoefler, "Atlahs: An application-centric network simulator toolchain for ai, hpc, and distributed storage," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 349–367. [Online]. Available: https://doi.org/10.1145/3712285.3759838

[17] S. Shen, L. Huang, M. Chrapek, T. Schneider, J. Dayal, M. Gajbe, R. Wisniewski, and T. Hoefler, "Llamp: Assessing network latency tolerance of hpc applications with linear programming," in *SC24: International Conference for High Performance Computing, Networking, Storage and Analysis*, 2024, pp. 1–18. [Online]. Available: https://doi.org/10.1109/SC41406.2024.00070

[18] D. De Sensi, T. De Matteis, K. Taranov, S. Di Girolamo, T. Rahn, and T. Hoefler, "Noise in the clouds: Influence of network performance variability on application scalability," *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, vol. 6, no. 3, pp. 1–27, 2022. [Online]. Available: https://doi.org/10.1145/3570609

[19] J. Dean and L. A. Barroso, "The tail at scale," *Commun. ACM*, vol. 56, no. 2, p. 74–80, Feb. 2013. [Online]. Available: https://doi.org/10.1145/2408776.2408794

[20] W. Won, T. Heo, S. Rashidi, S. Sridharan, S. Srinivasan, and T. 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)*. Los Alamitos, CA, USA: IEEE Computer Society, Apr. 2023, pp. 283–294. [Online]. Available: https://doi.org/10.1109/ISPASS57527.2023.00035

[21] S. Sridharan, T. Heo, L. Feng, Z. Wang, M. Bergeron, W. Fu, S. Zheng, B. Coutinho, S. Rashidi, C. Man, and T. Krishna, "Chakra: Advancing performance benchmarking and co-design using standardized execution traces," 2023. [Online]. Available: https://doi.org/10.48550/arXiv.2305.14516

[22] T. Hoefler, T. Schneider, and A. Lumsdaine, "Loggopsim: simulating large-scale applications in the loggops model," in *Proceedings of the*

19th ACM International Symposium on High Performance Distributed Computing, ser. HPDC '10. New York, NY, USA: Association for Computing Machinery, 2010, p. 597–604. [Online]. Available: https://doi.org/10.1145/1851476.1851564

[23] NVIDIA, "NCCL: NVIDIA collective communications library," https://developer.nvidia.com/nccl, 2024, accessed: 2025-03-15.

[24] Z. Hu, S. Shen, T. Bonato, S. Jeaugey, C. Alexander, E. Spada, J. Dinan, J. Hammond, and T. Hoefler, " Demystifying NCCL: An In-Depth Analysis of GPU Communication Protocols and Algorithms ," in *2025 IEEE Symposium on High-Performance Interconnects (HOTI)*. Los Alamitos, CA, USA: IEEE Computer Society, Aug. 2025, pp. 48–59. [Online]. Available: https://doi.org/10.1109/HOTI66940.2025.00024

[25] T. Hoefler, C. Siebert, and A. Lumsdaine, "Group operation assembly language - a flexible way to express collective communication," in *Proceedings of the 2009 International Conference on Parallel Processing*, ser. ICPP '09. USA: IEEE Computer Society, 2009, p. 574–581. [Online]. Available: https://doi.org/10.1109/ICPP.2009.70

[26] NVIDIA, "NVIDIA nsight systems user guide," https://docs.nvidia.com/nsight-systems/UserGuide/index.html, 2026, accessed: 2026-08-27.

[27] P. Patarasuk and X. Yuan, "Bandwidth optimal all-reduce algorithms for clusters of workstations," *J. Parallel Distrib. Comput.*, vol. 69, no. 2, p. 117–124, Feb. 2009. [Online]. Available: https://doi.org/10.1016/j.jpdc.2008.09.002

[28] M. Barnett, R. Littlefield, D. Payne, and R. Vandegeijn, "Global combine algorithms for 2-d meshes with wormhole routing," *Journal of Parallel and Distributed Computing*, vol. 24, no. 2, pp. 191–201, 1995. [Online]. Available: https://doi.org/10.1006/jpdc.1995.1018

[29] T. Hoefler, T. Schneider, and A. Lumsdaine, "Characterizing the influence of system noise on large-scale applications by simulation," in *Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '10. USA: IEEE Computer Society, 2010, p. 1–11. [Online]. Available: https://doi.org/10.1109/SC.2010.12

[30] Microsoft, "Npkit: Nccl profiling kit," https://github.com/microsoft/NPKit, 2021, accessed: 2026-04-08.

[31] D. Culler, R. Karp, D. Patterson, A. Sahay, K. E. Schauser, E. Santos, R. Subramonian, and T. von Eicken, "Logp: towards a realistic model of parallel computation," in *Proceedings of the Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, ser. PPOPP '93. New York, NY, USA: Association for Computing Machinery, 1993, p. 1–12. [Online]. Available: https://doi.org/10.1145/155332.155333

[32] A. Alexandrov, M. F. Ionescu, K. E. Schauser, and C. Scheiman, "Loggp: Incorporating long messages into the logp model for parallel computation," *Journal of Parallel and Distributed Computing*, vol. 44, no. 1, pp. 71–79, 1997. [Online]. Available: https://doi.org/10.1006/jpdc.1997.1346

[33] Mellanox Technologies, "Introducing 200G HDR InfiniBand Solutions," 2019, [Accessed 29-03-2024]. [Online]. Available: https://network.nvidia.com/sites/default/files/pdf/whitepapers/WP_Introducing_200G_HDR_InfiniBand_Solutions.pdf

[34] D. De Sensi, S. Di Girolamo, K. H. McMahon, D. Roweth, and T. Hoefler, "An in-depth analysis of the slingshot interconnect," in *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, 2020, pp. 1–14. [Online]. Available: https://doi.org/10.1109/SC41405.2020.00039

[35] D. De Sensi, L. Pichetti, F. Vella, T. De Matteis, Z. Ren, L. Fusco, M. Turisini, D. Cesarini, K. Lust, A. Trivedi, D. Roweth, F. Spiga, S. Di Girolamo, and T. Hoefler, "Exploring gpu-to-gpu communication: Insights into supercomputer interconnects," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'24)*, Nov. 2024, pp. 1–15. [Online]. Available: https://doi.org/10.1109/SC41406.2024.00039

[36] S. Potluri, K. Hamidouche, A. Venkatesh, D. Bureddy, and D. K. Panda, "Efficient inter-node mpi communication using gpudirect rdma for infiniband clusters with nvidia gpus," in *Proceedings of the 2013 42nd International Conference on Parallel Processing*, ser. ICPP '13. USA: IEEE Computer Society, 2013, p. 80–89. [Online]. Available: https://doi.org/10.1109/ICPP.2013.17

[37] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, Q. Le, G. Hinton, and J. Dean, "Outrageously large neural networks: The sparsely-gated mixture-of-experts layer," 2017. [Online]. Available: https://doi.org/10.48550/arXiv.1701.06538

[38] T. Hoefler, D. Roweth, K. Underwood, R. Alverson, M. Griswold, V. Tabatabaee, M. Kalkunte, S. Anubolu, S. Shen, M. McLaren, A. Kabbani, and S. Scott, "Data center ethernet and remote direct memory access: Issues at hyperscale," *Computer*, vol. 56, no. 7, p. 67–77, Jul. 2023. [Online]. Available: https://doi.org/10.1109/MC.2023.3261184

[39] Y. Zhu, H. Eran, D. Firestone, C. Guo, M. Lipshteyn, Y. Liron, J. Padhye, S. Raindel, M. H. Yahia, and M. Zhang, "Congestion control for large-scale rdma deployments," in *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, ser. SIGCOMM '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 523–536. [Online]. Available: https://doi.org/10.1145/2785956.2787484

[40] S. Peter, J. Li, I. Zhang, D. R. K. Ports, D. Woos, A. Krishnamurthy, T. Anderson, and T. Roscoe, "Arrakis: the operating system is the control plane," in *Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation*, ser. OSDI '14. USA: USENIX Association, 2014, p. 1–16. [Online]. Available: https://www.usenix.org/conference/osdi14/technical-sessions/presentation/peter

[41] T. Bonato, S. Abdous, A. Kabbani, A. Ghalayini, N. Gebara, T. Lam, A. Agarwal, T. Chen, Z. Yu, K. Taranov, M. Elhaddad, D. De Sensi, S. Ghorbani, and T. Hoefler, "Uno: A one-stop solution for inter- and intra-data center congestion control and reliable connectivity," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 1195–1210. [Online]. Available: https://doi.org/10.1145/3712285.3759884

[42] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat, "B4: experience with a globally-deployed software defined wan," in *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM*, ser. SIGCOMM '13. New York, NY, USA: Association for Computing Machinery, 2013, p. 3–14. [Online]. Available: https://doi.org/10.1145/2486001.2486019

[43] M. Besta and T. Hoefler, "Slim fly: a cost effective low-diameter network topology," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '14. IEEE Press, 2014, p. 348–359. [Online]. Available: https://doi.org/10.1109/SC.2014.34

[44] FS.com, "MQM9700-NS2F, NVIDIA® 64-Port NDR 400G InfiniBand Data Center Switch," 2026, accessed: 2026-04-09. [Online]. Available: https://www.fs.com/eu-en/products/194710.html

[45] Dell USA, "NVIDIA ConnectX-7 NDR 400G InfiniBand Adapter Card," 2026, accessed: 2026-04-09. [Online]. Available: https://www.dell.com/en-us/shop/nvidia-connectx-7-ndr-400g-infiniband-adapter-card/apd/ac478006/computer-chassis-components

[46] D. De Sensi, S. Di Girolamo, and T. Hoefler, "Mitigating network noise on dragonfly networks through application-aware routing," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '19. New York, NY, USA: ACM, 2019, pp. 1–32. [Online]. Available: https://doi.org/10.1145/3295500.3356196

[47] X. Wang, Q. Li, Y. Xu, G. Lu, D. Li, L. Chen, H. Zhou, L. Zheng, S. Zhang, Y. Zhu, Y. Liu, P. Zhang, K. Qian, K. He, J. Gao, E. Zhai, D. Cai, and B. Fu, "Simai: unifying architecture design and performance tuning for large-scale large language model training with scalability and precision," in *Proceedings of the 22nd USENIX Symposium on Networked Systems Design and Implementation*, ser. NSDI '25. USA: USENIX Association, 2025, pp. 541–558. [Online]. Available: https://www.usenix.org/conference/nsdi25/presentation/wang-xizheng-simai

[48] R. Underwood, J. Anderson, and A. Apon, "Measuring network latency variation impacts to high performance computing application performance," in *Proceedings of the 2018 ACM/SPEC International Conference on Performance Engineering*, ser. ICPE '18. New York, NY, USA: Association for Computing Machinery, 2018, pp. 68–79. [Online]. Available: https://doi.org/10.1145/3184407.3184427

[49] A. Patke, H. Qiu, S. Jha, S. Venugopal, M. Gazzetti, C. Pinto, Z. Kalbarczyk, and R. Iyer, "Evaluating hardware memory disaggregation under delay and contention," in *2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE, 2022, pp. 1221–1227. [Online]. Available: https://doi.org/10.1109/IPDPSW55747.2022.00210