

# LLAMP: Assessing Network Latency Tolerance of HPC Applications with Linear Programming

Siyuan Shen  
Department of Computer Science  
ETH Zürich  
Zürich, Switzerland  
siyuan.shen@inf.ethz.ch

Langwen Huang  
Department of Computer Science  
ETH Zürich  
Zürich, Switzerland  
langwen.huang@inf.ethz.ch

Marcin Chrapek  
Department of Computer Science  
ETH Zürich  
Zürich, Switzerland  
marcin.chrapek@inf.ethz.ch

Timo Schneider  
Department of Computer Science  
ETH Zürich  
Zürich, Switzerland  
timo.schneider@inf.ethz.ch

Jai Dayal  
SAIT Systems Architecture Lab  
Samsung Semiconductor  
San Jose, USA  
jai.dayal@samsung.com

Manisha Gajbe  
SAIT Systems Architecture Lab  
Samsung Semiconductor  
San Jose, USA  
m.gajbe@samsung.com

Robert Wisniewski  
Hewlett Packard Enterprise  
Ossining, USA  
bobww123@gmail.com

Torsten Hoefler  
Department of Computer Science  
ETH Zürich  
Zürich, Switzerland  
torsten.hoefler@inf.ethz.ch

**Abstract**—The shift towards high-bandwidth networks driven by AI workloads in data centers and HPC clusters has unintentionally aggravated network latency, adversely affecting the performance of communication-intensive HPC applications. As large-scale applications often exhibit significant differences in their network latency tolerance, it is crucial to determine the extent of network latency an application can withstand without significant performance degradation. Current approaches often rely on specialized hardware or simulators, which can be inflexible and time-consuming. We introduce LLAMP, a novel toolchain that offers an efficient analytical approach to evaluating HPC applications’ network latency tolerance using the LogGPS model and linear programming. Through our validation on a variety of MPI applications such as LULESH and MILC, we demonstrate our tool’s high accuracy, with relative prediction errors below 2%. Additionally, we include a case study of the ICON weather and climate model to illustrate LLAMP’s broad applicability in evaluating collective algorithms and network topologies.

**Index Terms**—Network latency tolerance, linear programming, MPI applications, high-performance computing

## I. INTRODUCTION

The growing demand for training large deep learning models has spurred the construction of advanced AI-focused data centers and supercomputing clusters, such as Meta’s \$800 million data center [1] and the upcoming Alps cluster at the Swiss National Supercomputing Center [2]. Driven by the needs of the AI industry, there have been notable improvements in cloud hardware, which makes cloud platforms increasingly appealing for running high-performance computing (HPC) applications due to the potential cost benefits over traditional on-premise HPC clusters [3], [4]. However, numerous studies [5]–[9] have shown, executing large HPC applications, particularly those

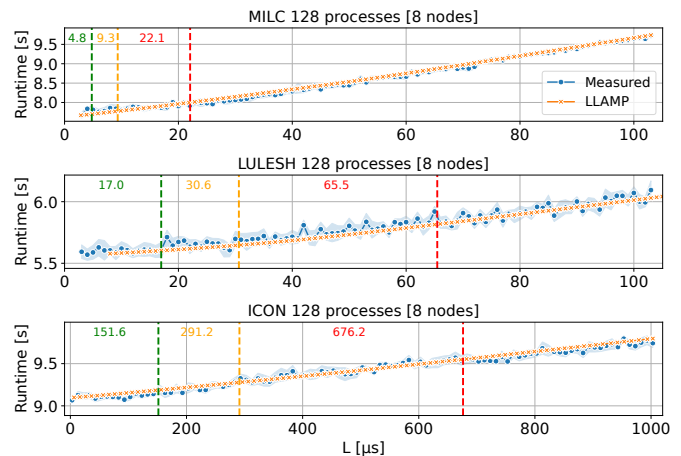


Fig. 1: An example demonstrating varying degrees of network latency tolerance among traditional HPC applications, namely MILC, LULESH, and ICON. The green, orange, and red zones correspond to the maximum network latencies before observing a performance degradation of 1%, 2%, and 5%, respectively. The comparison between measured and predicted runtime showcases the predictive accuracy of our toolchain. The tolerance intervals are calculated directly by our tool.

that are communication-intensive, can result in a suboptimal cost-performance ratio when migrated to the cloud, primarily due to increased network latency that often exceeds  $10 \mu\text{s}$  [5].

Moreover, parallel training of large deep learning models, which relies more on bandwidth than latency [9]–[13], has

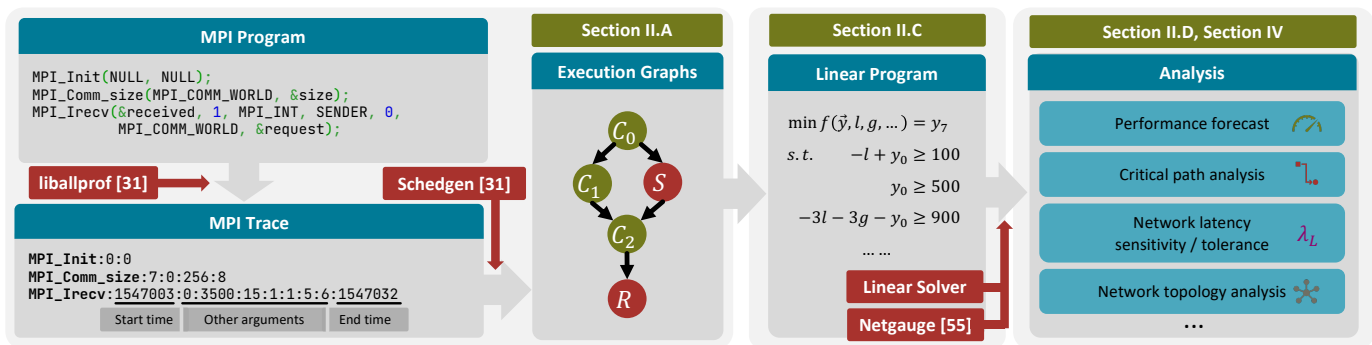


Fig. 2: High-level overview of the LLAMP toolchain.

prompted a rapid increase in network bandwidth. To accommodate this need, the network bandwidth has been growing exponentially in recent years [14]. This growth, accelerated by the adoption of higher frequencies and intricate signaling techniques like PAM4, has simultaneously led to higher bit error rates (BER) in transceivers. Consequently, this requires the implementation of complex forward error correction (FEC) mechanisms, as seen with the upcoming 800G and 1.6T IEEE P802.3df standards [15]. While existing fast FEC manages to keep latency around 50 ns, more complex FEC schemes in the future are expected to increase the decoding latency by more than 100 ns [16], [17], subsequently adding several hundred nanoseconds to the per-link latency.

With data centers and HPC clusters evolving towards high-bandwidth networks to conform with the need of the AI sector [18], [19], the trade-off between bandwidth and latency becomes more pronounced, and the implications of FEC-induced latency become a concern for the design and optimization of both HPC systems and applications. Addressing this latency issue is especially important in fields where the time to solution is crucial, such as weather forecasting simulations for climate change analysis or molecular dynamics simulations for in-depth exploration of COVID-19 at the molecular level [14].

In large-scale MPI applications, unique communication and computation patterns inherent in each application lead to significant variations in their network latency tolerance. For instance, in Fig. 1, we present the differences in network latency tolerance across three traditional HPC applications: MIMD Lattice Computation (MILC) [20], LULESH [21], and Icosahedral Nonhydrostatic Weather and Climate Model (ICON) [22]. The plots illustrate the impact of increasing network latency on each application’s runtime. MILC exhibits the lowest tolerance to network latency, indicating that as little as 20  $\mu\text{s}$  can adversely affect its performance. Conversely, ICON demonstrates the highest latency tolerance, able to withstand more than 650  $\mu\text{s}$  of network latency before performance degradation becomes apparent. These differences emphasize the importance of network configurations to suit the specific latency tolerance profiles of each application.

Given these considerations, it is imperative to determine the extent of network latency an application can withstand without significant performance degradation. Understanding this

threshold is key to designing applications that are both resilient and efficient, even under suboptimal network conditions.

However, existing methods for evaluating network latency tolerance have notable limitations. These strategies typically rely on one of three approaches, each with inherent drawbacks. The first method involves constructing elaborate performance models, demanding an in-depth understanding of the applications’ communication and computation behaviors [23], [24]. This requires extensive knowledge and expertise, making it inaccessible to many practitioners. The second strategy hinges on artificially injecting latency into network communication to observe its impact on application performance. To achieve this, one can exploit specialized hardware, which is often difficult to procure and also suffers from inflexibility [25], [26]. The third option is to rely on packet-level network simulators to predict the behaviors of applications under various conditions without the need for physical hardware. However, these simulators often require intricate configurations and can be time-consuming to execute, especially when simulating large-scale workloads [27]–[30]. A fundamental issue with the last two methods is the need for exhaustive parameter sweeps to ascertain an application’s latency tolerance.

In response to these challenges, we introduce *LLAMP* (LogGPS and Linear Programming based Analyzer for MPI Programs), a toolchain designed for efficient analysis and quantification of network latency sensitivity and tolerance in HPC applications. An overview of LLAMP is presented in Fig. 2. By leveraging the LogGPSim framework [31], LLAMP records MPI traces of MPI programs and transforms them into execution graphs. These graphs, through the use of the LogGPS model, are then converted into linear programs. They can be solved rapidly by modern linear solvers, allowing us to efficiently gather valuable metrics, such as the predicted runtime of programs, and critical path metrics. These insights are crucial for identifying an application’s network latency tolerance and its performance variability under diverse network configurations. LLAMP’s versatility is further demonstrated through a case study of the ICON climate model, exploring the effects of collective algorithms and network structures.

The source code of LLAMP is available at <https://github.com/spcl/llamp>, and an extended technical report can be accessed at <https://arxiv.org/abs/2404.14193>.

The primary contributions of this work are as follows:

- 1) We derive a novel analytical model that quantifies the network latency sensitivity and tolerance of MPI applications. This model leverages linear programming, execution graphs generated from application traces, and the LogGPS model, providing a mathematical foundation for our analysis.
- 2) We develop LLAMP, an open-source toolchain that allows us to efficiently forecast the performance of MPI applications and compute their network latency sensitivity and tolerance. This empowers architects to tailor infrastructure designs to application needs and enables software developers to make informed decisions regarding application deployment and optimization for reduced latency sensitivity.
- 3) We built a software-based latency injector capable of emulating flow-level network latency with high portability, facilitating large-scale latency injection experiments without specialized hardware or administrative privileges. Using this latency injector, we validate our model on a number of MPI applications from various domains, such as LULESH, HPCG, and MILC, among others. Furthermore, we demonstrate LLAMP’s broad applicability by conducting a case study on ICON, examining how collective algorithms and network topologies influence its performance.

## II. LLAMP TOOLCHAIN

### A. MPI Execution Graphs

In parallel computing, analyzing MPI programs through the lens of graphs proves to be a natural and insightful methodology. As MPI facilitates parallel execution by decomposing computational tasks into distinct units, dependency graphs provide a suitable representation of the relationships and communication patterns between these tasks. As exemplified in various other works [31]–[35], this approach has been validated and employed successfully.

We chose to base LLAMP on the work of Hoefer et al. [31], [35] and the reason is twofold. Firstly, the tracer and schedule generator included in the LogGOPSim toolchain allows us to easily generate execution schedules of MPI programs. Secondly, LogGOPSim stands out as a well-tested open-source framework with an accessible codebase [36], [37], facilitating extensibility and customization to suit our specific needs.

To elaborate, the schedule generator, **Schedgen**, parses the traces collected from the light-weight tracing library, **liballprof**, and converts them into execution graphs that capture the precedence of events and task dependencies in GOAL format [38]. An MPI execution graph is a directed acyclic graph (DAG) that contains three vertex types, namely *send*, *recv*, and *calc*, which denote point-to-point (p2p) sends and receives and computations, respectively. Edges connecting *send* and *recv* pairs are termed *communication edges*. By exploiting the difference in timestamps of consecutive MPI operations, Schedgen infers the amount of computation that occurred. In addition, Schedgen is able to substitute collective operations with p2p algorithms based on user specifications.

The LogGOPS model, a derivative of the LogP family [39], [40], serves to quantify the communication costs in parallel




Fig. 3: An example illustrating the transformation of blocking p2p operations into an execution graph, assuming that the eager protocol is used. (A) lists collected traces with only the start and end timestamps. (B) shows the corresponding space-time diagram. In (C), *calc* vertices are marked in green while *send* and *recv* vertices are in red.


applications. In this model,  $L$  denotes the maximum latency between two processors, which is the network latency we will focus on in this work. Parameter  $o$  represents the overhead on the CPU per message, and  $g$  is the gap between two consecutive messages.  $G$ , as introduced by Alexandrov et al. [41], models the gap per byte of a message, and it is equivalent to the inverse of bandwidth. Additionally,  $O$  is the CPU overhead per byte, and  $P$  represents the number of processes. The parameter  $S$  sets the size threshold for employing the *rendezvous protocol*, where messages smaller than  $S$  are sent immediately (i.e., the eager protocol), whereas larger messages require synchronization between sender and receiver before transmission. As highlighted by Hoefer et al. [35],  $O$  is commonly negligible due to high overlap. Therefore, we adopted the LogGPS model for our analysis.

In panel (A) of Fig. 3, we present the trace of a simple MPI program. Moving to panel (B), the cost of *calc* vertices are inferred from the start and end timestamps for each operation in the trace. For instance, the interval  $c_0 = t_0 - t_1$  on rank 1 indicates a computation period of  $c_0$  before initiating *MPI\_Send*. Through this process, Schedgen determines the computational workload and the operational dependencies within the application. Subsequently, by assigning specific LogGPS parameters to the sends and receives, we can effectively model an application’s behavior across different network configurations.

### B. Sensitivity Analysis

In this work, we explore the concept of network latency sensitivity. To frame our discussion, we will begin with an overview of sensitivity analysis (SA) formally. SA examines how a set of  $N$  input variables  $x = \{x_1, \dots, x_N\}$  influences the output  $y = \{y_1, \dots, y_D\}$  of  $y = g(x)$  where  $g: \mathbb{R}^N \rightarrow \mathbb{R}^D$  [42], [43]. When  $g$  is differentiable, derivative-based local SA can be performed by computing the partial derivative of  $y$  w.r.t. the  $i$ -th value in the input,  $x_i$ , denoted as  $\lambda_i = \left. \frac{\partial y}{\partial x_i} \right|_{x_0}$ , where  $\lambda_i$  is the sensitivity measure of  $x_i$ , and  $x_0 \in \mathbb{R}^n$  is the fixed point for evaluating the derivative [44], [45].

We start by exploring the impact of network latency,  $L$ , on the runtime of MPI programs. According to derivative-based



(a) Three possible paths through the DAG are highlighted with arrows of different patterns. (b) Values are assigned to parameters in a way that makes  $\partial T/\partial L = 1$ . The critical path is highlighted in orange. (c) If  $t_0$  is reduced to  $0.1 \mu s$ ,  $\partial T/\partial L$  will be dependent on the value of  $L$ . The relationship between  $T$ ,  $\partial T/\partial L$ , and  $L$  are plotted. The critical latency  $L_c$  is calculated to be  $0.385 \mu s$ .

Fig. 4: An example demonstrating that the network latency sensitivity of a program is determined by the number of messages along the critical path of the graph and is also dependent on the value of parameter  $L$  itself.

local SA, we can express the network latency sensitivity as

$$\lambda_L(\mathcal{G}, \theta) = \partial T(\mathcal{G}, \theta) / \partial L \quad (1)$$

where  $\mathcal{G}$  represents an execution graph,  $\theta$  denotes a specific configuration (i.e., a vector defining each parameter's value in LogGPS), and  $T(\mathcal{G}, \theta)$  computes the program runtime given  $\mathcal{G}$  and under the configuration  $\theta$ .  $\mathcal{G}$  and  $\theta$  are omitted if they are unambiguous. Intuitively,  $\lambda_L$  characterizes the variation in a program's runtime if  $L$  is increased by 1 unit.

In a parallel program, its runtime is determined by the critical path [32], [33], [46]. Therefore, to compute  $\lambda_L$ , it is crucial to derive an expression for the critical path as a function of  $L$ . To illustrate this, we refer to Fig. 4. Fig. 4a contains the execution graph for two ranks, where rank 0 executes MPI\_Send while rank 1 executes MPI\_Recv. Both ranks perform computation before and after the MPI operations, indicated by the green vertices connected to nodes  $S$  and  $R$  with costs  $c_0$  through  $c_3$ , respectively. There are in total three possible paths through the DAG, and without assigning values to the costs and the LogGPS parameters, the critical path can be any of them. Hence,  $T$  can be expressed formally as

$$T(\theta) = \max \left( \begin{array}{l} \underbrace{C_0 \rightarrow S \rightarrow C_1}_{t_0 + o + t_1}, \quad \underbrace{C_2 \rightarrow R \rightarrow C_3}_{t_2 + t_3 + o}, \\ \underbrace{C_0 \rightarrow S \rightarrow R \rightarrow C_3}_{t_0 + o + L + (s-1)G + t_3 + o} \end{array} \right) \quad (2)$$

where each term in the  $\max$  operator represents a distinct path. In Fig. 4b, we show that after substituting  $c_0 = c_1 = c_3 = 1 \mu s$ ,  $c_2 = 0.5 \mu s$ ,  $s = 4$ ,  $o = 0 s$ , and  $G = 5 \text{ ns}$  into Equation 2, it can be simplified to  $T = L + 2.015 \mu s$ , which leads to the result  $\lambda_L = \partial(L + 2.015) / \partial L = 1$ . This is designed to highlight that depending on vertices' costs and their precedence in the graph, some communication edges will stay on the critical path regardless of the value of  $L$ . In this scenario, due to a late sender, the edge  $S \rightarrow R$  remains on the critical path, making  $\lambda_L$  independent of the  $L$ 's value.

Nonetheless, if  $c_0$  is changed from  $1 \mu s$  to  $0.1 \mu s$  while all other variables stay the same, Equation 2 will evaluate to:  $T = \max(L + 1.115 \mu s, 1.5 \mu s)$ , in which case the result yielded by the  $\max$  operator, and consequently, the critical path, will be dependent on  $L$ . After plotting  $T = 1.5 \mu s$  and

$T = L + 1.115 \mu s$  as shown in Fig. 4c, we notice that when  $L \leq 0.385 \mu s$ ,  $T$  stays constant as the communication cost is overlapped. Once  $L$  is larger than this threshold,  $T$  starts to increase linearly. We refer to this point at which the critical path changes the critical latency, denoted as  $L_c$ . The same overlapping effect can also be observed within a rank when non-blocking operations are executed.

To generalize, the runtime  $T$  of an MPI program under the LogGPS model can be expressed as:

$$T(\mathcal{G}, \theta) = \max(a_0 L + C_0, a_1 L + C_1, \dots, a_n L + C_n) \quad (3)$$

where the  $i$ -th term in  $\max$  is the cost of a distinct path after simplification,  $a_i$  is the number of communication edges along the path and  $C_i$  is a constant representing all other costs. If two paths share the same  $a_i$ , the one with larger  $C_i$  will be kept. From Equation 3, we know that  $n$  is bounded by the length of the longest chain of messages in the graph.

We gained two insights from the formulation of  $T$ . Firstly, the number of messages along the critical path dictates  $\lambda_L$ . Secondly, the value of  $L$  has a second-order effect on the network latency sensitivity. As  $L$  increases, more communication edges that cannot be overlapped will lead to an increase in  $\lambda_L$ , and less latency tolerant a program will be. Therefore, it is crucial for network engineers to understand both the critical latencies and  $\lambda_L$  of MPI programs within an interval of interest. With this knowledge, they can make informed decisions regarding network configurations, improving program performance and resilience to varying levels of  $L$ .

1) *Generalization*: Extending our sensitivity analysis beyond network latency,  $L$ , we can apply similar principles to assess the impact of other parameters, such as  $G$  in LogGPS. For this purpose, consider the following expression:

$$T(\theta) = \max(s_0 G + C_0, s_1 G + C_1, \dots, s_m G + C_m) \quad (4)$$

Similar to Equation 2, each term in  $\max$  represents the cost of a path after simplification. In this case,  $s_i$  is approximately the number of bytes contained in messages along each path. Therefore, after evaluation, the bandwidth sensitivity measure  $\lambda_G(\mathcal{G}, \theta) = \frac{\partial T(\mathcal{G}, \theta)}{\partial G}$  can be seen as reflecting the total message size along the critical path. In essence, since  $T$  is determined by the critical path, as long as we can express the cost of each



path as a function of a parameter, we can derive the sensitivity measure for that parameter. Despite this generalization, our investigation will focus on exploring the impact of  $L$ .

### C. Linear Programming in LLAMP

Sensitivity measures of parameters can be computed through two conventional graph analysis approaches. The first approach involves traversing the graph twice: first to assign timestamps to all vertices for a specific  $\theta$ , and second to obtain the critical path and relevant metrics. Despite its linear time complexity (i.e.,  $O(|E|+|V|)$ ), it requires parameter sweeps to identify critical latencies. The second approach aims to find the general expression for  $T$  directly, such as Equations 3 and 4, enabling us to determine all critical latencies at once. Yet, since this requires iterating over all possible paths, it is generally intractable. An alternative would be using dynamic programming (DP) to store path-related information as a map for each vertex. For  $\lambda_L$ , the time complexity would be  $O(n|E| + |V|)$ , where  $n$  is the length of the longest message chain in the graph. While DP seems effective, when we tested it on a graph of LULESH with around 500K vertices, it took more than 4 hours to finish, which indicates that DP is also not scalable, even for graphs of moderate sizes.

1) *Linear Program Formulation*: While conventional graph analysis approaches for computing sensitivity measures exhibit inherent limitations, we present an alternative methodology that employs linear programming (LP). LP is a technique that maximizes or minimizes a linear function known as the *objective function* subject to a set of linear constraints [47]:

$$\begin{aligned} \text{Maximize: } & z = \mathbf{c}^T \mathbf{x} && z \in \mathbb{R}, \mathbf{c} \in \mathbb{R}^n, \mathbf{x} \in \mathbb{R}^n \\ \text{subject to: } & \mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq 0 && \mathbf{A} \in \mathbb{R}^{m \times n}, \mathbf{b} \in \mathbb{R}^m \end{aligned}$$

The canonical formulation of an LP, as illustrated above, involves  $\mathbf{A}$ ,  $\mathbf{b}$ , and  $\mathbf{c}$ , which all contain constants defined by the restrictions of the given problem. A solution  $\mathbf{x}$  is *feasible* if it satisfies all the constraints. The goal of solving an LP is to determine the values of *decision variables* in vector  $\mathbf{x}$  that lead to the optimal solution in the *feasible region*. A *basis* is a subset of  $\mathbf{A}\mathbf{x} \leq \mathbf{b}, \mathbf{x} \geq 0$  containing  $n$  linearly independent constraints. Geometrically, it uniquely determines a vertex on the polyhedron shaped by the feasible region. After solving an LP, a constraint is considered *tight* if the solution satisfies the constraint with equality, essentially defining a boundary of the polyhedron where the optimal solution is located.

To illustrate the relationship between execution graphs and LP, we will examine the running example presented in Fig. 4c. By substituting the values of the parameters into Equation 3 and re-organizing it, we gain the following formula:

$$T = \max(\underbrace{1.1}_{C_0 \rightarrow S \rightarrow C_1}, \max(\underbrace{L + 0.115}_{C_0 \rightarrow S \rightarrow R}, \underbrace{0.5}_{C_2 \rightarrow R})) + \underbrace{1}_{C_3} \quad (5)$$

In Equation 5, the costs of paths or vertices associated with the terms are labeled accordingly. To convert this expression to an LP problem, each *max* operator can be interpreted as two

constraints where each constraint corresponds to one term, and the latency  $L$  needs to be replaced with a decision variable. The linear program, after conversion, is shown as follows:

$$\min_{\ell, y_1, t} t \quad (6)$$

$$\text{s.t. (1) } y_1 \geq \ell + 0.115, \text{ (2) } y_1 \geq 0.5, \text{ (3) } t \geq 1.1, \text{ (4) } t \geq y_1 + 1$$

where  $y_1$  corresponds to the outcome of the inner *max* whereas  $t$  represents the result of the outer *max* and serves as the variable to be minimized in the objective function. Note that the LP is not formulated in its canonical form for easier understanding. To obtain  $T$  for a given  $L$ , one only needs to add the constraint  $\ell \geq L$ , and solve for the objective value. To distinguish between the decision variables denoting different LogGPS parameters and the constants that are assigned as their lower bounds, we use the lowercase letter in mathematical font to indicate the decision variables (e.g.,  $\ell, g, o$ ).




Fig. 5: Visualization of Equation 6. Dotted green lines define the linear constraints. Shaded areas mark the infeasible region. Blue lines highlight the borders of the feasible region.

Fig. 5 visualizes the LP in Equation 6. By adding  $l \geq 0.5$ , we aim to calculate the runtime when  $L = 0.5 \mu s$ . While the LP originally includes three decision variables, we have integrated the value of  $y_1$  into (4), reducing the problem to two dimensions for easier visualization. The visualization clearly pinpoints the result at the coordinate  $(0.5, 1.615)$ , demonstrating that when  $L = 0.5 \mu s$ , the runtime  $T = 1.615 \mu s$ .

To generalize, we consider the start and finish times of a vertex  $v$  as  $t_s(v)$  and  $t_e(v)$ . While traversing the graph in topological order,  $t_e(v) = t_s(v) + cost(v)$ , where  $cost(v)$  returns the execution time of  $v$ . If  $v$  has only one predecessor,  $u$ , then  $t_s(v) = t_e(u)$ . Conversely, if  $v$  has more than one predecessor, we introduce a new decision variable  $y_v$  and add the following set of constraints  $\{y_v \geq t_s(u) + cost(u, v) | \forall u \in V, (u, v) \in E\}$ . Since we only traverse the graph once, the time complexity of this method is  $O(|V| + |E|)$ . The space complexity is also  $O(|V| + |E|)$  as the number of variables and constraints in the LP are bounded by  $|V|$  and  $|E|$ , respectively. The complexity of solving the LP is discussed in Section II-D3.

### D. Performance Metrics

1) *Sensitivity Measures*: After converting the execution graph into an LP, we gain the ability not only to estimate an application's runtime but also to leverage the concept of *reduced cost* (RC) for assessing network latency sensitivity. The RC of a decision variable quantifies the amount at which

the objective value will change if the value of the variable varies by 1 unit. As illustrated in Fig. 5, the reduced cost for  $l$  equals 1, indicating a direct correlation where an increase in  $l$  by  $1 \mu s$  results in an increase in  $t$  by  $1 \mu s$ . Thus, the RC of  $l$  directly reflects  $\lambda_L$ . Since by solving an LP, we automatically calculate the RC for all decision variables, we can obtain sensitivity measures of parameters such as  $L$  and  $G$  much more efficiently, unlike graph analysis, which requires two separate traversals of the graph to deduce the same metric.

Constraints in LP problems can be viewed as an alternative representation of edges in graphs. Intuitively, *if a set of constraints are tight after optimization, their corresponding edges are on the critical path*. For instance, when we introduce  $l \geq 0.5$  to Equation 6 and solve for  $t$ , we observe that constraints (1) and (4) are tight, indicating that the path they represent,  $C_0 \rightarrow S \rightarrow R \rightarrow C_3$  in Fig.4a, is the critical path. As a basis closely correlates with constraints, if the optimal basis of an LP remains constant, the critical path of its corresponding execution graph will be unchanged. Given any  $L$  as the lower bound of  $l$ , we can compute its *range of feasibility*,  $L_{fl} \leq L \leq L_{fu}$ , where  $L_{fl}$  and  $L_{fu}$  are the feasibility lower bound and upper bound respectively. Within this range, both the critical path and  $\lambda_L$  would remain the same. Note that a linear solver produces the range of feasibility of all variables after optimization. The difference between our approach and graph analysis is that every time the LP is solved, it provides us with information about the region around the optimal basis. Thus, we can explore an interval more efficiently.

Derived from  $\lambda_L$ , which indicates the number of messages along the critical path as discussed in Section II-B, the product  $L \cdot \lambda_L$  reveals the total time consumed by network latency on the critical path. Consequently, we define  $\rho_L = (L \cdot \lambda_L)/T$ , termed as the *L ratio*.  $\rho_L$  highlights what fraction of the critical path's execution time is due to network latency. By calculating this ratio, we gain insight into the overall impact of network latency on the performance of the application.

2) *Network Latency Tolerance*: Beyond assessing network latency sensitivity, a significant benefit of LP lies in its ability to directly calculate the network latency tolerance of applications. Importantly, in Fig. 1, different tolerance values were not derived from iterating over predicted or actual application runtimes; instead, they were directly computed by solving LPs.

To achieve this, it requires two simple adjustments to the original LP. Firstly, we change the objective function from minimizing  $t$  to maximizing  $l$ . Following this, an additional constraint is introduced to set the maximum allowable execution time  $t$ . Solving the LP under these conditions yields the precise maximum tolerable  $L$ , beyond which the application's runtime exceeds a predefined threshold. We thus define the  $x\%L$  tolerance of an application as the highest  $L$  tolerable before experiencing more than  $x\%$  degradation in performance.

Fig. 6 showcases the modified LP applied to our running example. Our goal here is to determine the maximum network latency,  $L$ , that keeps the application's runtime below  $2 \mu s$ . Hence, the model incorporates a constraint  $t \leq 2 \mu s$  and




Fig. 6: The alternative optimization problem that tries to maximize  $l$  given an upper bound of the decision variable  $t$ . The shaded regions are infeasible.  $t \geq 1.1$  is not shown.

optimizes for  $l$ . As illustrated, the LP efficiently identifies the optimal solution as  $0.885 \mu s$ . This underscores LP's effectiveness as a tool for network architects and software engineers, providing a direct measure of network latency tolerance without the need for exhaustive parameter sweeps.

3) *Advantages of Linear Programming*: While transforming an execution graph into a linear program has linear time and space complexities, some doubts might be raised regarding the efficiency of this approach considering that solving a linear program is known to have exponential time complexity in the worst case with the simplex algorithm [48], [49]. In theory, one might expect traditional simulation or graph analysis methods to be more scalable. However, the reality is different.





Fig. 7: Runtime of the linear solver used by LLAMP vs. LogGOPSim across various applications. The number in the brackets lists the number of events in the execution graph.

The runtime comparison between LLAMP utilizing Gurobi [51], a state-of-the-art linear solver, and LogGOPSim is presented in Fig.7. Our choice of LogGOPSim for comparison is based on two key factors. Firstly, LogGOPSim stands out as one of the most efficient and scalable simulators currently available [28], [31]. Secondly, the graphs used to generate the LP models in our analysis are created by Schedgen, a component also utilized by LogGOPSim, ensuring a fair comparison. Despite LogGOPSim's notable speed and scalability, it is consistently outperformed by the linear solver, often by a factor of more than  $6\times$ , regardless of the problem size. This discrepancy arises mainly because the presolve phase of the linear solver efficiently eliminates all redundant constraints with advanced heuristics [52], [53]. This drastically reduces the solve time. Moreover, the barrier algorithm, also referred to as the interior point method [48], [54], employed by the solver allows the problems to be easily parallelized. Consequently, in practice, employing LP for performance forecasting is more efficient.



(a) Intended outcome of latency injection in which  $\Delta L$  latency is added to the base latency  $L_0$ .  
 (b) Adding delay to sends without any additional threads. (Underwood et al. [50])

(c) Adding delay on the receiver's side with a progress thread.

(d) Adding delay to receives with progress and delay threads. (Our work)

Fig. 8: (A) shows the expected result when  $\Delta L$  is added to the network latency in a scenario where the sender ( $R_0$ ) performs two consecutive eager `MPI_Sends` while the receiver ( $R_1$ ) posts two `MPI_Recv`s prior to the sends.  $t_{R_0}$  and  $t_{R_1}$  represent runtime of the ranks, respectively. The orange vertical lines indicate the time at which each send starts. The arrows denote the transmission of entire messages.  $L_0$  represents the network's base latency. For brevity, the bandwidth cost for message transmission,  $(s-1)G$ , is simplified to  $B$ . Subsequent panels depict various network latency injector implementations and their effects. Expressions in each panel that match those in (A) are marked in green, while those that differ are in red.

### III. VALIDATION

As LLAMP forecasts application performance, we can artificially introduce latency to network communication so as to compare its actual impact with LLAMP's predictions. In this section, we present two major contributions: a software-based latency injector for precise, scalable flow-level latency injection without specialized hardware or admin privileges, and a demonstration of LLAMP's predictive accuracy.

#### A. Network Latency Injector

Emulating network latency accurately is a complex task. For instance, take the situation shown in Fig. 8, where our goal is to inject an extra delay,  $\Delta L$ , into the network. We expect  $R_0$  to complete sending messages by timestamp  $2o$ , and  $R_1$  to receive them at  $3o + L_0 + B + \Delta L$ , assuming both start at timestamp 0. The simplest method to add latency is to delay the send operation by  $\Delta L$ . As seen in (B), this unintentionally delays both  $R_0$  and  $R_1$ . Note that even with an additional progress thread to handle the sending, one `MPI_Send` would still have to wait until  $o + \Delta L$  before returning, delaying the next send. Underwood et al.'s work [50] exploited this basic approach by hooking their latency injector into the `post_send` function in `libibverbs` library. However, as we have shown here, it would introduce unwanted delays to consecutive send operations.

The approach in (C) adds a progress thread to the receiver to process the delay, freeing the sender from the wait. Nonetheless, this method faces issues when  $\Delta L$  exceeds  $o$ , which is often the case in practice. The progress thread is still handling the first message's delay when the second arrives. This leads to the second message being delayed to  $o + L_0 + B + 2\Delta L$ , not the expected  $2o + L_0 + B + \Delta L$ . To concurrently process delays for multiple messages, the progress thread would need to track release times while polling the receive queue, which can greatly reduce the accuracy and resolution of  $\Delta L$ .

Our solution, depicted in (D), utilizes a *delay thread* to precisely manage  $\Delta L$  for each receive. When the progress

thread receives a message, it marks it with a timestamp,  $t_m$ , and passes it to the delay thread. This thread queues the message, releasing it to the application only when the current time matches  $t_m + \Delta L$ . This enables us to precisely emulate network latency, achieving the intended delay effect.

We implemented our solution in MPICH and UCX. MPICH was chosen since it already has asynchronous progress threads, unlike Open MPI. UCX was favored over libfabric for its simpler API, allowing easier customization. While one can place the delay queue on the sender side, this would require capturing send requests and manipulating data buffers. This would introduce overhead from memory copies and demand extensive modifications to the MPICH source code, potentially compromising its integrity. In conclusion, our solution stands out among other methods for its portability across all UCX transports. Moreover, it allows users to conduct cluster-wide experiments without requiring special privileges or changes to existing infrastructure.

#### B. Experimental Setup

All the experiments detailed in this section were conducted on a 188-node test-bed cluster maintained by the Swiss National Supercomputing Center (CSCS). The cluster features a fat-tree topology built on 18 Mellanox SX6036 switches. Each node is powered by a 20-core Intel Xeon CPU E5-2660 v2, equipped with 32 GB DDR3 RAM and a ConnectX-3 56 Gbit/s NIC, running CentOS 7.3. We used MPICH 4.1.2, UCX 1.16.0. The stack and all applications were compiled with GCC 11.4.0. To precisely measure the network parameters critical for the LogGPS model, we employed *Netgauge* 2.4.6 [55]. Aggregating measurements across the cluster yielded network parameters:  $L = 3.0 \mu\text{s}$ ,  $G = 0.018 \text{ ns}$ , and  $S = 256 \text{ KB}$ . Observing that  $o > g$  across all data sizes, we opted to omit  $g$  from the analysis. To enhance the accuracy of runtime predictions, we computed average packet sizes for each application from the traces and matched the  $o$  value using *Netgauge*'s outputs. The applications were

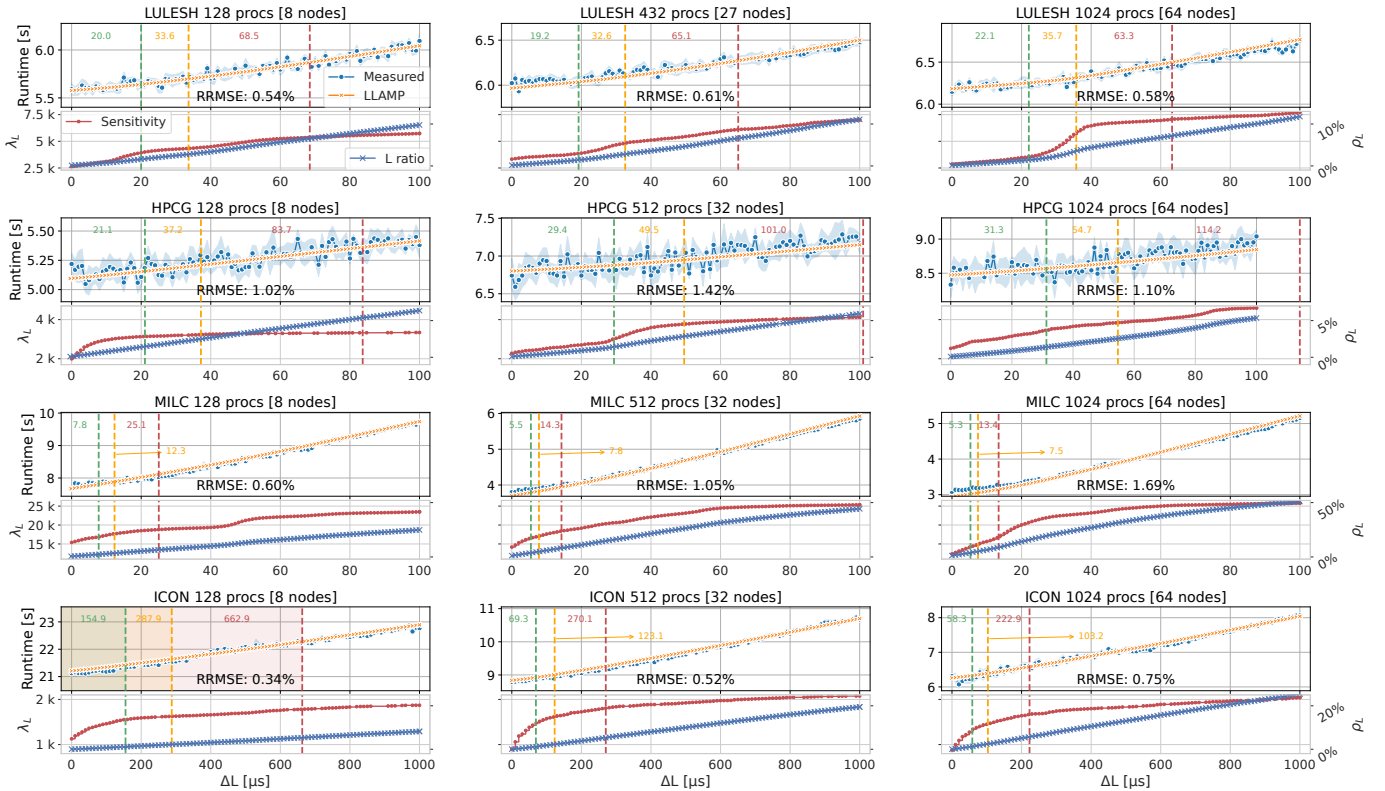


Fig. 9: Results for 4 selected applications, detailing the analysis across varying node and process counts. The upper section of each subplot compares actual runtimes against LLAMP predictions, with different colors and labels indicating 1%, 2%, and 5%  $L$  tolerance.  $RRMSE$  values highlight LLAMP’s prediction accuracy relative to measured runtimes. The lower section illustrates the variation in latency sensitivity,  $\lambda_L$ , and the latency ratio,  $\rho_L$ , against changes in  $\Delta L$ .

executed in a hybrid MPI+OpenMP mode in which every node runs one MPI rank and 16 OpenMP threads. For each variation in network latency,  $\Delta L$ , we conducted 10 runs per application, averaging the runtimes to produce the final results. The LLAMP toolchain was executed on a machine with an AMD EPYC 9654 96-Core 3.7 GHz and 375 GB of memory.

### C. Validation Results

Validation experiments were conducted on seven HPC applications across diverse domains, employing a variety of node configurations. Collecting runtime data required over 45 hours of cluster compute time. LLAMP demonstrated its efficiency by generating equivalent results, along with additional metrics, in 4 to 5 hours, which included time for tracing and generating LP models. Under the same condition, LogGOPSim is projected to need over 12 hours for simulation, an estimation that does not take into account LogGOPSim’s inability to evaluate  $\lambda_L$ . Due to space limits, we showcase the outcomes graphically from the four applications that best represent our findings in Fig. 9. Detailed results for all evaluated applications are listed in Table I.

The plots demonstrate LLAMP’s capability to accurately forecast the runtime of applications across various  $\Delta L$ , evidenced by the fact that relative root mean square errors

( $RRMSE$ ) [56] are consistently below 2%. A noticeable variation in the measured runtimes for HPCG suggests that it might be inherently more susceptible to system and network noise. Additionally, a slight bias is observed for MILC on setups involving 32 and 64 nodes, particularly within the 0 to 20  $\mu$ s  $\Delta L$  range. This bias may stem from MILC’s extensive use of persistent operations, in contrast to the `MPI_Send` and `MPI_Recv` operations utilized by Netgauge for measuring LogGP parameters. The differing overheads associated with persistent operations could contribute to this discrepancy.

Based on the definition of  $\lambda_L$  (i.e.,  $\partial T / \partial L$ ), the bottom plots allow us to assess how sensitive an application is to network latency for any given  $\Delta L$ . The visualizations clearly mark the intervals where  $\lambda_L$  increases most rapidly and where it reaches a plateau, indicating a convergence of the number of messages along the critical path toward the longest message chain in the execution graph. For example, the top right plot illustrates that LULESH, when executed on 64 nodes, has a  $\lambda_L$  that is stable between 0  $\mu$ s and 20  $\mu$ s. This shows that LULESH’s performance remains relatively unaffected by variations in  $L$  within this interval. Developers are encouraged to optimize their applications to improve communication-computation overlap, aiming for a flatter  $\lambda_L$  curve.

LULESH and HPCG were evaluated under weak scaling,



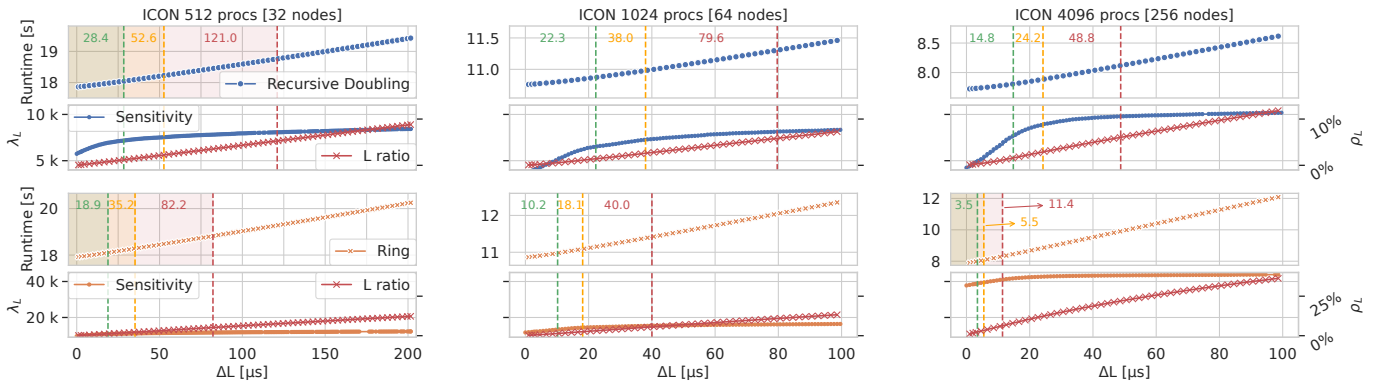


Fig. 10: Comparative analysis of ICON’s network latency sensitivity and tolerance when employing recursive doubling and the ring algorithm for `MPI_Allreduce`. The top plots illustrate ICON’s performance prediction with recursive doubling for allreduce. The corresponding bottom series showcases performance outcomes utilizing the ring algorithm for the same operation.

showing that their network latency tolerance remained relatively stable as the number of nodes increased. Interestingly, HPCG’s latency tolerance improved, likely due to the optimized overlap between communication and computation. On the other hand, MILC (`su3_rmd`) and ICON were subjected to strong scaling, revealing a substantial decrease in latency tolerance as the applications scaled up. This decline is attributed to the diminishing computational workload per node in strong scaling scenarios, reducing the amount of computation that may be overlapped. MILC, in particular, showed a significantly lower tolerance, highlighting its great reliance on communication.

#### IV. ICON CASE STUDY

Having validated LLAMP, we now turn our attention to applying it to a real-world application. For this purpose, we selected the Icosahedral Nonhydrostatic Weather and Climate Model (ICON) as our case study. ICON is a global atmospheric simulation framework designed for weather forecasting and climate studies, leveraging a novel icosahedral grid for higher precision and computational efficiency. Central to ICON is its dynamical core that solves the nonhydrostatic equations of motion. Our choice of ICON is motivated by its widespread adoption among prominent European weather services, including MeteoSwiss [57] and the German Weather Service (DWD) [58]. Given the importance of climate studies, ICON stood out as a perfect candidate to showcase the versatile functionalities of our toolchain.

The case study was conducted on the Piz Daint supercomputer, operated by CSCS. For this study, we deployed ICON version 2.6.7, compiled with Cray MPICH 12.0.3. The LogGPS parameters measured for the cluster were  $L = 1.4\mu s$ ,  $G = 0.013ns$ , and  $S = 256KB$ . ICON was executed in hybrid mode, utilizing 16 OpenMP threads per node. The experiments were conducted across three different scales: 32, 64, and 256 nodes. The values of  $o$  for these configurations were measured as  $8.5\mu s$ ,  $7.4\mu s$ , and  $6.03\mu s$ , respectively.

Application	# Proc / # Nodes	Time [min]	Size [MB]	# Trace Events	RMSE [s]	RRMSE [%]
LULESH	128/8	3	28	895 K	0.031	0.54
	432/27	12	137	4.55 M	0.038	0.61
	1024/64	40	353	13.5 M	0.037	0.58
HPCG	128/8	2.5	21	738 K	0.053	1.02
	512/32	15	133	4.83 M	0.099	1.42
	1024/64	35	333	12.2 M	0.095	1.10
MILC [su3_rmd]	128/8	10	69	2.79 M	0.052	0.60
	512/32	45	280	12.1 M	0.050	1.05
	1024/64	120	539	23.6 M	0.068	1.69
ICON	128/8	1.5	9.7	362 K	0.074	0.34
	512/32	4.5	51	1.89 M	0.051	0.52
	1024/64	12	102	3.82 M	0.053	0.75
LAMMPS	128/8	0.5	3.9	155 K	0.083	1.40
	512/32	2	16	596 K	0.102	1.46
	1024/64	3	32	1.21 M	0.128	1.65
OpenMX	128/8	0.6	4.6	271 K	0.204	0.88
	512/32	4.3	64	1.78 M	0.097	0.55
CloverLeaf	128/8	4.1	6.1	162 K	0.037	0.70

TABLE I: Validation results across all chosen applications. The third column shows the total time overhead required to produce the  $\lambda_L$  results as  $\Delta L$  varies from  $0\mu s$  to  $100\mu s$ . The fourth column displays the trace size produced by liballprof, while the fifth column lists the number of trace events in the execution graphs generated by Schedgen. The RMSE and RRMSE columns indicate the Root Mean Square Error and Relative Root Mean Square Error, respectively, between the predicted and measured runtimes.

1) *Impact of Collective Algorithms*: Firstly, we demonstrate LLAMP’s ability to analyze how different implementations of collective algorithms impact ICON’s performance. Given ICON’s reliance on allreduce for data exchange in its dynamical core, we explored the effect of changing its allreduce from recursive doubling to the ring algorithm. As discussed in Section II-A, this can be achieved by altering the scheduling for allreduce in Schedgen. The results are presented in Fig. 10,

where ICON was traced once per node configuration.

Our observations reveal that ICON’s performance becomes increasingly sensitive to  $L$  when employing the ring allreduce, a consequence of dependent sends and receives in this algorithm. This effect intensifies with the scaling of the application. Notably, at 256 nodes, ICON’s 5%  $L$  tolerance using ring allreduce stands at  $11.4 \mu s$ , in contrast to its recursive doubling counterpart, which exhibits a  $4\times$  latency tolerance. Furthermore, despite similar trends in their  $\lambda_L$  curves across varying scales, the magnitude of  $\lambda_L$  for the ring allreduce significantly exceeds that of recursive doubling, indicating a significant increase in latency sensitivity. At 256 nodes, ICON’s  $\rho_L$  surpasses 25% for a  $\Delta L = 100 \mu s$ , doubling the  $\rho_L$  observed with recursive doubling. This suggests that ring allreduce should be used cautiously, especially when bandwidth and network congestion are less critical.

Through LLAMP, we highlight how software engineers can assess the impacts of collective algorithms. This approach enables a more informed decision-making process in optimizing collective communications for HPC applications.

2) *Impact of Network Topology*: LLAMP’s utility extends beyond the evaluation of collective algorithms, it can also analyze various network topologies. This capability is illustrated in an example in Fig. 11. The core concept revolves around substituting the latency of all wires with a decision variable  $l_{wire}$ . Given a constant switch latency, denoted as  $d_{latency}$ , the cost to transmit a message between two nodes can be formulated as  $(h+1) \cdot l_{wire} + h \cdot d_{latency}$ , where  $h$  represents the number of hops. This allows us to model the communication cost between any two nodes as a function of  $l_{wire}$  and the number of hops as defined by the network topology. By setting  $l_{wire} \geq L$ , we effectively focus on wire latency instead of the end-to-end latency between nodes. Consequently, by calculating  $\partial T / \partial L$ , we can determine how sensitive an application is to changes in wire latency. This adjustment shifts the focus of analysis toward understanding the influence of the topology and the wire latency on application performance.




Fig. 11: An example demonstrating that by replacing the cost of communication edges of the execution graph with appropriate decision variables, we can investigate the impact of the wire latency as well as the network topology on the overall performance of an application.

In this case study, we leverage LLAMP to assess the impact of two predominant network topologies, namely Fat Tree [59] and Dragonfly [60], on ICON’s performance. Based on the

work of Zambre et al. [61], we set initial  $l_{wire}$  to 274 ns and  $d_{switch}$  to 108 ns. For the Fat Tree topology, we chose a three-tier design where each switch has a radix of 16 ( $k = 16$ ). The selected dragonfly topology configuration is characterized by  $g = 8$ ,  $a = 4$ , and  $p = 8$ . We disregard  $h$  in our calculations, assuming that routing is always minimal. Moreover, in our analysis, we assume nodes are densely packed within both topologies. For example, nodes 0 to 7 are clustered within the same pod in the Fat Tree configuration and under a single switch for the Dragonfly topology.

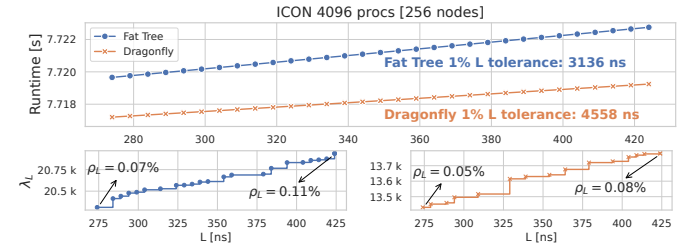


Fig. 12: Comparison of the impact of the Fat Tree and Dragonfly topologies on the performance of ICON.

We evaluate the performance impact of these two topologies by using the execution graph of the 256-node setup (with recursive doubling allreduce) and present the result in Fig. 12. Considering the potential for a more complex FEC to increase the latency per wire by over 100 ns as mentioned in the introduction, we selected an interval ranging from the base latency of 274 ns to 424 ns. Our results show that ICON under Dragonfly exhibits a marginally higher network latency tolerance compared to the Fat Tree. This advantage is attributed to Dragonfly’s lower average number of hops between nodes under our configurations. The key takeaway is that despite Dragonfly’s slightly superior network latency tolerance, the performance of ICON under both topologies remains largely unaffected by the anticipated latency increases from more complex FEC. This is demonstrated by the requirement for the *per-link latency* to exceed 3000 ns before ICON’s performance degrades by 1%.

This study, while hypothetical, showcases LLAMP’s adaptability in assessing the impact of network topologies. It highlights its utility for architects aiming to refine their systems for optimal performance with specific HPC applications.

## V. RELATED WORK

a) *Trace Replay*: Numerous tools exist for forecasting MPI applications’ performance based on collected traces. PHANTOM [62], [63] and LogGOPSim [31] utilize communication models from the LogP family, offering fast simulation times. PSINS [64] presents a choice among three communication models of varying complexity, with the PMAc model [65] standing out for its accuracy. Hermanns et al. [66] aim at reenacting communication operations from traces to identify load imbalances. SMPI [67] and the work of Desprez et al. [68], rely on SimGrid [69], a discrete event simulator, for more precise performance predictions during online simulations.

Kenny et al. [70] incorporate SST [71] and PISCES, a packet-level model, for architectural simulations, employing Bayesian inference for simulator validation. Eyerman et al. [72] propose integrating profiling, node simulation, and high-level network simulation with SST. Despite their capabilities, these simulators often rely on complex discrete event and packet-level simulators that are usually time-consuming to execute. One significant limitation they share, when compared to LLAMP, is the necessity for extensive parameter sweeps to ascertain each application’s network latency tolerance.

b) *Critical Path Analysis*: Graph-based analysis of MPI applications is a well-established approach. Schulz [32] and Böhme et al. [33] have contributed to identifying wait-states and load imbalances by employing critical path analysis and trace replay. Chen et al. [46] further refined this approach by introducing the concept of critical-path candidates, identifying potential critical paths through profiling-based instruction and communication counts. Schmitt et al. [73] expanded these mostly rely on traditional graph analysis, LLAMP introduces a completely new and unique perspective on viewing graphs and extracting information from the critical path.

c) *Network Latency Sensitivity and Tolerance*: Numerous studies have investigated how network latency affects application performance. Efforts by Underwood et al. [50], Gao et al. [74], Patke et al. [25], Link Gradients [75], and Richar Paul Martin’s thesis [76] have predominantly involved manually injecting latency to networks and conducting extensive trials with varied latency levels to identify patterns. Kerbyson et al. [23] utilized performance models of three specific applications for similar analyses. While insightful, it demands considerable effort and deep application knowledge, limiting their generality. Our work stands out by mathematically defining network latency sensitivity/tolerance and offering a scalable, analytical method to compute these metrics across a broad spectrum of MPI applications.

## VI. EXTENSIONS AND DISCUSSION

Beyond the capabilities we have shown, LLAMP holds extensive potential for a broad spectrum of analyses, offering users a flexible framework to explore a variety of performance metrics. LLAMP’s versatility extends to examining the sensitivity and tolerance of additional critical LogGPS model parameters, such as bandwidth  $G$ . We demonstrate LLAMP’s ability to perform bandwidth sensitivity analysis in Fig. 13 by setting  $L$  at  $3\mu s$  and varying the lower bound of the decision variable that represents  $G$  from  $0.02ns$  (50 GB/s) to  $0.1ns$  (10 GB/s). The result shows that  $\lambda_G$  exhibits a notable increase from  $14.52 \times 10^6$  to  $14.61 \times 10^6$  when  $G$  reached approximately  $0.07ns$  (13.8 GB/s). Beyond this point,  $G$  remained stable. This suggests that reducing network bandwidth below 13.8 GB/s causes the data volume on the critical path to rise from 14.52 MB to 14.61 MB. LLAMP showed that this change had a minimal impact on the overall runtime of LULESH.

Moreover, LLAMP’s support for *heterogeneous network models*, such as HLogGP, unlocks the potential to assess




Fig. 13: Bandwidth sensitivity analysis of LULESH.

pairwise network latency sensitivities across MPI ranks. Leveraging this, we devised a new process placement algorithm to optimize the mapping of MPI ranks onto physical processors. These functionalities exemplify LLAMP’s versatility and how it empowers users to tailor the toolchain to their specific needs.

There are aspects of LLAMP that can still be improved. Firstly, the trace-based nature of our analysis means that LLAMP relies on a static snapshot of application behavior to produce the performance metrics, which may not capture the variability of live execution. As exemplified by the works of Nikitenko et al. [77] and Hoeﬂer et al. [35], an application’s performance during tracing might be influenced by system noise and network congestion. To overcome this, we can incorporate a more detailed model, such as LogGOPSC [78], to account for network contention or statistics to account for the variability of network behavior and system noise, allowing LLAMP to predict a range of outcomes and their probabilities.




Fig. 14: Runtime of NAMD, a molecular dynamics simulation built on charm++, showing measured vs. predicted runtimes. The legend indicates  $\Delta L$  at which the traces were recorded.

Currently, LLAMP’s is not fully equipped to handle the intricacies of dynamically scheduled operations seen in frameworks like *charm++* [79], which dynamically adjust based on runtime conditions. However, LLAMP still manages to reveal the inherent latency tolerance within these frameworks. As illustrated in Fig. 14, LLAMP demonstrates how *charm++* proactively adjusts its communication schedule to enhance resilience to increased network latency. In the future, we plan to explore a general strategy that can analyze and quantify network latency sensitivity and tolerance across a variety of parallel programming models. This includes not only those that use dynamic scheduling, but also models that diverge from the MPI standard, such as Legion [80].

## VII. CONCLUSION

In this work, we introduce LLAMP, an innovative toolchain that integrates the LogGPS model with linear programming

to analyze network latency sensitivity and tolerance in HPC applications. By uncovering the intricate connection between the critical paths in a program’s execution graphs and LP, we provide a new method for efficiently gathering performance metrics. LLAMP’s flexibility allows users to utilize metrics such as  $\lambda_L$  and  $\rho_L$  in ways tailored to their specific needs.

Moreover, the development of a specialized latency injector that accurately emulates flow-level latency enables us to validate LLAMP’s predictions across a variety of HPC applications, demonstrating an error margin of less than 2%. To illustrate LLAMP’s practicality, we conducted a case study with ICON, a leading climate modeling tool in Europe, demonstrating its utility in addressing the needs of both software engineers and network architects.

#### ACKNOWLEDGEMENTS

The authors would like to thank Andrei Ivanov, Lukas Gianinazzi, and Alexandru Calotoiu for their helpful discussions. We also wish to thank the Swiss National Supercomputing Center (CSCS) for their support of this work. This work was financially supported by the 2023 Global Research Outreach Program of the Samsung Advanced Institute of Technology (SAIT) and Samsung R&D Center America, Silicon Valley (SRA-SV) under the supervision of the System Architecture Laboratory (SAL). This work received support from the PASC project DaCeMI: Harnessing future hardware using Data-Centric ML Integration. This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955776 and a donation from Intel. ChatGPT was used for light editing of the authors’ original text.

#### REFERENCES

[1] K. Wagner, “Meta is building new \$800 million ai-focused data center in indiana,” *Bloomberg*.

[2] CSCS, “Cscs, hewlett packard enterprise and nvidia announce world’s most powerful ai-capable supercomputer,” *Swiss National Supercomputing Center*.

[3] G. Guidi, M. Ellis, A. Buluç, K. Yelick, and D. Culler, “10 years later: Cloud computing is closing the performance gap,” in *Companion of the ACM/SPEC International Conference on Performance Engineering, ICPE ’21*, (New York, NY, USA), p. 41–48, Association for Computing Machinery, 2021.

[4] A. G. Carlyle, S. L. Harrell, and P. M. Smith, “Cost-effective hpc: The community or the cloud?,” in *2010 IEEE Second International Conference on Cloud Computing Technology and Science*, pp. 169–176, 2010.

[5] 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,” *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 6, dec 2022.

[6] A. Gupta and D. Milojicic, “Evaluation of hpc applications on cloud,” in *2011 Sixth Open Cirrus Summit*, pp. 22–26, 2011.

[7] A. Marathe, R. Harris, D. K. Lowenthal, B. R. de Supinski, B. Rountree, M. Schulz, and X. Yuan, “A comparative study of high-performance computing on the cloud,” in *Proceedings of the 22nd International Symposium on High-Performance Parallel and Distributed Computing, HPDC ’13*, (New York, NY, USA), p. 239–250, Association for Computing Machinery, 2013.

[8] R. R. Expósito, G. L. Taboada, S. Ramos, J. Touriño, and R. Doallo, “Performance analysis of hpc applications in the cloud,” *Future Generation Computer Systems*, vol. 29, no. 1, pp. 218–229, 2013. Including Special section: AIRCC-NetCoM 2009 and Special section: Clouds and Service-Oriented Architectures.

[9] D. Tomić, Z. Car, and D. Ogrizović, “Running hpc applications on many million cores cloud,” in *2017 40th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*, pp. 209–214, 2017.

[10] M. Naumov, J. Kim, D. Mudigere, S. Sridharan, X. Wang, W. Zhao, S. Yilmaz, C. Kim, H. Yuen, M. Ozdal, K. Nair, I. Gao, B.-Y. Su, J. Yang, and M. Smelyanskiy, “Deep learning training in facebook data centers: Design of scale-up and scale-out systems,” 2020.

[11] S. Athlur, N. Saran, M. Sivathanu, R. Ramjee, and N. Kwatra, “Varuna: scalable, low-cost training of massive deep learning models,” in *Proceedings of the Seventeenth European Conference on Computer Systems, EuroSys ’22*, (New York, NY, USA), p. 472–487, Association for Computing Machinery, 2022.

[12] K. Z. Ibrahim, T. Nguyen, H. A. Nam, W. Bhimji, S. Farrell, L. Oliker, M. Rowan, N. J. Wright, and S. Williams, “Architectural requirements for deep learning workloads in hpc environments,” in *2021 International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, pp. 7–17, 2021.

[13] S. Barrachina, A. Castelló, M. Catalán, M. F. Dolz, and J. I. Mestre, “Using machine learning to model the training scalability of convolutional neural networks on clusters of gpus,” *Computing*, vol. 105, no. 5, pp. 915–934, 2023.

[14] 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, pp. 67–77, 2023.

[15] J. D’Ambrosia, “Ieee p802.3df™ defines architecture holistically to achieve 800 gb/s and 1.6 tb/s ethernet,” *IEEE Standards Association*, 2022.

[16] C. Liu, “100+ Gb/s Ethernet Forward Error Correction (FEC) Analysis,” tech. rep., 2019.

[17] S. Bates, “Forward error correction for 400g - ieee 802,” 2013.

[18] T. Hoefler, T. Bonato, D. De Sensi, S. Di Girolamo, S. Li, M. Hedges, J. Belk, D. Goel, M. Castro, and S. Scott, “Hammingmesh: A network topology for large-scale deep learning,” in *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–18, 2022.

[19] R. Mayer and H.-A. Jacobsen, “Scalable deep learning on distributed infrastructures: Challenges, techniques, and tools,” *ACM Comput. Surv.*, vol. 53, feb 2020.

[20] C. Bernard, M. C. Ogilvie, T. A. Degrand, C. E. Detar, S. A. Gottlieb, A. Krasnitz, R. Sugar, and D. Toussaint, “Studying quarks and gluons on mimd parallel computers,” *Int. J. High Perform. Comput. Appl.*, vol. 5, p. 61–70, dec 1991.

[21] I. Karlin, J. Keasler, and R. Neely, “Lulesh 2.0 updates and changes,” Tech. Rep. LLNL-TR-641973, August 2013.

[22] G. Zängl, D. Reinert, P. Ripodas, and M. Baldauf, “The icon (icosahedral non-hydrostatic) modelling framework of dwd and mpi-m: Description of the non-hydrostatic dynamical core,” *Quarterly Journal of the Royal Meteorological Society*, vol. 141, no. 687, pp. 563–579, 2015.

[23] D. Kerbyson, “A look at application performance sensitivity to the bandwidth and latency of infiniband networks,” in *Proceedings 20th IEEE International Parallel & Distributed Processing Symposium*, pp. 7 pp.–, 2006.

[24] G. Bauer, S. Gottlieb, and T. Hoefler, “Performance modeling and comparative analysis of the milc lattice qcd application su3\_rmd,” in *2012 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (ccgrid 2012)*, pp. 652–659, 2012.

[25] 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)*, pp. 1221–1227, 2022.

[26] E. Rosenthal and E. Leon, “Characterizing application sensitivity to network performance,” 11 2014.

[27] A. Bhatlele, N. Jain, M. Mubarak, and T. Gamblin, “Analyzing cost-performance tradeoffs of hpc network designs under different constraints using simulations,” in *Proceedings of the 2019 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation, SIGSIM-PADS ’19*, (New York, NY, USA), p. 1–12, Association for Computing Machinery, 2019.

[28] M. Besta, M. Schneider, S. D. Girolamo, A. Singla, and T. Hoefler, “Towards million-server network simulations on just a laptop,” *CoRR*, vol. abs/2105.12663, 2021.



- [29] G. Xu, H. Ibeid, X. Jiang, V. Svilan, and Z. Bian, "Simulation-based performance prediction of hpc applications: A case study of hpl," in *2020 IEEE/ACM International Workshop on HPC User Support Tools (HUST) and Workshop on Programming and Performance Visualization Tools (ProTools)*, pp. 81–88, 2020.
- [30] M. Mubarak, C. D. Carothers, R. B. Ross, and P. Carns, "Enabling parallel simulation of large-scale hpc network systems," *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 1, pp. 87–100, 2017.
- [31] 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, HPDC '10*, (New York, NY, USA), p. 597–604, Association for Computing Machinery, 2010.
- [32] M. Schulz, "Extracting critical path graphs from mpi applications," in *2005 IEEE International Conference on Cluster Computing*, pp. 1–10, 2005.
- [33] D. Böhme, F. Wolf, B. R. de Supinski, M. Schulz, and M. Geimer, "Scalable critical-path based performance analysis," in *2012 IEEE 26th International Parallel and Distributed Processing Symposium*, pp. 1330–1340, 2012.
- [34] R. Pereira, A. Roussel, P. Carribault, and T. Gautier, "Investigating dependency graph discovery impact on task-based mpi+openmp applications performances," in *Proceedings of the 52nd International Conference on Parallel Processing, ICPP '23*, (New York, NY, USA), p. 163–172, Association for Computing Machinery, 2023.
- [35] T. Hoefler, T. Schneider, and A. Lumsdaine, "Characterizing the influence of system noise on large-scale applications by simulation," in *SC '10: Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–11, 2010.
- [36] S. Levy, B. Topp, K. B. Ferreira, D. Arnold, T. Hoefler, and P. Widener, "Using simulation to evaluate the performance of resilience strategies at scale," in *High Performance Computing Systems. Performance Modeling, Benchmarking and Simulation* (S. A. Jarvis, S. A. Wright, and S. D. Hammond, eds.), (Cham), pp. 91–114, Springer International Publishing, 2014.
- [37] R. Riesen, C. Vaughan, and T. Hoefler, "What if mpi collective operations were instantaneous?," *Cray User Group (CUG)*, 2006.
- [38] T. Hoefler, C. Siebert, and A. Lumsdaine, "Group operation assembly language - a flexible way to express collective communication," in *2009 International Conference on Parallel Processing*, pp. 574–581, 2009.
- [39] D. Culler, R. Karp, D. Patterson, A. Sahay, K. E. Schauer, E. Santos, R. Subramonian, and T. von Eicken, "Logp: towards a realistic model of parallel computation," *SIGPLAN Not.*, vol. 28, p. 1–12, jul 1993.
- [40] J. A. Rico-Gallego, J. C. Díaz-Martín, R. R. Manumachu, and A. L. Lastovetsky, "A survey of communication performance models for high-performance computing," *ACM Comput. Surv.*, vol. 51, jan 2019.
- [41] A. Alexandrov, M. F. Ionescu, K. E. Schauer, and C. Scheiman, "Loggp: incorporating long messages into the logp model—one step closer towards a realistic model for parallel computation," in *Proceedings of the Seventh Annual ACM Symposium on Parallel Algorithms and Architectures, SPAA '95*, (New York, NY, USA), p. 95–105, Association for Computing Machinery, 1995.
- [42] E. Borgonovo and E. Plischke, "Sensitivity analysis: A review of recent advances," *European Journal of Operational Research*, vol. 248, no. 3, pp. 869–887, 2016.
- [43] A. Saltelli, S. Tarantola, F. Campolongo, and M. Ratto, *Sensitivity Analysis in Practice: A Guide to Assessing Scientific Models*. Wiley, 2004.
- [44] E. Borgonovo, "Sensitivity analysis of model output with input constraints: A generalized rationale for local methods," *Risk Analysis*, vol. 28, no. 3, pp. 667–680, 2008.
- [45] F. Pianosi, K. Beven, J. Freer, J. W. Hall, J. Rougier, D. B. Stephenson, and T. Wagener, "Sensitivity analysis of environmental models: A systematic review with practical workflow," *Environmental Modelling & Software*, vol. 79, pp. 214–232, 2016.
- [46] J. Chen and R. M. Clapp, "Critical-path candidates: scalable performance modeling for mpi workloads," in *2015 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, pp. 1–10, 2015.
- [47] The Editors of Encyclopaedia Britannica, "linear programming." <https://www.britannica.com/science/linear-programming-mathematics>, Feb 2024. Accessed 25 March 2024.
- [48] *Linear Programming Algorithms*, pp. 71–97. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008.
- [49] V. Klee and G. J. Minty, "How good is the simplex algorithm," 1970.
- [50] 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, ICPE '18*, (New York, NY, USA), p. 68–79, Association for Computing Machinery, 2018.
- [51] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2023.
- [52] Wunderling, Roland, "Advanced Gurobi Algorithms," 2022.
- [53] E. D. Andersen and K. D. Andersen, "Presolving in linear programming," *Mathematical Programming*, vol. 71, no. 2, pp. 221–245, 1995.
- [54] S. Bradley, A. Hax, and T. Magnanti, *Applied Mathematical Programming*. Addison-Wesley Publishing Company, 1977.
- [55] T. Hoefler, T. Mehlan, A. Lumsdaine, and W. Rehm, "Netgauge: A Network Performance Measurement Framework," in *Proceedings of High Performance Computing and Communications, HPCC '07*, vol. 4782, pp. 659–671, Springer, Sep. 2007.
- [56] M. Despotovic, V. Nedic, D. Despotovic, and S. Cvetanovic, "Evaluation of empirical models for predicting monthly mean horizontal diffuse solar radiation," *Renewable and Sustainable Energy Reviews*, vol. 56, pp. 246–260, 2016.
- [57] MeteoSwiss, "Icon-22," 2019.
- [58] DWD, "Icon (icosahedral nonhydrostatic) model."
- [59] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," in *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication, SIGCOMM '08*, (New York, NY, USA), p. 63–74, Association for Computing Machinery, 2008.
- [60] J. Kim, W. J. Dally, S. Scott, and D. Abts, "Technology-driven, highly-scalable dragonfly topology," in *2008 International Symposium on Computer Architecture*, pp. 77–88, 2008.
- [61] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," in *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication, SIGCOMM '08*, (New York, NY, USA), p. 63–74, Association for Computing Machinery, 2008.
- [62] J. Zhai, W. Chen, and W. Zheng, "Phantom: predicting performance of parallel applications on large-scale parallel machines using a single node," in *Proceedings of the 15th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP '10*, (New York, NY, USA), p. 305–314, Association for Computing Machinery, 2010.
- [63] W. Chen, J. Zhai, J. Zhang, and W. Zheng, "Loggpo: An accurate communication model for performance prediction of mpi programs," *Science in China Series F: Information Sciences*, vol. 52, no. 10, pp. 1785–1791, 2009.
- [64] M. M. Tikir, M. A. Laurenzano, L. Carrington, and A. Snavely, "Psin: An open source event tracer and execution simulator for mpi applications," in *Euro-Par 2009 Parallel Processing* (H. Sips, D. Epema, and H.-X. Lin, eds.), (Berlin, Heidelberg), pp. 135–148, Springer Berlin Heidelberg, 2009.
- [65] A. Snavely, L. Carrington, N. Wolter, J. Labarta, R. Badia, and A. Purkayastha, "A framework for performance modeling and prediction," in *SC '02: Proceedings of the 2002 ACM/IEEE Conference on Supercomputing*, pp. 21–21, 2002.
- [66] M.-A. Hermanns, M. Geimer, F. Wolf, and B. J. Wylie, "Verifying causality between distant performance phenomena in large-scale mpi applications," in *2009 17th Euromicro International Conference on Parallel, Distributed and Network-based Processing*, pp. 78–84, 2009.
- [67] A. Degomme, A. Legrand, G. S. Markomanolis, M. Quinson, M. Stillwell, and F. Suter, "Simulating mpi applications: The smpi approach," *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 8, pp. 2387–2400, 2017.
- [68] F. Desprez, G. S. Markomanolis, M. Quinson, and F. Suter, "Assessing the performance of mpi applications through time-independent trace replay," in *2011 40th International Conference on Parallel Processing Workshops*, pp. 467–476, 2011.
- [69] H. Casanova, A. Giersch, A. Legrand, M. Quinson, and F. Suter, "Versatile, scalable, and accurate simulation of distributed applications and platforms," *Journal of Parallel and Distributed Computing*, vol. 74, pp. 2899–2917, June 2014.
- [70] J. P. Kenny, K. Sargsyan, S. Knight, G. Michelogiannakis, and J. J. Wilke, "The pitfalls of provisioning exascale networks: A trace replay analysis for understanding communication performance," in *High Per-*

- formance Computing* (R. Yokota, M. Weiland, D. Keyes, and C. Trinitis, eds.), (Cham), pp. 269–288, Springer International Publishing, 2018.
- [71] A. F. Rodrigues, K. S. Hemmert, B. W. Barrett, C. Kersey, R. Oldfield, M. Weston, R. Risen, J. Cook, P. Rosenfeld, E. Cooper-Balis, and B. Jacob, “The structural simulation toolkit,” *SIGMETRICS Perform. Eval. Rev.*, vol. 38, p. 37–42, mar 2011.
  - [72] S. Eyerman, W. Heirman, K. D. Bois, and I. Hur, “Accurate and scalable many-node simulation,” 2024.
  - [73] F. Schmitt, R. Dietrich, and G. Juckeland, “Scalable critical path analysis for hybrid mpi-cuda applications,” in *2014 IEEE International Parallel & Distributed Processing Symposium Workshops*, pp. 908–915, 2014.
  - [74] P. X. Gao, A. Narayan, S. Karandikar, J. Carreira, S. Han, R. Agarwal, S. Ratnasamy, and S. Shenker, “Network requirements for resource disaggregation,” in *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, OSDI’16, (USA), p. 249–264, USENIX Association, 2016.
  - [75] S. Chen, K. R. Joshi, M. A. Hiltunen, W. H. Sanders, and R. D. Schlichting, “Link gradients: Predicting the impact of network latency on multitier applications,” in *IEEE INFOCOM 2009*, pp. 2258–2266, 2009.
  - [76] R. P. Martin, *A Systematic Characterization of Application Sensitivity to Network Performance*. PhD thesis, EECS Department, University of California, Berkeley, 1999.
  - [77] D. A. Nikitenko, F. Wolf, B. Mohr, T. Hoefler, K. S. Stefanov, V. V. Voevodin, A. S. Antonov, and A. Calotoiu, “Influence of noisy environments on behavior of hpc applications,” *Lobachevskii Journal of Mathematics*, vol. 42, no. 7, pp. 1560–1570, 2021.
  - [78] Y. Baicheng, Z. Yi, X. Limin, H. Jiantong, and W. Zhaokai, “Loggops: A parallel computation model extending network contention into loggops,” in *2019 IEEE International Conference on Cluster Computing (CLUSTER)*, pp. 1–2, 2019.
  - [79] L. V. Kale and S. Krishnan, “Charm++: a portable concurrent object oriented system based on c++,” *SIGPLAN Not.*, vol. 28, p. 91–108, oct 1993.
  - [80] M. Bauer, S. Treichler, E. Slaughter, and A. Aiken, “Legion: Expressing locality and independence with logical regions,” in *SC ’12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, pp. 1–11, 2012.