VEDRFOLNIR: Diagnosing RDMA Network Performance Anomalies in Collective Communications

Yuxuan Chen, *Graduate Student Member, IEEE*, Menghao Zhang, *Member, IEEE*, Jingyuan Zhu, Yuan Yang, Xiao Li, Fangzheng Jiao, *Graduate Student Member, IEEE*, Shicheng Wang, Chunming Hu, and Mingwei Xu

***Abstract*—Collective communication has become increasingly crucial as large language models rapidly evolve, leveraging RDMA networks for inter-node connectivity to achieve high throughput and low latency. However, RDMA inevitably faces Network Performance Anomalies (NPAs) due to mechanisms like line-rate start and PFC flow control. Collective communication usually utilizes multiple flows simultaneously, which puts new challenges for efficient diagnosis of NPAs. Existing works fall short in analyzing co-flow patterns, which cannot diagnose RDMA NPAs in collective communication accurately and efficiently. To address this, we propose VEDRFOLNIR, an accurate and efficient diagnosis system for RDMA NPAs in collective communication. VEDRFOLNIR (1) constructs a waiting graph through algorithm decomposition, (2) adaptively detects anomalies while efficiently collecting diagnostic data, and (3) precisely analyzes performance bottlenecks and identifies root causes. We implement an open-source prototype of VEDRFOLNIR. Evaluation results show that VEDRFOLNIR can achieve accurate diagnosis results and reduce 60%-98% overhead compared to the state of the art.**

***Index Terms*—RDMA, collective communication, network performance anomaly, diagnosis.**

I. INTRODUCTION

WITH the rapid growth of large language models (LLMs), collective communication has become an essential component that determines the efficiency of data exchange and parallel computation among GPUs. RDMA (Remote Direct Memory Access), with its high bandwidth and low latency characteristics, is widely used in data transmission for collective communication. As data center network infrastructure evolves, more clusters are transitioning from InfiniBand (IB) networks to more versatile high-speed Ethernet, commonly using the RDMA over Converged Ethernet version 2 (RoCEv2) protocol [1] for data exchange [2], [3], [4], [5], especially in areas like LLM training [6], [7], [8], [9], [10], [11], [12]. For example, Meta's 16,000-GPU cluster [6] utilizes RoCEv2 for LLM training.

To ensure high performance, in RoCEv2, PFC (Priority Flow Control) [13] is typically used to achieve a lossless network, along with congestion control algorithms such as DCQCN [14] and Swift [15] to manage flow control and ensure transmission efficiency. Complex collective communication operations performed on this basis introduce additional complexity for diagnosing performance anomalies. On one hand, collective communication is dynamic during execution, with data dependencies between flows and potential changes in the flows across different stages. On the other hand, RDMA mechanisms such as PFC flow control and line-rate start make network congestion more frequent and complex. As a result, collective communication is prone to encounter various types of network anomalies (e.g., flow contention, PFC backpressure), leading to degraded performance.

Existing research faces significant challenges in accurately and efficiently diagnosing RDMA Network Performance Anomalies (NPAs) in collective communication. Firstly, existing diagnostic approaches predominantly focus on single-flow level analysis (e.g., Hawkeye [16], [17], SpiderMon [18], Sonata [19], and *Flow [20]), neglecting complex communication paradigms such as collective communication. In contrast to the static nature of single-flow communication, flows in collective communication are dynamic and may evolve over time, potentially having dependencies on other flows as the communication progresses. For instance, in the Halving and Doubling [21] algorithm, the destination of a flow can change multiple times, and such changes depend on the data transmitted by other flows. Current diagnostic methods fail to capture these intricate dynamic behaviors and the provenance between flows within the collective communication, rendering them inadequate for accurately analyzing RDMA NPAs in collective

Received 4 February 2026; revised 1 August 2026; accepted 17 October 2026; approved by IEEE TRANSACTIONS ON NETWORKING Editor K. Lee. Date of publication 30 October 2026; date of current version 6 November 2026. This work was supported in part by the National Key Research and Development Program of China under Grant 2024YFB4505604, in part by the National Natural Science Foundation of China under Grant 62402025, in part by the Open Research Fund of State Key Laboratory of Internet Architecture under Grant HLW2025ZD01, and in part by the Fundamental Research Funds for the Central Universities. An earlier version of this paper was presented at the Conference of INFOCOM 2026 [DOI: 10.1109/INFOCOM59046.2026.11571395]. *(Corresponding author: Menghao Zhang.)*

Yuxuan Chen and Menghao Zhang are with the School of Software, Beihang University, Beijing 100191, China, and also with the State Key Laboratory of Internet Architecture, Tsinghua University, Beijing 100084, China (e-mail: yuxuanchen@buaa.edu.cn; zhangmenghao0503@gmail.com).

Jingyuan Zhu, Fangzheng Jiao, and Chunming Hu are with the School of Software, Beihang University, Beijing 100191, China (e-mail: jingyuanzhu@buaa.edu.cn; jiaofz@buaa.edu.cn; hucm@buaa.edu.cn).

Yuan Yang, Xiao Li, Shicheng Wang, and Mingwei Xu are with State Key Laboratory of Internet Architecture, Tsinghua University, Beijing 100084, China (e-mail: yangyuan_thu@mail.tsinghua.edu.cn; li-x23@mails.tsinghua.edu.cn; wangshicheng1225@163.com; xumw@tsinghua.edu.cn).

Digital Object Identifier 10.1109/TON.2026.3748812

2998-4157 © 2026 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission.
See https://www.ieee.org/publications/rights/index.html for more information.

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

communication. Second, due to the involvement of numerous nodes and flows in collective communication, achieving accurate diagnosis with minimal overhead presents another significant challenge. Existing approaches collect telemetry data across all switches (e.g., NetSight [22] and PINT [23]) or at all times (e.g., Hawkeye [16], [17]). While ensuring accuracy, these methods introduce substantial communication and computational overhead, leading to potential performance bottlenecks. Thus, minimizing the impact on network performance while maintaining diagnostic accuracy remains a critical issue that requires significant attention.

To address the above challenges, we propose VEDRFOLNIR, an accurate and efficient RDMA network performance anomaly diagnosis system for collective communication. In terms of data collection, VEDRFOLNIR monitors collective communication performance information on the host side and performs network tracing when performance anomalies occur. For diagnostics, VEDRFOLNIR primarily utilizes the waiting relationships between flows to construct a waiting relationship graph for analyzing performance bottlenecks in collective communication, and builds network provenance graphs for root cause analysis and flow impact assessment. VEDRFOLNIR is built with the following three key modules. First, VEDRFOLNIR characterizes collective communication performance by decomposing the communication algorithm into steps and defining a waiting relationship graph (§III-B). Second, VEDRFOLNIR employs a step-aware adaptive detection mechanism to reduce redundant data collection and an interval PFC awareness mechanism to precisely diagnose anomalies (§III-C). Third, VEDRFOLNIR uses the collected data to construct waiting and network provenance graphs for comprehensive root cause analysis (§III-D). We implement a prototype on NS3, which is open-source at GitHub [24]. Evaluations show that VEDRFOLNIR can accurately diagnose collective communication NPAs while reducing overhead up to 98% compared to Hawkeye or full polling. The additional CPU overhead caused by VEDRFOLNIR is negligible.

Contributions. We analyze the complexity that collective communication introduces to RDMA NPA diagnosis (§II) and propose VEDRFOLNIR, an accurate and efficient RDMA NPA diagnosis system for collective communication (§III). We implement a prototype using NS-3 simulation and experimental results show that VEDRFOLNIR can accurately diagnose performance issues while reducing substantial overhead (§IV).

II. Background and Motivation

Collective communication adds significant complexity to the diagnosis of NPAs, making existing work inadequate for effective diagnosis. VEDRFOLNIR focuses on diagnosing collective communication performance anomalies resulting from RDMA NPAs, and aims to answer the following key question: *what network performance anomaly is responsible for the degraded collective communication?*

A. Network Complexity of Collective Communication

The network complexity of collective communication primarily consists of two main aspects:

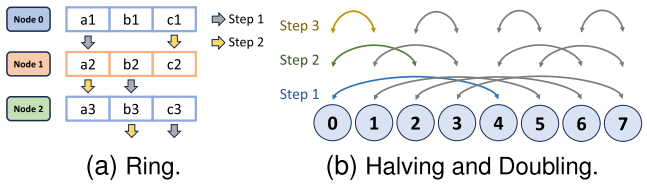


Fig. 1. Examples of collective communication algorithms.

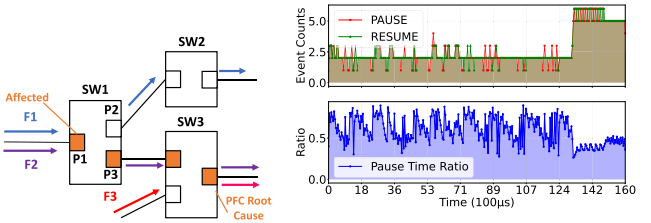


Fig. 2. Complexity introduced by RDMA.

Co-flow pattern in collective communication. Collective communication typically employs specific algorithms (e.g., Ring, Halving and Doubling) to enhance communication efficiency (Figure 1), thereby introducing additional complexity for network anomaly diagnosis. First, unlike single-flow applications, collective communication requires collaborative efforts from multiple flows originating at multiple nodes. Second, collective communication changes dynamically during execution. The specific flows involved differ across distinct algorithmic phases. Third, there are data dependencies between flows. Host B may receive the data transmitted by flow A before it can start flow B to transmit the data to other nodes.

Complexity Introduced by RDMA. The complexity of RDMA primarily stems from the PFC flow control mechanism in RoCEv2. RDMA requires PFC to achieve a lossless network, thereby ensuring superior transmission performance. With PFC, when a switch queue exceeds a certain threshold, it sends a PAUSE frame to its upstream devices. After that, the upstream device stops transmitting on the link until it receives a RESUME frame. The following two characteristics of PFC pose challenges for diagnosis: 1) Propagation. When encountering network congestion, PFC exerts a backpressure effect, cascading PAUSE frames upstream, which can lead to congestion over a wider network area. In the example shown in Figure 2a, the PFC propagation triggered by a culprit flow (F3) can even affect a flow (F1) that does not share forwarding paths with it, increasing the difficulty of root cause localization [14], [17]. 2) Intermittency. The pauses caused by PFC are not necessarily sustained; rather, they may repeatedly switch between pause and resume states over a period (Figure 2b). As shown in Figure 2a, if the culprit flow (F3) is paused at a certain moment, flow F2's congestion is temporarily alleviated and its rate resumes. In this case, our previous method [25] relying on instantaneous PFC detection would fail to identify flow F1 as PFC-affected, missing the root cause.

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

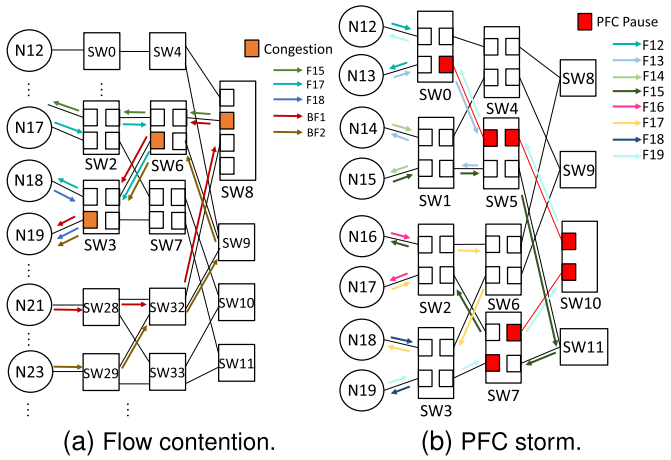


Fig. 3. Anomaly cases.

Apart from PFC flow control, other mechanisms in RDMA also make network congestion more likely to occur. For example, RDMA does not have a slow start process like TCP, but instead sends at line speed initially, which makes network congestion more frequent and transient given the shallow buffers of commercial RDMA switches [3]. Additionally, RDMA hardware struggles to achieve fine-grained load balancing [26].

B. NPAs in Collective Communication

Network performance anomalies are complex and diverse. Figure 3 illustrates two fundamental types of anomalies: 1) Flow contention. In large-scale clusters, many tasks share the same infrastructure, making flow contention in the network inevitable. Figure 3a shows an instance of collective communication using the Ring algorithm among 8 nodes (Nodes 12-19), accompanied by two background flows (BF1 and BF2). Due to path overlap, the background flows collide with the collective communication flows, causing congestion at multiple switch ports and degrading transmission performance. 2) PFC storms. A PFC storm refers to the sustained injection of PFC frames, often caused by hardware bugs, which triggers widespread propagation of PFC and leads to network-wide performance anomalies. Figure 3b displays the flows of an 8-node collective communication. Persistent PFC injection occurs at a port on Switch 0, causing Flow 12 to be halted across multiple switches, severely impacting performance.

In real network environments, flow contention and PFC propagation often occur in combination. Depending on their different causes or characteristics, some studies [17], [18] have also summarized the following other common anomalies: 1) Load imbalance. Due to misjudgments in load balancing policies (e.g., ECMP), traffic from collective communication or other flows may not be distributed ideally across the network topology, leading to contention between different flows. 2) Loops. Due to asynchrony among switches during network reconfiguration, some switches may form a forwarding loop. A flow may be forwarded from a switch port and, after multiple hops, return to the same port. 3) PFC backpressure by flow contention. When flow contention rapidly fills the switch queues, the PFC is triggered, potentially resulting in the situation illustrated in Figure 2a. 4) PFC deadlock. When PFC propagation forms a cycle, all flows at the cyclic locations are halted, resulting in a deadlock.

C. Existing Solutions Fall Short

Lack of analysis capability for co-flow patterns. Existing work [16], [17], [18], [19], [20], [27], [28] primarily focuses on single-flow analysis and overlooks more complex paradigms such as collective communication, leading to insufficient co-flow analysis. Due to the complexity of collective communication, an anomaly in one flow is likely to delay the start time of another flow that depends on it. Therefore, diagnosing collective communication performance must consider multiple flows simultaneously and analyze them based on their interrelationships. However, existing host-based approaches [29], [30], [31] mainly focus on inferring anomalies in the network, while switch-side methods [18], [19], [20], [27] primarily use programmable switches to directly collect network data. These approaches focus solely on the network, with analysis conducted at the flow or packet level, without considering potential relationships at the application level.

Telemetry collection incurs significant redundant overhead. Some works [22], [23] collect telemetry information from all switches, ensuring accuracy but introducing substantial communication or analysis overhead. Other works [18], [32], while analyzing queue contention on only a subset of switches, fail to account for the PFC chaining propagation characteristics in RDMA, resulting in inaccurate information. Hawkeye [17] achieves single-flow-level accuracy and efficiency by devising a PFC tracing method, but it still lacks design considerations for collective communication. For example, when Hawkeye is applied to collective communication anomaly detection, its mechanically set anomaly detection trigger mechanism can cause repeated triggers at multiple communication nodes within a short period, collecting redundant data and introducing significant overhead.

Lack of precise root cause diagnosis at the collective communication level. Existing diagnostic approaches can be categorized into three types: flow statistics-based [19], [20], [27], [28], [31], [33], queue measurement-based [32], [34], and provenance-based [16], [17], [18], [35]. First, methods based on flow statistics and queue measurements contribute primarily by statistical analysis of flows or queues, but they struggle with precise anomaly localization (e.g., detecting switches related to loops or PFC deadlocks). Second, provenance-based methods address the aforementioned limitations, but still fail to consider the complexities of collective communication scenarios, unable to analyze the performance bottlenecks and locate the root causes.

III. VEDRFOLNIR DESIGN

A. Overview

To diagnose collective communication performance anomalies resulting from RDMA NPAs accurately and efficiently, VEDRFOLNIR consists of the following key steps. 1) To capture the dynamics of collective communication and the data

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

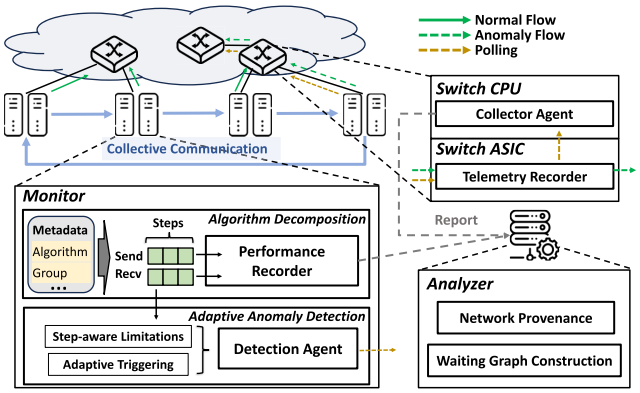


Fig. 4. VEDRFOLNIR framework.

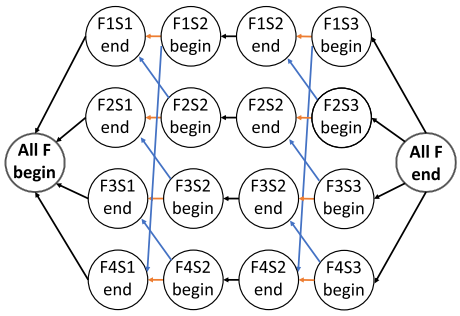


Fig. 5. Waiting graph of a Ring reduce-scatter communication.

dependencies of flows, VEDRFOLNIR algorithmically decomposes collective communications at a per-step granularity and constructs a directed graph to describe the entire collective communication process using the waiting relationships of flows (§III-B). 2) To ensure accuracy with low overhead, VEDRFOLNIR employs a step-aware adaptive mechanism to limit unnecessary detection triggers, and leverages coordinated data collection through host monitoring and network telemetry to efficiently and accurately support diagnostics of collective communication (§III-C). 3) To enable effective root cause analysis, VEDRFOLNIR constructs multi-level graphs for bottleneck analysis and root cause localization, and evaluates the primary contributing flows to flow contention (§III-D).

Figure 4 shows the overall architecture of VEDRFOLNIR. First, the monitor (located on each host) pre-executes algorithmic decomposition and monitors performance metrics of the flows by steps during collective communication in real time. These metrics are reported to the analyzer. Second, when performance anomalies occur, VEDRFOLNIR adaptively adjusts the trigger for anomaly detection in monitors. Monitors send notification packets based on their wait relationships to control the frequency of anomaly detection. During detection, the selected hosts initiate the process by transmitting polling query packets. These packets trigger the collection of telemetry data from all relevant switches involved in the anomaly. The collected telemetry data is then reported to the analyzer. Finally, the analyzer constructs both a waiting graph and network provenance graphs. Through comprehensive analysis of these graphs, VEDRFOLNIR provides structured diagnostic results.

B. Algorithm Decomposition and Graphical Description

Algorithm decomposition. VEDRFOLNIR decomposes the algorithm into multiple steps, describing the distinct phases of various flows in a collective communication. For a flow originating from a particular node, the transmitted data chunk or the destination address changes between each step. For example, in Figure 1a demonstrating the Ring algorithm, the data chunk transmitted by flow F0 (originating from Node 0) changes from chunk A to chunk C during the transition from step 1 to step 2. Similarly, Figure 1b shows that in the Halving and Doubling algorithm, the destination of flow F0 shifts from Node 4 in step 1 to Node 2 in step 2. When either of these two transitions occurs, we say that F0 has gone through a step. Other collective communication algorithms can also be decomposed using a similar approach. The fundamental idea is to break them down based on the data dependencies between flows, where each step of a flow serves as a unit that either waits for or is waited for by other flows. In practice, step definitions can be finalized prior to the actual execution of communication, as the algorithm is usually predetermined. Taking NCCL as an example, the initialization phase preceding communication already determines the algorithm selection, after which the communication process commences.

Waiting graph definition. After the step splitting of the flows, we construct a waiting graph to describe the collective communication process. Taking the Ring reduce-scatter illustrated in Figure 5 as an example, we define 1) vertices: the start or end of each step of each flow. F_iS_j represents step j of flow i. Specifically, the end of the final step for all flows serves as the graph's source, while the start of the first step for all flows serves as the sink. 2) directed edges: indicating waiting dependencies between vertices. For example, due to data dependency, the start of F2S2 needs to wait for the end of F1S1, which is represented by a blue edge; meanwhile, the start of F2S2 needs to wait for the end of the previous step F2S1, which is represented by an orange edge; the relationship between the start and the end of step execution within F2S2 itself is represented by a dark edge. 3) weight: indicating the waiting time. The weight of the light-colored (i.e., blue, orange) edge is 0, and the weight of the dark edge is the execution time of this step of the flow.

C. Logging Workflow

VEDRFOLNIR monitors collective communication performance metrics on the host side (§III-C1). Upon detecting performance anomalies, VEDRFOLNIR adaptively triggers network telemetry (§III-C2) and retrieves network information recorded in the switches (§III-C3), enabling subsequent anomaly diagnosis.

1) Performance Monitoring: On the host side, VEDRFOLNIR monitors the waiting state and records performance information.

Waiting status awareness. As discussed in Section III-B, VEDRFOLNIR decomposes the collective communication process using the waiting relationships of flows. However, in

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

TABLE I
WAITING STATE DETERMINATION

Index Comparison (state)	Meaning
Send Steps == Recv Steps (waiting)	The next send step waits for the current receive to complete.
Send Steps < Recv Steps (non-waiting)	Do not wait for others, execute the next send step as soon as current step is finished.

actual execution, the waiting relationship occurs selectively. As shown in Figure 5, the start of F1S2 waits for both the end of F1S1 and F4S1, but due to the sequence of the two events, F1S2 actually waits for only one of them. To enable real-time awareness of the waiting status, VEDRFOLNIR's monitor maintains two queues: a Send Step Queue (SSQ) and a Receive Step Queue (RSQ). During the algorithm decomposition, the flow originating from the current host is divided into multiple steps. The target of each sending step is sequentially enqueued into the SSQ. Correspondingly, the source of the required data for each receiving step (i.e., the source of flow which the next send step waits) is sequentially enqueued into the RSQ. By tracking the index of the currently active step within these queues, VEDRFOLNIR can determine the waiting status of each flow (Table I).

Performance recording. To construct the waiting graph and analyze performance bottlenecks in collective communication, the monitor on each host collects performance metrics for its local flows. Upon completion of each flow step, the host reports its 5-tuple, data volume transferred, start time, end time, and the source host of the flow it is waiting for.

2) Step-Aware Adaptive Anomaly Detection: To address the issues of high redundant overhead and lack of multi-flow analysis in previous work (e.g., Hawkeye [16], [17] and SpiderMon [18]), VEDRFOLNIR primarily improves in two aspects: 1) Step-aware detection trigger constraints. Under ideal conditions, the execution time of a single step in a collective communication flow is typically short (e.g., approximately 30 ms for transmitting 300 MB of data at 100 Gbps). However, our testing indicates that, in commonplace flow contention, the communication flow execution time can easily experience a two-fold or greater delay. Previous approaches lack temporal management for detection triggers, repeatedly collecting telemetry within the same step and thus generating substantial redundant overhead. Hence, step-level detection granularity is more appropriate. 2) Adaptive detection opportunity allocation. As defined by the waiting graph (§III-B), the overall performance (i.e., total execution time) of collective communication is determined by its critical path, and the affected steps on this path tend to be prolonged. It is appropriate to allocate more detection opportunities to the affected steps based on their wait-for dependencies.

For detection triggering, VEDRFOLNIR primarily makes decisions dynamically based on network topology and user configurations, specifically including the following two aspects. First, host-side triggering based on RTT thresholds. Unlike Hawkeye's fixed RTT threshold settings, VEDRFOLNIR recalculates RTT thresholds based on network topology before each step initiation. This avoids inappropriate RTT settings caused by changes in flow paths (e.g., the algorithm in Figure 1b) or failure to distinguish between different flows. When the monitor detects that the RTT surpasses the threshold, an anomaly detection is triggered. Second, triggering constraints. As shown in Figure 6, to avoid redundancy, VEDRFOLNIR uses configurable parameters to define the number of anomaly detection triggers for each step of a collective communication flow. Based on the estimated FCT (flow completion time, which can be calculated from the network topology or set empirically), it sets a triggering time interval limit. This ensures that triggering opportunities are evenly distributed, thus ensuring coverage as comprehensively as possible.

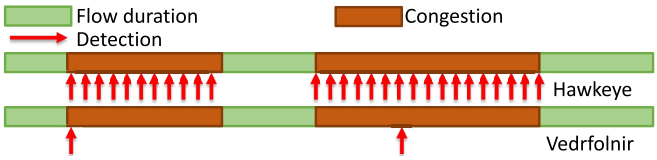


Fig. 6. Examples of detection triggers.

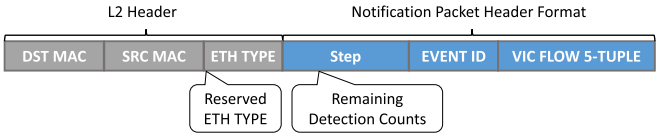


Fig. 7. Notification packet format.

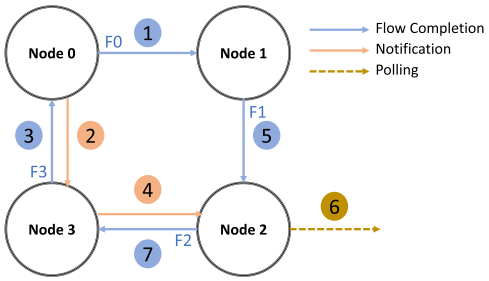


Fig. 8. Example of the adaptive mechanism. The numbers represent the order of events.

For the adaptive mechanism, VEDRFOLNIR mainly achieves adaptation by allowing monitors to observe the waiting states of their respective flows. Specifically, when a certain step of a flow is completed, its monitor sends a notification packet to inform the monitor of the flow waiting for it (if any), transferring its remaining detection opportunities. The notification packet contains detection count information (Figure 7) and is assigned the highest priority to avoid being affected by network congestion. An example is illustrated in Figure 8. F0 completes first, and node 0 transfers its detection opportunities to host 3 by sending a notification packet. Subsequently, F3 completes, and the opportunities are further transferred to host 2. Before F2 completes, F1 finishes earlier. As a result, the slowest host, host 2, obtains the most detection opportunities. In this way, VEDRFOLNIR strives to collect telemetry data for

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

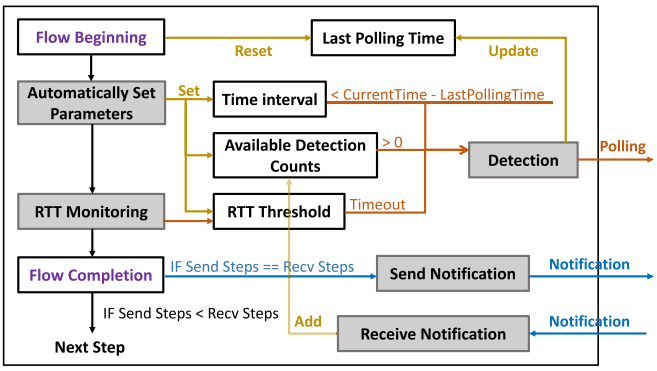


Fig. 9. Algorithmic flow of one step.

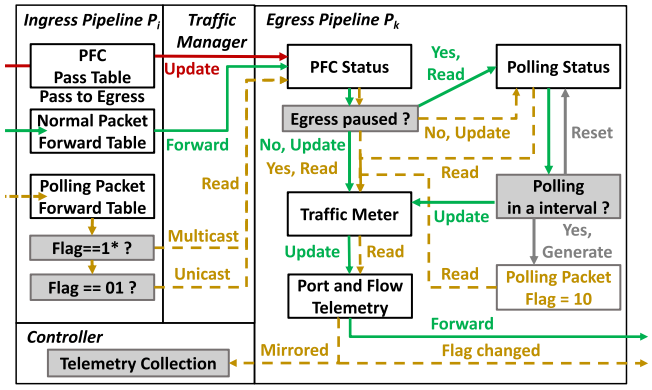


Fig. 10. PFC causality analysis.

the relatively slowest flow at each step, while theoretically ensuring an upper bound on overhead.

VEDRFOLNIR implements the aforementioned design on the host side through a detection agent, which is responsible for the dynamic updating of various parameters, the sending and receiving of notification packets, the judgment of triggering conditions, and other functions. The specific algorithmic flow is illustrated in Figure 9.

3) Telemetry Collection: Figure 10 illustrates the overall telemetry collection process. To address the complexities introduced by PFC, VEDRFOLNIR (1) adopts a telemetry collection mechanism similar to Hawkeye [17] to accurately localize the root cause of PFC (addressing Fig. 2a); (2) introduces an interval PFC awareness mechanism to avoid triggering detection precisely during PFC recovery periods (addressing Fig. 2b).

Telemetry logging mechanism. For telemetry data logging, switches periodically record flow-level telemetry data (e.g., flow' 5-tuple, packet count, queue depth) and port-level telemetry data (e.g., traffic volume between ports, number of packets paused by PFC per port, egress queue depth), while also sensing port PFC status. In the detection process, upon receiving a polling packet, the switch data plane propagates it along the flow path and the PFC propagation path. Simultaneously, the switch controller assists in data collection, reporting the telemetry information to the analyzer. The forwarding of the polling packet is controlled by two flag bits in its header: when the lower bit is set to 1, the polling packet is forwarded along the victim flow path; when the higher bit is set to 1, the polling packet is forwarded to the PFC propagation path for root cause tracing.

Interval PFC awareness mechanism. Unlike Hawkeye's high-frequency detection triggering, the trigger mechanism of VEDRFOLNIR may lead to the detection failure issue mentioned in §II-A. Given the high-frequency switching nature of PFC states, it is necessary to shift PFC awareness from an instantaneous basis to an interval basis. As shown in Figure 10, when a polling packet is received, if the current egress queue is not paused, the polling state is updated and the current timestamp is recorded. Within a subsequent time interval (e.g., 100 microseconds), if the egress queue is paused by PFC, the paused traffic will trigger an additional detection process. A polling packet will be generated in the data plane with its higher flag bit set to 1, further initiating telemetry collection and PFC root cause tracing.

D. Comprehensive Root Cause Analysis

Considering the complex flow patterns of collective communication, the root cause diagnosis of VEDRFOLNIR should be able to answer the following questions: 1) Where are the performance bottlenecks in collective communication? Which specific anomalous flows are slowing down the collective communication? 2) What is the underlying network root cause of the performance anomaly? Which flows or switches are involved? 3) If other flows are causing the decrease in collective communication performance, how significant is the impact of these flows on collective communication performance?

1) Graph Construction: VEDRFOLNIR constructs a waiting graph to analyze the performance bottlenecks of collective communication and identify critical flows. For each step of the collective communication, it constructs provenance graphs to locate network root causes and evaluate the impact of flows.

Waiting graph of collective communication. The definition of the waiting graph is described in §III-B. During runtime, the analyzer queues the collected data entries in order of their completion time and constructs the waiting graph sequentially according to the queue order. Based on the real-time collected data regarding waiting relationships and their corresponding timestamps, the analyzer constructs corresponding nodes and edges. Algorithm 1 demonstrates the detailed algorithm for graph construction. For simplicity, upon determining that a node is not being waited for (i.e., has an in-degree of zero), the analyzer can recursively prune nodes with an in-degree of zero. After constructing the complete waiting graph, the analyzer calculates the critical path, which represents the performance bottleneck of the collective communication.

Provenance graph of network. VEDRFOLNIR constructs the provenance graphs based mainly on packet-level waiting relationships between flows and ports. Previous works [16], [17], [18] are unable to evaluate the contribution between flows while performing PFC traceability. To support identifying the primary contributing flows, VEDRFOLNIR redefines the graph construction based on Hawkeye [16], [17]. The vertex set V consists of F (flows), P (ports), and CF (collective communication flows), where $CF \in F$ and $F \cup P = V$. The

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

Algorithm 1 The Waiting Graph Construction

Data: L_{CF} - A list containing flow data, where each item represents a single step of a communication flow
Result: G - The waiting graph

Function `build_waiting_graph(`L_{CF}`)`**:**
  Sort L_{CF} in the order of flow completion
  Init graph G with "Begin" and "End" nodes
  Init N_{pre} to record the predecessor node of a given node
  For each host id i, $N_{pre}[i] \leftarrow$ BeginNode("Begin")
  Init W to record the non-predecessor wait relations
  $n_s \leftarrow$ the largest $sendStep$ obtained from L_{CF}
  for cf **in** L_{CF} **do**
    if $cf.sendStep$!= n_s **then**
      $N_{curr} \leftarrow$ EndNode("$F_{cf.sourceID}S_{cf.sendStep}$")
      G.AddNode(N_{curr})
    else
      $N_{curr} \leftarrow$ EndNode("End")
    if $W.get(cf.sourceID)$ is None **then**
      G.AddEdge($N_{pre}[cf.sourceID] \rightarrow N_{curr}$, Weight=$cf.fct$)
    else
      $N_b \leftarrow$ BeginNode("$F_{cf.sourceID}S_{cf.sendStep}$")
      G.AddNode(N_b)
      G.AddEdge($N_b \rightarrow w[cf.sourceID]$, Weight=0)
      G.AddEdge($N_{curr} \rightarrow N_b$, Weight=$cf.fct$)
    if $cf.sendStep == cf.receiveStep$ **then**
      $W[cf.sourceID] \leftarrow$ EndNode("$F_{cf.waitID}S_{cf.receiveStep}$")
    else
      $W[cf.sourceID] \leftarrow$ None
    $N_{pre}[cf.sourceID] \leftarrow N_{curr}$
  return G

directed edge set E describes waiting relationships between vertex, which are categorized into the following three types:

(1) $e(f,p)$: Describes the waiting relationship from a flow to a port. Essentially, a flow waits on a port due to congestion caused by contention among multiple flows at that port. We use the queue depth encountered by the flow's packets while queued at the port as the edge weight. Suppose that a packet pkt of flow f_i enters the queue, and define the number of packets of flow f_j in front of it as $x_j(pkt)$. Then, the waiting weight of flow f_i for f_j is defined as $w(f_i,f_j) = \sum_{pkt \in f_i} x_j(pkt)$, and the weight of the flow for the port is $w(f_i,p) = \sum_{j \neq i} w(f_i,f_j)$.

(2) $e(p,f)$: Describes the waiting relationship from a port to a flow. We mainly use this type of edge to quantify the contribution of the flow to port congestion. Similarly, we use the number of packets contributed by the flow to the port queue depth as the weight of the edge. Assuming that the queue depth of port p is $qdepth(p)$ and the number of packets detected within a certain period is $pkt_num(p)$, the weight of the port for the flow is calculated as $w(p,f_i) = \frac{pkt_num(f_i)}{pkt_num(p)} \times qdepth(p)$.

(3) $e(p_i,p_j)$: Describes the waiting relationship between ports. We use this type of edge to capture PFC causal relationships. When a PFC pause event occurs, an egress port p_j on a downstream switch may halt an egress port p_i on an upstream switch, establishing a waiting relationship from p_i to p_j. Since the measure of the wait is based on the queue depth, in this case, we consider p_i to be waiting for all packets in the p_j queue. However, PFC backpressure can propagate along multiple paths. In order to accurately evaluate the contribution ratio of the path from p_j to p_i, we define the weight from p_i to p_j as the traffic ratio $w(p_i,p_j) = \sum_k \frac{meter(p_i,p_j)}{meter(p_k,p_j)}$.

2) Anomaly Breakdown: Similar to Hawkeye [17] and SpiderMon [18], VEDRFOLNIR diagnoses specific root causes based on the signatures of different anomaly types. Similarly, VEDRFOLNIR can extend the range of detectable anomaly types by adding more signature definitions. Below, we provide examples of diagnosing two anomaly types: (1) *Flow contention.* VEDRFOLNIR identifies flow contention by checking the port vertices in the network provenance graph. When $\exists\, p \in P, \{f_i, cf\} \subseteq F \wedge \{e(f_i,p), e(cf,p)\} \subseteq E \wedge f_i \neq cf$, it is considered that f_i contends with the collective communication flow cf at port p. (2) *PFC backpressure.* VEDRFOLNIR identifies PFC backpressure by tracing the PFC spreading path. This occurs when $\exists\ p \in P, cf \in F, e(cf,p) \in E \wedge \exists\ p_j \in P, e(p,p_j) \in E$. Following the PFC spreading path. VEDRFOLNIR can pinpoint the root cause location of the PFC event. PFC storms typically also exhibit backpressure phenomena and is identified analogously.

Regarding the use of provenance graphs, users can define the anomalies to be reported based on the features of these graphs. A provenance graph can depict the locations and quantities of flow contentions, as well as the locations and propagation paths of PFC, thereby clarifying the progression of anomalies and the flow-level causal order. To add a new type, such as incast, characterized by multiple flow nodes connecting to a specific edge switch node, users may define its features using nodes and edges of the provenance graph. By taking the provenance graph as input, the system determines whether an incast-type anomaly exists in the output. However, anomaly types serve merely as output formats. Identifying the underlying root cause requires direct use of the provenance graph itself.

3) Contributor Rating: In scenarios where collective communication contends with other flows, VEDRFOLNIR is capable of quantifying the contribution of these contending flows to identify the main contributors. In this way, VEDRFOLNIR can make recommendations on the priority of the actions of the network operations personnel.

Evaluating contribution to a collective communication flow. First, starting from all collective communication flows (CF), we obtain the largest connected subgraph of the network provenance graph, then all flows $f \notin CF$ belong to the evaluation object. Algorithm 2 demonstrates the detailed algorithm for flow-level computation. For each flow, we calculate the contribution of the flow to a vertex by traversing along the reverse direction of directed edges (i.e., the direction of being waited for). Suppose that the neighboring vertices of flow f_i are p_1 and p_2, where p_2 is downstream of p_1. Then the impact score of f_i is $R(f_i,p_2) = w(p_2,f_i)$, $R(f_i,p_1) = w(p_1,f_i) + R(f_i,p_2) \times w(p_1,p_2)$. Generally, the contribution of a flow to a port is:

$$R(f_i,p_j) = w(p_j,f_i) + R(f_i,p_k) \times w(p_j,p_k), e(p_j,p_k) \in E \quad (1)$$

Given a collective communication flow cf, its neighboring vertices set P_{cf}, and $R(f_i,p_k)$ for any $p_k \in P_{cf}$. Since cf

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

Algorithm 2 Rating for Flow-Level Impact

Data: G - Network provenance graph; V_P - Port vertex in graph G; V_{CF} - Communication flow vertex in graph G; V_F - Other flow vertex in graph G
Result: R - Rating result

Function `parse_network_graph(G)`:
  Init graph H for computational convenience
  Copy $G.vertex$ to H
  for $v_i \to v_j$ **in** $G.edges$ **do**
    if $v_i \in V_p, v_j \in V_F$ ***or*** $v_i, v_j \in V_P$ ***or*** $v_i \in V_{CF}, v_j \in V_p$ **then**
      H.AddEdge($v_j \to v_i$)
  Sort vertex in H according to topological order
  Init R_P, R_{CF} for results record
  for v **in** $topo_order$ **do**
    for vs **in** $v.successors$ **do**
      if $v \in V_F$ **then**
        $R_P[vs][v] \leftarrow R_P[vs][v] + Weight(v \to vs)$
      else if $v, vs \in V_P$ **then**
        for all flow v_f recorded in $R_P[v]$, $R_P[vs][v_f] \leftarrow R_P[vs][v_f] + R_P[v][v_f] \times Weight(v \to vs)$
      else if $v \in V_P, vs \in V_{CF}$ **then**
        for all flow v_f recorded in $R_P[v]$, if $edge\ v_f \to vs$ exists, $R_{CF}[vs][v_f] \leftarrow R_{CF}[vs][v_f] + R_P[v][v_f] - Weight(v_f \to v) + Weight(v_f \to vs)$ else $R_{CF}[vs][v_f] \leftarrow R_{CF}[vs][v_f] + R_P[v][v_f]$
      else
        do nothing
  return R_{CF}

and f_i may directly cause flow contention at port p_k, in this case, the impact of f_i on cf is $w(cf, f_i)$ rather than $w(p_k, f_i)$. Therefore, the contribution of flow f_i to cf is calculated as:

$$R(f_i, cf) = \sum_{p_k \in P_{cf}} [(w(cf, f_i) - w(p_k, f_i)) \times \mathbb{I}(e(f_i, p_k) \in E) + R(f_i, p_k)] \quad (2)$$

where $\mathbb{I}(\cdot)$ is the indicator function that determines whether f_i and cf contend at p_k; $w(cf, f_i)$ can be derived via a replay algorithm [17].

Evaluating contribution to collective communication. Since the performance of collective communication is mainly determined by the flow on the critical path in the waiting graph, we obtain the impact score of flow f_a to the overall collective communication by weighting the sum of the critical flow scores $R(f_a, cf_i)$ at step i. Let the execution time of step i be $exec_time(i)$ and the expected execution time be $expect_time(i)$, the contribution of flow f_a to the overall collective communication is:

$$R(f_a) = \sum_i R(f_a, cf_i) \times \beta(i) \times \frac{exec_time(i) - expect_time(i)}{\sum_k (exec_time(k) - expect_time(k))} \quad (3)$$

where $expect_time(i)$ can be derived theoretically or empirically. $\beta(i)$ is a weight determined by the flow size. Since each polling operation only captures transient snapshots of network state, sometimes smaller flows may occasionally accumulate a large number of packets locally during polling windows, leading to inflated scoring values despite their negligible actual impact. To compensate for this inaccuracy, we introduce $\beta(i)$ to apply a bounded correction. It is computed as follows:

$$\beta(i) = \frac{flow_size(i)}{\sum_k flow_size(k)},\ flow_k \in critical\ path \quad (4)$$

By scoring each flow, VEDRFOLNIR finds the main contributors to flow contention.

IV. EVALUATION

The prototype of VEDRFOLNIR in NS-3 simulator is implemented with about 1000 lines of additional C++ code and the source code is publicly available at Github [24]. The VEDRFOLNIR prototype in real testbed is implemented with about 500 lines of additional C++/Python codes based on the system of our partner company. Through both large-scale simulation and real testbed experiments, we evaluate VEDRFOLNIR to answer the following questions: (1) Can VEDRFOLNIR diagnose RDMA NPAs in collective communications accurately and efficiently (§IV-B)? (2) How effective is design techniques in VEDRFOLNIR (§IV-C)? (3) How does VEDRFOLNIR perform under real-world cases (§IV-E)?

A. Experimental Setup

Topology. We set up a standard Fat-Tree (K= 4) topology with 20 switches on NS3. The link bandwidth is 100 Gbps and the link delay is 2 us. Our real testbed experiments use four H100s connected via RoCE NICs, each machine having a total of 192 CPU cores and 50 GB of memory.

Workload. We simulate an empirical collective communication workload derived from an analysis of LLM training [36]. Specifically, 97% of collective communication operations are AllReduce or AllGather, each with a data size of 360 MB per traffic. For simplicity, we adopt 360 MB per flow in a similar way and use the redRing and Halving-and-Doubling algorithms for AllGather as the workload.

Anomaly construction. We construct four anomaly scenarios for comprehensive testing. `1) Flow contention` are constructed by injecting flows into the existing workload. We construct 60 cases, with the number of flows uniformly distributed in 1-6, sizes randomly chosen in 20 MB-1 GB, and start times randomly distributed in 0-200 ms. These flows are placed randomly but deliberately set to collide with collective communication flows. For diagnosis, we define detecting all injected flows as a true positive, detecting only some flows as a false positive, and failing to detect any anomaly as a false negative. `2) Incast` includes 60 cases in a similar way to flow contention. Specifically, the number of flows changes to 3-8, with random positions but multiple flows targeting the same node. Flow sizes are randomly selected in 20-200MB, and all flows start simultaneously. The diagnostic criteria are the same as for flow contention. `3) PFC storm` is constructed by injecting continuous PFC frames. In 40 cases, the injection point is selected from the switch ports along the paths of 4 collective communication flows. The start time is random in 0-150 ms, and the duration is random in 10-100 ms. For diagnosis, tracing to the source port where the

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

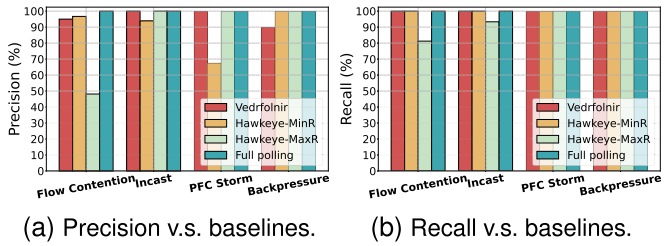


Fig. 11. Diagnosis accuracy under the ring workload.

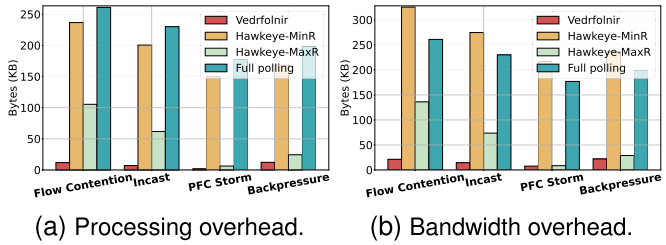


Fig. 12. Diagnosis overhead under the ring workload.

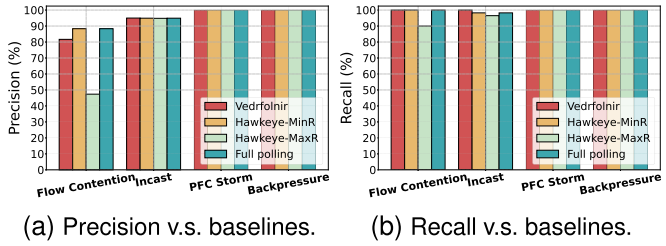


Fig. 13. Diagnosis accuracy under the HD workload.

PFC occurred is defined as a true positive, merely reporting the presence of PFC is defined as a false positive, and failing to detect any anomaly is defined as a false negative. 4) PFC Backpressure is similar to PFC storms. It includes 60 cases. The unique feature is that the PFC does not originate on the collective communication path but affects collective communication through multi-hop propagation. We generate PFC using incast traffic and design propagation paths partially overlapping collective communication flows. The diagnostic criteria are the same as for PFC storms.

Baseline systems. We compare VEDRFOLNIR with two baseline systems. 1) Hawkeye [17]: A state-of-the-art RDMA NPA diagnosis system that collects telemetry data from PFC-related paths and diagnoses root causes. When anomaly detection, Hawkeye sets a fixed RTT threshold for all flows, whereas the RTTs of different flows in collective communication vary. For accurate evaluation, we use Hawkeye-MaxR to denote setting Hawkeye's RTT threshold to 120% of the maximum RTT, and Hawkeye-MinR to denote setting the RTT threshold to 120% of the minimum RTT. 2) Full polling: Continuously collects telemetry data from all switches in the network.

B. Diagnosis Effectiveness

Precision and recall. Our initial analysis focuses on the precision and recall of different solutions. For each scenario, we demonstrate the precision and recall with optimal parameters. The experimental results are presented in Figure 11a, Figure 11b, Figure 13a and Figure 13b. VEDRFOLNIR demonstrates excellent precision and recall across all scenarios. In flow contention, Hawkeye-MaxR tends to easily overlook flows with small RTTs, resulting in poor performance. Additionally, due to its detection-triggering design, Hawkeye is prone to collecting several pieces of telemetry data within tens of microseconds, though only one of them may actually be needed. In the source code of Hawkeye, to avoid excessive processing overhead, only one piece of data is retained every 50 microseconds, which can inadvertently lead to valid data being discarded. As a result, it can be observed that Hawkeye-MinR collects more redundant data and performs even worse. VEDRFOLNIR does not have this issue due to its restricted detection design.

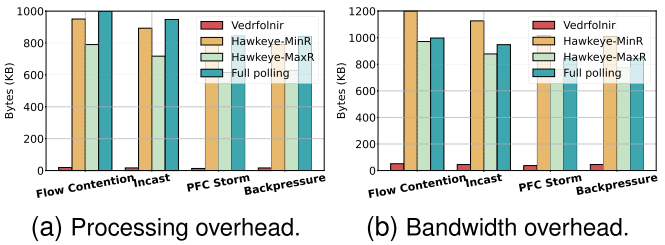


Fig. 14. Diagnosis overhead under the HD workload.

Processing overhead. We utilize the size of telemetry packets collected for diagnostic purposes to evaluate the processing overhead of different approaches. Figure 12a and Figure 14a show the average processing overhead for all samples in each scenario. As shown in Figure 12, VEDRFOLNIR demonstrates the lowest processing overhead across all tested scenarios, with telemetry data volume consistently maintained around 10KB, achieving 60%–98% savings compared to Hawkeye. Due to the low RTT threshold, Hawkeye-MinR continuously triggers detection for flows with long RTTs, resulting in high overhead. For the reason illustrated in Figure 6, Hawkeye-MaxR also performs poorly. However, Hawkeye-MaxR shows better performance in PFC anomalies because Hawkeye measures RTT through ACKs for each packet. When persistent PFC halts an entire flow, no packets are sent, and thus no detection is triggered, leading to lower overhead. In contrast, Full polling demonstrates the upper bound of overhead by continuously triggering detection during collective communication.

Bandwidth overhead. We measure the additional bandwidth overhead introduced by different methods during monitoring. The bandwidth overhead here includes polling during detection, notification packets, and switch telemetry reports. As illustrated in Figure 12b redand Figure 14b, VEDRFOLNIR still achieves the lowest bandwidth overhead across all scenarios. Compared to processing overhead, the bandwidth overhead of VEDRFOLNIR and Hawkeye is higher. However, Hawkeye's overhead increases more noticeably due to repeated anomaly detection. Full polling simulates the scenario where switches continuously and autonomously report telemetry data, and thus excluding detection overhead.

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

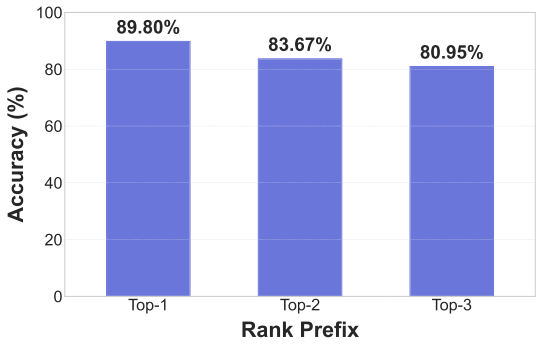


Fig. 15. Top-k accuracy of flow contribution ranking.

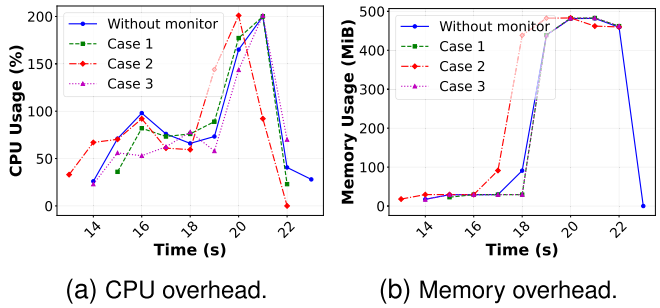


Fig. 16. Overhead in real testbed.

Contributor rating validation. To validate whether the contributor rating reflects the actual impact of background flows, we conduct per-flow removal experiments on the 4-flow contention cases under the standard Fat-Tree (K=4) topology. For each case, we first record VEDRFOLNIR's ratings with all injected flows, then rerun the case while removing one background flow at a time, using the reduction in the overall collective communication completion time as the empirical contribution of that flow. Overall, VEDRFOLNIR ranks 88.81% of background flows consistently with their empirical performance impact. As shown in Figure 15, the top-1, top-2, and top-3 ordered ranking accuracies of cases are 89.80%, 83.67%, and 80.95%, respectively. We do not evaluate the top-4 because, due to the randomness in flow generation, the fourth-ranked flow in many cases does not lie on the critical path and has negligible impact. These results indicate that VEDRFOLNIR's contributor rating can effectively distinguish the relative impact of different background flows, especially for identifying the most influential contributors.

Real testbed evaluation. Compared to Hawkeye, VEDRFOLNIR introduces an additional host-side monitor, which incurs extra CPU overhead during collective communications. We conduct real testbed experiments using NCCL, executing a 4-node AllGather operation with 1 GB transmission volume. Performance collection, notification packet transmission, and data reporting are triggered during the corresponding phases. Figure 16 illustrates the computational and memory overhead. When comparing three test cases with scenarios without the monitor, the additional CPU and memory overhead is practically negligible.

C. Ablation Study

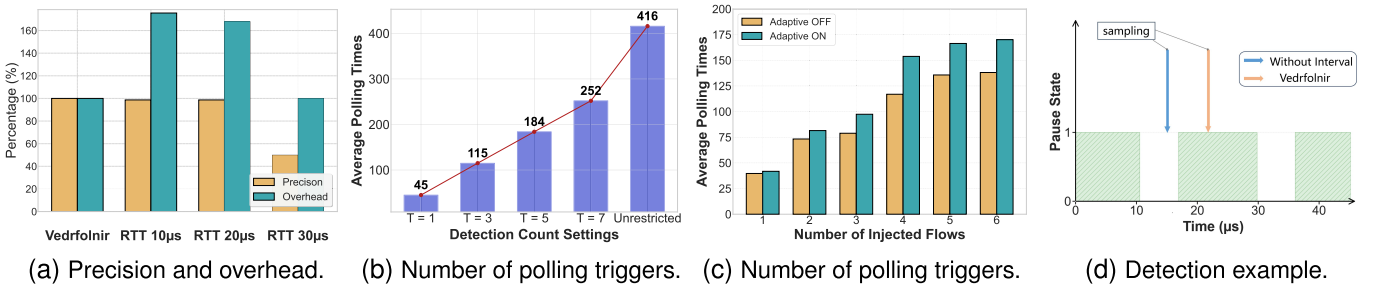
Step-aware mechanism. We validate the effectiveness of the step-aware mechanism from two aspects: 1) Step-grained RTT thresholds settings. In collective communication, RTT varies not only across different flows but may also change across different steps. Figure 17a presents the precision and overhead of VEDRFOLNIR under different fixed RTT thresholds. The experiment is conducted in the flow contention scenario, limiting the number of detections per step to a maximum of 3. 2) Detection triggering count allocation. We evaluate different settings for the number of detections to observe the polling frequency. Figure 17b shows the average results under different configurations in the flow contention scenario. Compared to unrestricted triggering similar to Hawkeye, VEDRFOLNIR achieves significant gains in overhead reduction.

Adaptive mechanism. The design objective of the adaptive mechanism is to increase the frequency of detection triggered for critical flows while ensuring the theoretical overhead upper bound is maintained. Figure 17c illustrates the change in detection frequency after enabling the adaptive mechanism in the flow contention scenario. The RTT threshold is set to 180%, and the per-step detection counts is set to 3. Since detection is only triggered when performance anomalies occur, when the number of interfering flows is high, the collective communication steps experience more severe slowdowns, and the adaptive mechanism correspondingly triggers additional detection more easily. This mechanism can provide more information for diagnosing critical flows and also helps avoid detection failures in some corner cases. In this example, compared to the 98% accuracy achieved with the adaptive mechanism enabled, the accuracy drops to 96.7% when the adaptive mechanism is disabled.

RTT threshold and detection counts. We evaluate two important detection parameters in VEDRFOLNIR: the RTT threshold and the number of detections per step. Figure 18 shows the precision and recall of VEDRFOLNIR for each scenario under different parameter settings. It can be observed that the packet-based RTT timeout triggering method is sensitive to both PFC and incast anomalies, demonstrating robust performance. Theoretically, a higher RTT threshold results in slower detection response after an anomaly occurs. In the flow contention, the 240% RTT setting fails to provide timely detection response for anomalies, leading to poor performance, particularly noticeable in cases with small flows. Increasing the number of detections enhances diagnostic accuracy. Additional detections help locate root causes separately when multiple anomalies happen to coincide in a single flow step. However, due to the random nature of test case generation, such scenarios are infrequent, resulting in limited observable gains. Nevertheless, in the flow contention with RTT threshold set to 120%, raising detection frequency does yield improved diagnostic accuracy. As a trade-off between overhead and accuracy, we recommend setting detection counts as high as possible while keeping the overhead within acceptable limits.

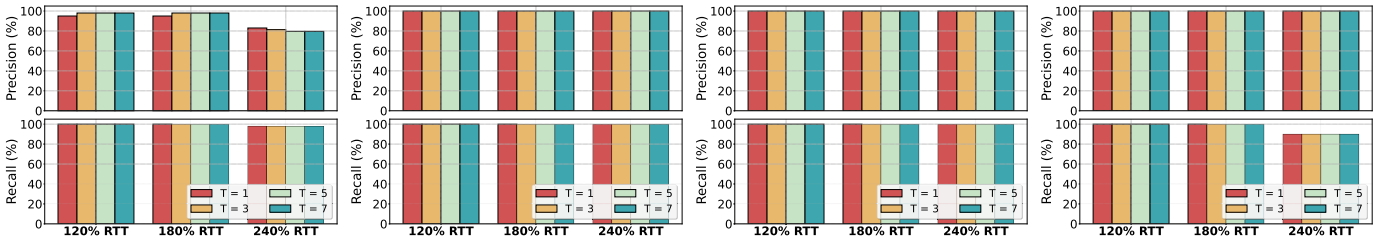
Interval PFC Awareness. We conduct experiments in the PFC backpressure scenario to validate the effectiveness of

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.



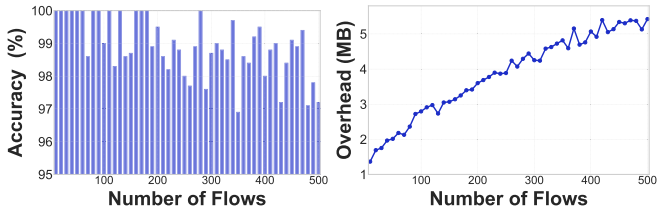
(a) Precision and overhead. (b) Number of polling triggers. (c) Number of polling triggers. (d) Detection example.

Fig. 17. Ablation studies of our methods.



(a) Flow contention. (b) Incast. (c) PFC storm. (d) PFC backpressure.

Fig. 18. Precision & recall over different RTT thresholds and detection counts.



(a) Accuracy. (b) Overhead.

Fig. 19. VEDRFOLNIR performance at scale.

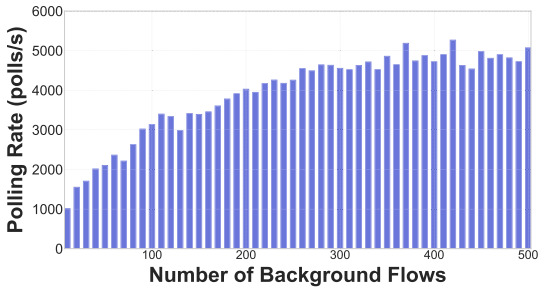


Fig. 20. Polling rate at scale.

interval PFC awareness. The RTT threshold for VEDRFOLNIR is set to 120%, with one detection per step. Compared to the 100% precision of VEDRFOLNIR (Figure 18d), the precision without interval PFC awareness is 63%. Figure 17d illustrates why interval PFC awareness is effective. In this example, the interval mechanism enables detection to wait until the queue is paused by PFC before triggering, thereby avoiding ineffective detection.

D. Performance at Scale

We set up experiments to observe the performance of VEDRFOLNIR at a larger scale. We configure a k=6 fat-tree network topology, comprising 45 switches and 288 hosts, and perform collective communications using the 144-node Ring algorithm. In 50 test cases, background flows are randomly injected to conflict with the collective communication traffic, with the number of background flows uniformly distributed between 50 and 500. The parameter for VEDRFOLNIR is set to a 180% RTT threshold and 3 detection counts per step. Figure 19a illustrates VEDRFOLNIR's diagnostic accuracy, showing that even with a high number of flows, VEDRFOLNIR can still diagnose over 95% of background flows. Figure 19b illustrates the telemetry processing overhead of VEDRFOLNIR's diagnostics. As the number of flows grows, the flow information included in the data collected by VEDRFOLNIR also increases. The processing overhead per switch port exhibits a positive linear relationship with the number of flows per port. Due to VEDRFOLNIR's real-time and restricted triggering mechanism, much redundant overhead is avoided.

Since the packet-level NS-3 simulation at this scale already takes more than one day to complete, further scaling the simulation exceeds our practical evaluation budget. To further evaluate the applicability of VEDRFOLNIR, we analyze the storage overhead of the waiting graph and provenance graph, the aggregate polling rate, and the notification overhead. The memory footprint associated with the waiting graph depends on the scale of the collective communication. Under the current fixed configuration of 144 nodes, the storage required for simulated runs remains approximately 1.36 MB. The raw provenance telemetry overhead increases from 2.62 MB with 10 background flows to 19.86 MB with 500 background flows, with a maximum of 20.17 MB across all cases.

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

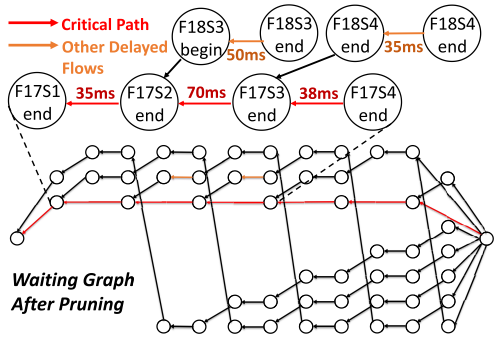


Fig. 21. Waiting graph in the case study.

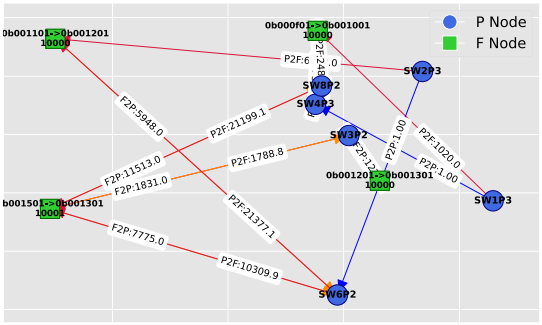


Fig. 22. Network provenance of 10-20 ms in the case study.

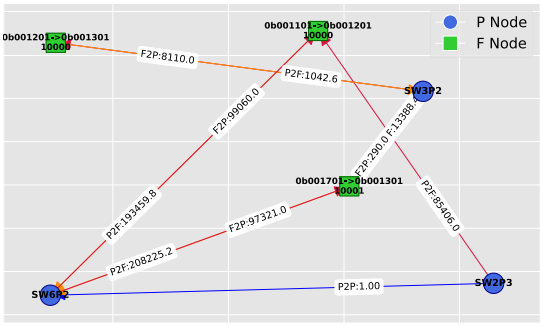


Fig. 23. Network provenance of 130-140 ms in the case study.

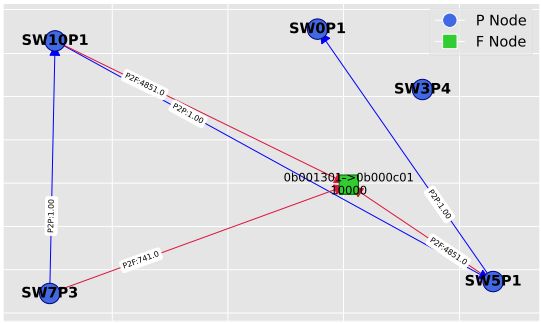


Fig. 24. Network provenance of pfc storm in the case study.

TABLE II
INJECTED FLOWS

Flow	Path	StartTime	Size	Rating in Fig. 23	Final rating
BF1	21→19	11ms	9M	0/0	698
BF2	23→19	59ms	457M	248489/8110	78876

We compute the polling frequency as the total number of polling events divided by the total elapsed time. As shown in Figure 20, the polling frequency increases from 1.01K polls/s to 5.07K polls/s, with a maximum of 5.27K polls/s. The polling frequency does not grow linearly with network size, because communication paths in collective operations have limited overlap and we configure each switch to respond to at most one polling request within any 500 μs interval. The notification overhead is determined by the number of collective communication steps. For 144-node Ring AllGather, the total notification overhead is approximately 1005 KB. These results show that, at the largest scale we can practically simulate, the storage, polling, and notification overheads of VEDRFOLNIR remain manageable.

E. Case Study

We demonstrate the effectiveness of VEDRFOLNIR through comprehensive root cause diagnosis of the anomalies mentioned in §II-B.

Flow contention. We diagnose the flow contention anomaly in Figure 3a. In a k= 4 fat-tree topology, an 8-node (N12–N19) Ring algorithm collective communication is executed. The collective communication consists of 8 flows (F12–F19), each step with a size of 360M. After the collective communication starts, two interference flows are injected sequentially, with their timing and sizes shown in Table II, and their paths illustrated in Figure 3a. Our diagnosis results include the following three aspects: (1) Waiting graph. Figure 21 shows a schematic of the waiting graph. The graph visually displays the execution time of each flow step and computes the critical path. In this example, both F17 and F18 are significantly affected, with F17 contributing the most to the collective communication delay. (2) Network provenance. Based on the waiting graph, we further examined the network provenance results for the preceding 200 ms. Figure 22 shows the network provenance graph generated in our experiment for the 10–20 ms interval. It can be seen that BF1 (0b001501 → 0b001301) conflicts with F17 (0b001101 → 0b001201) at SW6P2, and the congestion is propagated to SW2P3 via PFC, further affecting F17. Meanwhile, BF1 conflicts with F18 (0b001201 → 0b001301) and F15 at different locations. Figure 23 shows the graph for the 130–140 ms interval, which corresponds to the execution time of F17S3, the slowest flow step in the waiting graph. (3) Contribution evaluation. Following the method described in §III-D, we score the contributions of BF1 and BF2. Table II presents the scores of the two flows for F17/F18 in Figure 23, as well as the final comprehensive scores. From Figure 21, it is evident that F17 is more severely affected (70 ms) than F18 (50 ms), which is reflected in the edge weights and scores in Figure 23. Since BF1 has already finished execution by 130–140 ms and has no impact, its score is 0. Finally, the comprehensive score is calculated according to Formula 3, identifying BF2 as the primary contributor to the collective communication delay.

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

PFC Storm. We diagnose the PFC storm anomaly in Figure 3b. The topology and settings are the same as in Figure 3b. PFC is continuously injected at Switch 0. Our diagnosis results are shown in Figure 24. Through the blue edges representing PFC wait-relations, the network provenance graph reveals the root cause location of the PFC storm, its propagation path, and the affected victim flow (F19, 0b001301 → 0b000c01).

V. Discussion

Non-network root cause diagnosis. Beyond network performance anomalies, collective communication anomalies may stem from other sources such as slow computation. In collective communication algorithms, compute operations such as reduce are interleaved between receive and send operations. We focus primarily on network diagnosis and, therefore, omit the representation of compute operations within the waiting graph. However, the waiting graph formalism is inherently applicable to compute operations and supports step-by-step performance diagnosis. Currently, VEDRFOLNIR provides a novel perspective for diagnosing network performance anomalies and can seamlessly integrate with host-based diagnostic systems to achieve a comprehensive root cause diagnosis. We leave full integration of these two methodologies as future work.

Extensibility to anomaly types. Our current evaluations cover four anomaly types, yet VEDRFOLNIR enables straightforward incorporation of additional types based on their distinctive signatures. Although VEDRFOLNIR's adaptive mechanism may be ineffective for anomalies such as PFC deadlocks (which stall multiple collective flows), its flexibility allows simple fixes, e.g., immediately triggering an investigation when a flow remains stalled for an extended period.

Extensibility to asymmetric operations. For asymmetric collectives such as Broadcast, Reduce, Scatter, and Gather, VEDRFOLNIR does not need to directly apply the waiting graph mechanism. Unlike collective communication algorithms with intricate dependencies, these asymmetric operations lack complex synchronization logic between flows. VEDRFOLNIR treats these operations as standard peer-to-peer communications for diagnostic purposes, while its polling and network provenance mechanisms remain fully applicable.

Extensibility of algorithm decomposition. VEDRFOLNIR primarily characterizes execution flows through an algorithm's wait-dependency relationships. Given that synchronization is fundamental to most collective communications, VEDRFOLNIR applies broadly across nearly all collective algorithms.

Applicability in large-scale scenarios. The computational and storage complexity of the waiting graph is $O(N_n S)$, where N_n denotes the number of nodes in the set communication and S represents the number of steps. The computational and storage complexity of the tracing graph is $O(N_s T)$, where N_s denotes the number of relevant switches and T represents the number of switch reports. In large-scale scenarios, the centralized analyzer may become a bottleneck, requiring further design to address storage demands.

VI. Related Work

RDMA performance anomalies. Hostping [37] detects performance anomalies in different PCIe and NVLink pathways within the machine by performing internal loopback tests. R-Pingmesh [38], through full mesh connectivity testing, diagnoses performance anomalies between machines. However, most performance anomalies are derived from internal machine issues or misconfigurations, and they do not address the impact of complex traffic patterns in the network.

Model training performance anomalies. Holmes [39], Aegis [40], and Mycroft [41] detect anomalous nodes in the training cluster by analyzing the patterns of collective communication calls. The goal of this detection is to identify problematic nodes or links, considering both computational and communication performance. However, there is no in-depth analysis of the root causes of network performance anomalies.

Root cause diagnosis. Both Hawkeye [17] and SpiderMon [18] can perform root cause localization based on different anomaly types. However, when evaluating contributions, SpiderMon treats all detected flows equally and cannot evaluate the contribution of individual flows. Hawkeye traces root causes via PFC causality but fails to quantify the impact of contributions under collective communications.

VII. Conclusion

The co-flow pattern of collective communications introduces new complexities for RDMA NPA diagnosis. We present VEDRFOLNIR, an accurate and efficient anomaly diagnosis system for RDMA NPAs in collective communications. It employs algorithm decomposition to characterize collective communication processes, step-aware adaptive detection to minimize redundant overhead, and comprehensive diagnosis to identify the root causes. Evaluation demonstrates that VEDRFOLNIR accurately generates diagnostic results with low overhead.

References

[1] *InfiniBand Architecture Specification Release 1.4 Annex A17: RoCEv2*, InfiniBand Trade Association, Beaverton, OR, USA, 2020.
[2] Y. Gao et al., "When cloud storage meets RDMA," in *Proc. NSDI*, 2021, pp. 519–533.
[3] C. Guo et al., "RDMA over commodity Ethernet at scale," in *Proc. ACM SIGCOMM Conf.*, Aug. 2016, pp. 202–215.
[4] W. Bai, S. S. Abdeen, A. Agrawal, K. K. Attre, and P. Bahl, "Empowering Azure storage with RDMA," in *Proc. NSDI*, pp. 49–67.
[5] A. Singhvi et al., "1RMA: Re-envisioning remote memory access for multi-tenant datacenters," in *Proc. Annu. Conf. ACM Special Interest Group Data Commun. Appl.*, Technol., Archit., Protocols Comput. Commun., Jul. 2020, pp. 708–721.
[6] A. Gangidi et al., "RDMA over Ethernet for distributed training at meta scale," in *Proc. ACM SIGCOMM Conf.*, Aug. 2024, pp. 57–70.
[7] J. Cao et al., "Crux: GPU-efficient communication scheduling for deep learning training," in *Proc. ACM SIGCOMM Conf.*, Aug. 2024, pp. 1–15.
[8] Y. Jiang, Y. Zhu, C. Lan, B. Yi, Y. Cui, and C. Guo, "A unified architecture for accelerating distributed DNN training in heterogeneous GPU/CPU clusters," in *Proc. 14th USENIX Symp. Operating Syst. Design Implement. (OSDI)*, 2020, pp. 463–479.
[9] Z. Jiang et al., "MegaScale: Scaling large language model training to more than 10,000 GPUs," in *Proc. NSDI*, 2024, pp. 745–760.
[10] W. Wang et al., "TopoOpt: Co-optimizing network topology and parallelization strategy for distributed training jobs," in *Proc. NSDI*, 2022, pp. 739–767.

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.

[11] J. Xue, Y. Miao, C. Chen, M. Wu, L. Zhang, and L. Zhou, "Fast distributed deep learning over RDMA," in *Proc. 14th EuroSys Conf.*, Mar. 2019, pp. 1–14.
[12] K. Qian et al., "Alibaba HPN: A data center network for large language model training," in *Proc. ACM SIGCOMM Conf.*, 2024, pp. 691–706.
[13] (2011). IEEE 802.1 Qbb - Priority-Based Flow Control. [Online]. Available: https://1.ieee802.org/dcb/802-1qbb/
[14] Y. Zhu et al., "Congestion control for large-scale RDMA deployments," in *Proc. ACM Conf. Special Interest Group Data Commun.*, Aug. 2015, pp. 523–536.
[15] G. Kumar et al., "Swift: Delay is simple and effective for congestion control in the datacenter," in *Proc. ACM SIGCOMM Conf. Appl. Technol. Archit. Protocols Comput. Commun. (SIGCOMM)*, 2020, pp. 514–528.
[16] X. Li, S. Wang, M. Zhang, Z. Wang, M. Xu, and J. Yang, "POSTER: RDMA network performance anomalies diagnosis with hawkeye," in *Proc. ACM SIGCOMM Conf., Posters Demos*, Aug. 2024, pp. 1–3.
[17] S. Wang et al., "Hawkeye: Diagnosing RDMA network performance anomalies with PFC provenance," in *Proc. ACM SIGCOMM Conf.*, Sep. 2025, pp. 481–495.
[18] W. Wang, X. C. Wu, P. Tammana, A. Chen, and T. S. E. Ng, "Closed-loop network performance monitoring and diagnosis with SpiderMon," in *Proc. USENIX NSDI*, Apr. 2022, pp. 267–285.
[19] A. Gupta, R. Harrison, M. Canini, N. Feamster, J. Rexford, and W. Willinger, "Sonata: Query-driven streaming network telemetry," in *Proc. Conf. ACM Special Interest Group Data Commun.*, Aug. 2018, pp. 357–371.
[20] J. Sonchack, O. Michel, A. J. Aviv, E. Keller, and J. M. Smith, "Scaling hardware accelerated network monitoring to concurrent and dynamic queries with *flow," in *Proc. USENIX ATC*, 2018, pp. 823–835.
[21] R. Thakur, R. Rabenseifner, and W. Gropp, "Optimization of collective communication operations in MPICH," *Int. J. High Perform. Comput. Appl.*, vol. 19, no. 1, pp. 49–66, Feb. 2005.
[22] N. Handigol, B. Heller, V. Jeyakumar, David Mazières, and N. McKeown, "I know what your packet did last hop: Using packet histories to troubleshoot networks," in *Proc. USENIX NSDI*, 2014, pp. 71–85.
[23] R. B. Basat, S. Ramanathan, Y. Li, G. Antichi, M. Yu, and M. Mitzenmacher, "PINT: Probabilistic in-band network telemetry," in *Proc. Annu. Conf. ACM Special Interest Group Data Commun. Appl.*, Technol., Archit., Protocols Comput. Commun., Jul. 2020, pp. 662–680.
[24] Vedrfolnir. (2026). *Vedrfolnir: Diagnosing RDMA Network Performance Anomalies in Collective Communications*. [Online]. Available: https://github.com/Networked-System-and-Security-Group/Vedrfolnir/tree/TON26
[25] Y. Chen, M. Zhang, X. Li, F. Jiao, and C. Hu, "Vedrfolnir: RDMA network performance anomalies diagnosis in collective communications," in *Proc. ACM SIGCOMM Posters Demos*, Sep. 2025, pp. 10–12.
[26] C. H. Song, X. Z. Khooi, R. Joshi, I. Choi, J. Li, and M. C. Chan, "Network load balancing with in-network reordering support for RDMA," in *Proc. ACM SIGCOMM Conf.*, Sep. 2023, pp. 816–831.
[27] Y. Li, R. Miao, C. Kim, and M. Yu, "FlowRadar: A better NetFlow for data centers," in *Proc. USENIX NSDI*, 2016, pp. 311–324.
[28] P. Tammana, R. Agarwal, and M. Lee, "Distributed network monitoring and debugging with SwitchPointer," in *Proc. USENIX NSDI*, 2018, pp. 453–466.
[29] M. Ghasemi, T. Benson, and J. Rexford, "Dapper: Data plane performance diagnosis of TCP," in *Proc. Symp. SDN Res.*, Apr. 2017, pp. 61–74.
[30] C. Guo et al., "Pingmesh: A large-scale system for data center network latency measurement and analysis," in *Proc. ACM Conf. Special Interest Group Data Commun.*, Aug. 2015, pp. 139–152.
[31] M. Moshref, M. Yu, R. Govindan, and A. Vahdat, "Trumpet: Timely and precise triggers in data centers," in *Proc. ACM SIGCOMM Conf.*, Aug. 2016, pp. 129–143.
[32] Y. Lei, L. Yu, V. Liu, and M. Xu, "PrintQueue: Performance diagnosis via queue measurement in the data plane," in *Proc. ACM SIGCOMM Conf.*, Aug. 2022, pp. 516–529.
[33] B. Arzani et al., "007: Democratically finding the cause of packet drops," in *Proc. USENIX NSDI*, 2018, pp. 419–435.
[34] X. Chen et al., "Fine-grained queue measurement in the data plane," in *Proc. 15th Int. Conf. Emerg. Netw. Experiments Technol.*, Dec. 2019, pp. 15–29.
[35] Y. Wu, A. Chen, and L. T. X. Phan, "Zeno: Diagnosing performance prob- lems with temporal provenance," in *Proc. USENIX NSDI*, 2019, pp. 395–420.
[36] H. Liao et al., "UB-mesh: A hierarchically localized nD-fullmesh datacenter network architecture," 2025, *arXiv:2503.20377*.
[37] K. Liu et al., "Hostping: Diagnosing intra-host network bottlenecks in RDMA servers," in *Proc. USENIX NSDI*, 2023, pp. 15–29.
[38] K. Liu et al., "R-pingmesh: A service-aware RoCE network monitoring and diagnostic system," in *Proc. ACM SIGCOMM Conf.*, Aug. 2024, pp. 554–567.
[39] Z. Yao, M. C. Hu, Pengbo, and X. Jia "Holmes: Localizing irregularities in LLM training with mega-scale GPU clusters," in *Proc. USENIX NSDI*, 2025, pp. 523–540.
[40] J. Dong, K. Qian, and P. Zhang, "Evolution of aegis: Fault diagnosis for AI model training service in production," in *Proc. USENIX NSDI*, 2025, pp. 865–881.
[41] Y. Deng et al., "Mycroft: Tracing dependencies in collective communication towards reliable LLM training," in *Proc. ACM SIGOPS 31st Symp. Operating Syst. Princ.*, Oct. 2025, pp. 254–269.

Yuxuan Chen (Graduate Student Member, IEEE) is currently pursuing the B.S. degree with the School of Software, Beihang University. His research interests include RDMA networks, collective communication, and LLM systems.

Menghao Zhang (Member, IEEE) received the B.S. and Ph.D. degrees in computer science from Tsinghua University, China. He is currently an Associate Professor with the School of Software, Beihang University. His research interests include programmable networks, high-performance networks, networked systems, and network security.

Jingyuan Zhu is currently pursuing the B.S. degree in software engineering with the School of Software, Beihang University. His research interests include collective communication and RDMA networks.

Yuan Yang received the B.Sc., M.Sc., and Ph.D. degrees from Tsinghua University. He is currently an Associate Professor with the Department of Computer Science and Technology, Tsinghua University. His research interests include computer network architecture, routing protocols, and algorithms.

Xiao Li received the B.S. degree in computer science from Tsinghua University, China, where he is currently pursuing the Ph.D. degree. His research interests include programmable networks and high-performance networks.

Fangzheng Jiao (Graduate Student Member, IEEE) is currently pursuing the Ph.D. degree with the School of Software, Beihang University. His research interests include large-scale model infrastructure and stability of distributed systems.

Shicheng Wang received the B.S. and Ph.D. degrees in computer science from Tsinghua University. His research interests include high-performance networks, programmable networks, and network systems.

Chunming Hu received the Ph.D. degree in computer science from Beihang University, Beijing, China, in 2006. He is currently a Professor with the School of Software, Beihang University, Beijing. His research interests include distributed systems, system virtualization, data management, and processing systems.

Mingwei Xu received the B.S. and Ph.D. degrees from Tsinghua University. He is currently a Full Professor with the Department of Computer Science and Technology, Tsinghua University. His research interests include computer network architecture, high-speed router architecture, and network security.

Authorized licensed use limited to: BEIHANG UNIVERSITY. Downloaded on November 06,2026 at 00:05:16 UTC from IEEE Xplore. Restrictions apply.