



Study of Workload Interference with Intelligent Routing on Dragonfly

Yao Kang

Department of Computer Science
Illinois Institute of Technology
Chicago, USA
ykang17@hawk.iit.edu

Xin Wang

Department of Computer Science
Illinois Institute of Technology
Chicago, USA
xwang149@hawk.iit.edu

Zhiling Lan

Department of Computer Science
Illinois Institute of Technology
Chicago, USA
lan@iit.edu

Abstract—Dragonfly interconnect is a crucial network technology for supercomputers. To support exascale systems, network resources are shared such that links and routers are not dedicated to any node pair. While link utilization is increased, workload performance is often offset by network contention. Recently, intelligent routing built on reinforcement learning demonstrates higher network throughput with lower packet latency. However, its effectiveness in reducing workload interference is unknown. In this work, we present extensive network simulations to study multi-workload contention under different routing mechanisms, intelligent routing and adaptive routing, on a large-scale Dragonfly system. We develop an enhanced network simulation toolkit, along with a suite of workloads with distinctive communication patterns. We also present two metrics to characterize application communication intensity. Our analysis focuses on examining how different workloads interfere with each other under different routing mechanisms by inspecting both application-level and network-level metrics. Several key insights are made from the analysis.

Index Terms—HPC, interconnect network, Dragonfly, network interference

I. INTRODUCTION

High-performance computing (HPC) systems rely on efficient and scalable interconnect networks to support unprecedented system size at a reasonable cost. Recent high-radix, low-diameter Dragonfly topology has demonstrated its high capability on the current Top500 list [1] [2]. The next generation Slingshot interconnects also adopt Dragonfly topology to support future exascale HPC systems [3].

Dragonfly topology achieves extreme performance by arranging network resources of routers and links into fully connected groups. As a result, Dragonfly is a *hierarchical topology* that uses global links for inter-group connection and local links to connect routers in the same group. The hierarchical design makes Dragonfly a low-diameter system such that any point in the network can be accessed by crossing the interconnect hierarchy in three hops. The fully connected inter-group structure also makes Dragonfly a *high path diversity* network with abundant routing path possibilities. A packet can either be minimally forwarded from its source group to the destination group or be relayed at an intermediate group in the system. Larger systems with more groups further increase the degree of path diversity. While high path diversity

provides more flexibility in choosing the packet forwarding path, it poses a significant challenge in routing design.

Due to the sharing nature of Dragonfly, traffic flows from different applications have a high probability of competing for a portion or the entire network forwarding path, thus leading to *workload interference*. Adaptive routing is widely deployed on production Dragonfly systems to deal with high path diversity and network congestion [3]–[5]. When forwarding a packet, adaptive routing [1], [6] chooses a path between either a minimal path or a non-minimal path via an intermediate group according to local information such as router port queue occupancy. Although heuristic based adaptive routing can mitigate the workload interference in some degree, several empirical studies conducted on production Dragonfly systems have indicated that HPC applications can see a slowdown as much as 70% [7], [8]. Communication interference among different workloads can lead to many negative impacts, e.g., longer application execution time and wasted resource cycles, all of which lead to low system productivity.

In order to address workload interference, various approaches have been studied. Several efforts have demonstrated that careful job placement such as *contiguous placement* can help mitigate workload interference [9], [10]. Contiguous placement helps isolate an application from others by placing application processes in the same group. While contiguous placement is shown to be effective for reducing workload interference through network simulation studies, it may cause local hot spots [11]. In addition, contiguous placement is impractical in practice as it can cause severe system fragmentation. For example, external fragmentation occurs when there is a sufficient number of compute nodes available for a job; however, they cannot be allocated to the job because these compute nodes are not in a contiguous partition. Another research approach is *application-aware routing* [12]. The main idea is to build an analytical model to estimate network congestion and use it to adjust adaptive routing bias at the application’s runtime. Continuous network monitoring is needed for performance modeling, which causes undesirable overhead to application performance. Besides, the model may not be accurate which further impacts the effectiveness of this approach. Other more sophisticated methods such as *congestion control*, with support of additional hardware and/or software, are also proposed to

be used in combination with adaptive routing algorithms [3], [13].

In this study, we investigate workload interference by presenting *a quantitative analysis of intelligent routing for mitigating workload interference on Dragonfly systems*. Q-adaptive routing is an intelligent routing method built on reinforcement learning for Dragonfly [14]. Preliminary results indicate that Q-adaptive routing provides superior performance than adaptive routing under either well balanced or extremely imbalanced network traffics. However, its effectiveness for reducing workload interference is an open problem.

Quantitative analysis of workload interference under routing on Dragonfly is challenging. Dragonfly is a proprietary network technology, and is targeted for large-scale supercomputers such as those listed on the Top500 list. Although dozens of Dragonfly systems are available at national supercomputer centers, routing configuration is part of system configuration, which is impossible for general users to make changes at will [10], [11], [15]. Following the common practice in the literature [16], [17], we use high-fidelity flit-level network simulation in this study.

Our study makes *four contributions*. **First**, we enhance the well-known Structural Simulation Toolkit (SST) toolkit for the quantitative analysis (Section III). Several SST core and element modules are extended to support workload interference analysis at the flit level. The extended SST toolkit is available as open source software on GitHub [18]. **Second**, in order to represent realistic HPC workloads, nine applications with distinct communication patterns are selected and developed for the study. These applications include both conventional HPC workloads and emerging machine learning workloads. They are also available on GitHub, along with the enhanced SST toolkit. Moreover, we present two metrics to formally characterize communication intensity of HPC applications (Section IV). We use the augmented SST toolkit to systematically analyze workload interference through pairwise workload and mixed workload analysis on a 1,056-node Dragonfly system. **Third**, even with high-fidelity network simulation, how to effectively study workflow interference with different routing mechanisms is quite challenging. In particular, an overwhelming amount of application-level and network-level data series can be collected. In this study, we propose two sets of quantitative metrics for interference analysis: (1) *application-level metrics* such as application communication throughput, average application packet latency, and application communication time; and (2) *network-level metrics* such as router port stall time, link traffic size and congestion index (Section V-VI). These metrics are used in various simulation configurations to examine when and how interference occurs. **Fourth**, several key insights are obtained through the analysis, some being listed below:

- Application communication intensity can be well captured by two metrics, that it, message injection rate and peak communication ingress volume. An application with a lower message injection rate tends to be interfered by those with a higher rate, and an application with

high peak ingress volume can cause severe workload interference to other applications.

- Compared to adaptive routing, Q-adaptive routing can greatly reduce workload interference by saving up to 42.63% communication time and mitigating up to 70.80% communication performance variation caused by network contention. System-wide network analysis shows that Q-adaptive routing optimizes overall network utilization by balancing system traffic, eliminating hot spots and hence avoiding network congestion.
- For collective communication operations, tail latency control becomes important since the operation completes only after all have received the messages. Moreover, workload interference can be hidden by the application with long computation time and short communication time.

II. BACKGROUND AND RELATED WORK

A. Dragonfly Topology

Dragonfly is a *hierarchical topology* by arranging network resources of routers and links into high-radix groups [1]. An example is shown in Figure 1. The groups are fully connected by global links and routers in the same group are connected through local links. Although groups are always all-to-all connected to assure system connectivity, the intra-group connection structure is not fixed and can have various forms. For example, routers within the same group can also be all-to-all connected for higher system throughput or be arranged into a two-dimensional array or a tree to support a larger system size [5], [19], [20]. In this study, we focus on the Dragonfly with a fully connected intra-group connection since this structure will be deployed on the next-generation HPC systems [3].




Fig. 1. The hierarchical structure of 1,056-node Dragonfly system.

With a fully connected inter-group and intra-group structure, Dragonfly is a low-diameter topology. Any packet can be delivered within three hops by crossing: (1) one local link in the source group to reach the router that has global link connection to the destination group, (2) one global link crossing groups, (3) and one local link in the destination group to the destination router. Since two local links are required with one global link, local and global links are typically provisioned in a ratio of 2:1 on each router to make sure global link bandwidth is not wasted due to the insufficient local links [1].

B. Routing on Dragonfly

Unlike fat-tree topologies, the hierarchical design and fully shared network resources make Dragonfly routing unique. Static routing such as minimal routing that performs well on other topologies is not ideal for Dragonfly due to the limited number of global links between source and destination groups. Hence, Dragonfly relies on adaptive routing to let routers sense local congestion and bypass hot spots by forwarding packets non-minimally through an intermediate group if necessary. **Adaptive routing** is a class of heuristic methods that use router's port queue occupancy to estimate network congestion. When forwarding a packet, a router first randomly selects two minimal paths and two non-minimal paths and checks their corresponding port queue occupancy [4]. When the best minimal path queue occupancy is less than twice of the best non-minimal path queue occupancy, the packet is minimally forwarded, otherwise, the non-minimal path is used.

Depending on where the dynamic routing decision is made, adaptive routing algorithms have three widely used solutions: Universal Globally-Adaptive Load-balanced routing (UGAL) lets the source router make a one-time routing decision with two variants. In case a packet enters an intermediate group, **UGALg** (UGAL_group) forwards it minimally towards the destination group, whereas **UGALn** (UGAL_node) tries to avoid local link congestion by having the packet firstly visit a random router in that intermediate group [21]. **PAR** (Progressive Adaptive Routing) is similar to UGALn, except that it allows routers in the source group to revise the previously made minimal routing decision in case local congestion is encountered at downstream routers [6]. Although adaptive routing is good at avoiding near-end congestion, they are oblivious to the network condition at a few hops away hence may cause network congestion.

Q-adaptive Routing is a recent Dragonfly routing solution leveraging reinforcement learning (RL) techniques [14]. In Q-adaptive routing, each router learns the overall network condition, records the knowledge in a light-weight two-level Q-table, and uses the table for packet forwarding. Through the process of sending packets and receiving feedback signals from the neighbors, every router learns the system-wide network condition and stores the information in the Q-table. Similar to adaptive routing algorithms, Q-adaptive routing forwards packets either minimally or non-minimally. In other words, unlike adaptive routing which makes routing decisions based on local information, Q-adaptive routing uses *the learned overall system condition and global path congestion* to direct packet delivery (Figure 2). In addition, Q-adaptive routing is fully distributed such that routers make their independent decisions without requiring any shared information.

In this study, we investigate workload interference under the aforementioned routing mechanisms on Dragonfly systems.

C. Related Works

Chunduri et al. studied Dragonfly network interference on production systems and observed as much as 7x performance variability for MPI collectives and 70% variability for MILC

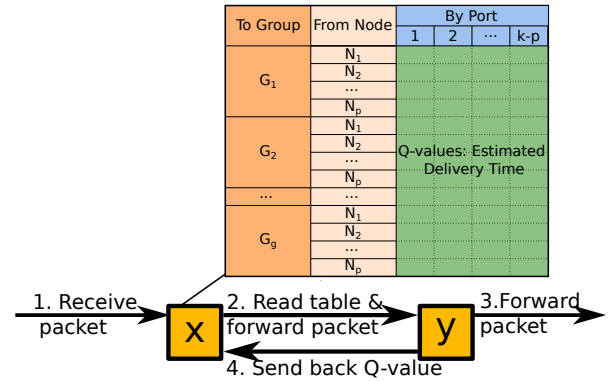


Fig. 2. Q-adaptive routing and its two-level Q-table per router. Router X follows four steps to forward a packet and update its table. Through the process of sending packets and receiving feedback signals from the neighbors, every router learns the system-wide network condition and records the knowledge in its two-level Q-table for packet forwarding.

application due to the network contention [7]. Simulation studies show that network interference also appears on Dragonfly variants under different intra-group connection structures [9], [22]. Yang et al. observed the bully effect on the Dragonfly system that communication intensive application affects less intensive applications and purposed using contiguous placement to reduce network interference [10]. Jain et al. proposed to use random placement to improve system throughput by reducing network hot-spots on Dragonfly systems [11]. In addition to job placement, QoS is another approach to mitigate network interference by separating traffic flows of different applications or communication types into isolated channels [13], [23], [24]. At the routing level, Smith et al proposed adaptive flow-aware routing to mitigate network interference [25]. De Sensi et al proposed to use dynamic adaptive routing bias to constrain the contention effect [12]. A more complicated congestion control mechanism is also proposed such that when congestion happens, the message generation rate is throttled to drain the network system [3], [26].

This work presents a quantitative analysis of different routing methods with respect to workload interference mitigation, hence complementing the above related interference studies.

III. NETWORK SIMULATION DESIGN

Workload interference can be analyzed through application tracing and profiling. However, the data collected in the trace is limited to the given application, and the information of other co-running applications is not included. Although analyzing application trace or profile can unveil the interference effect of the studied application, it fails to answer how and when workload interference occurs.

High-fidelity network simulation is a viable approach to study workload interference [10], [27]. In this study, we leverage the well-known Structural Simulation Toolkit (SST) for interference analysis [28]. The SST toolkit contains a number of core and element libraries. In this work, we enhance three SST libraries, namely Ember, Firefly, and Merlin, for the quantitative network study. *SST/Ember* contains a collection of

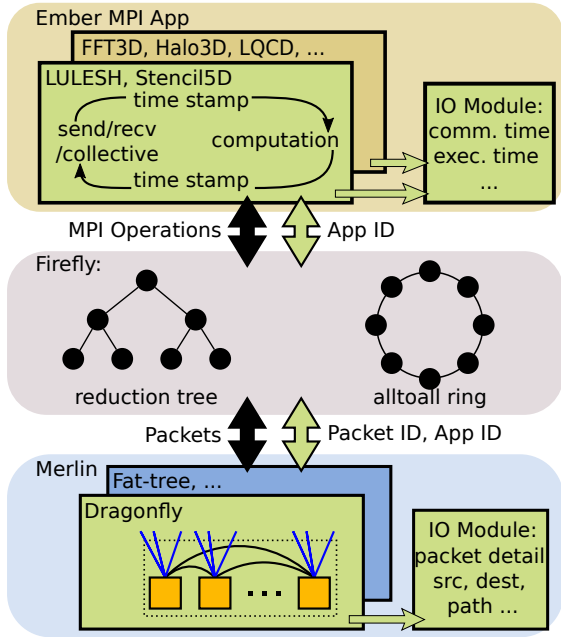


Fig. 3. Enhancing SST for workload interference analysis. Our enhancements are shaded in green.

MPI motifs that represent real-world application communication behaviors. The communication generated by SST/Ember is passed to *SST/Firefly*, a state machine based MPI implementation layer that simulates blocking, non-blocking communication, and MPI collective operations with eager and rendezvous protocols. Finally, the Firefly MPI messages are packetized and transmitted to *SST/Merlin*, which provides high-fidelity, flit-level network simulation including Dragonfly topology.

Figure 3 depicts the interplay of these SST libraries, along with our enhancements. Specifically, we implement an IO module that can record any desired performance counter at any frequency. For the purpose of simulation efficiency, the IO module can be flexibly configured to coalesce multiple write operations into one action to balance the trade-off between IO efficiency and system memory usage. The designed IO module enables us to investigate every detail of a simulated system at both application and network level. At the application level, all the studied MPI applications are updated to accurately timestamp each iteration’s starting, ending time, and time spent on different messaging operations. The applications are connected with the IO module to record all the metrics for post-simulation analysis. We also enhance SST/Merlin Dragonfly topology by connecting it with the IO module to record every packet’s detailed information including source, destination, sending, receiving time, and forwarding path. Moreover, we enhance the bridge between SST/Ember and SST/Merlin by passing application’s job ID to provide fine-grained per-application network link usage statistics.

In this study, we focus on a 33-group, 1,056-node Dragonfly system as shown in Figure 1. Each group contains eight fully connected routers with a total of 32 global links. Each router

TABLE I
LIST OF APPLICATIONS WITH THEIR COMMUNICATION PATTERNS.

| Pattern | App | Total Msg (MB) | Execution time (ms) | Injection Rate (GB/s) | Peak Ingress Volume |
|----------------|-----------|----------------|---------------------|-----------------------|---------------------|
| Random | UR | 11829.48 | 13.31 | 888.48 | 3.07KB |
| Sweep | LU | 13713.22 | 13.71 | 999.88 | 30.0KB |
| Alltoall | FFT3D | 15781.09 | 12.53 | 1259.35 | 51.68KB |
| Stencil | Halo3D | 47769.10 | 10.85 | 4403.81 | 1.15MB |
| | LQCD | 11924.31 | 13.79 | 864.70 | 4.60MB |
| | Stencil5D | 9833.95 | 13.70 | 717.87 | 14.0MB |
| Allreduce | CosmoFlow | 2373.84 | 13.65 | 173.86 | 2.25MB |
| | DL | 9714.44 | 11.86 | 819.12 | 2.30MB |
| Stencil +Sweep | LULESH | 17900.12 | 12.34 | 1450.78 | 1.95MB |
| | | | | | 14.91KB |

hosts four compute nodes using 128B flits and 512B packets. To ensure network bandwidth is not limited by the flying credits of credit-based flow control, each router port can buffer up to 30 packets. Local and global links are configured with 200Gb/s bandwidth according to the Slingshot system [3]. Considering the difference in link length, flit transmission latency is 30ns and 300ns for local and global links to keep the 1:10 ratio as used in previous works [1], [6], [21]. The evaluated adaptive routing algorithms are configured with zero bias towards the minimal path and Q-adaptive routing with the same hyperparameters as in [14].

IV. APPLICATIONS AND COMMUNICATION PATTERNS

Although existing studies show that communication-intensive applications can bully other applications on Dragonfly [10], *there is no formal definition of communication intensity*. We propose two metrics to characterize application communication patterns. Both metrics can be measured or be derived from the application source code.

- *Message injection rate* is defined as an average value derived from an application’s total message size and its execution time. It is measured as the average bandwidth requirement of an application assuming its packets are steadily injected into the network.
- *Peak ingress volume* is defined as an application’s peak bandwidth requirement, meaning the amount of message that the network is expected to process in a short time. It is measured as the consecutive message size injected into the network by the application.

In order to well capture representative workloads on production systems, we carefully select nine applications based on their communication patterns. They cover a broad range of workloads, which are grouped into five representative communication patterns commonly observed in the traditional scientific and the emerging machine learning applications. Some applications are from the SST package, and some are

developed in this work. Table I summarizes these applications, along with their communication patterns.

- **Random:** Random is a typical one-to-one communication pattern in the scientific field to study network performance. The Uniform Random (*UR*) is mainly used as a background application to mimic a system under a balanced network load with each process sending messages to random targets.
- **Sweep:** The Lower-Upper (*LU*) Gauss-Seidel solver application from the NAS Parallel Benchmark suite is featured by a sweep communication pattern [29]. LU arranges its processes as a 2D square and initiates the network communication from one square corner. All processes send data to their downstream neighbors after receiving messages from their upstream peers along both square dimensions. As a result, the communication behaves like a wavefront sweeping from one corner of the square to the other diagonal corner stage by stage, which causes long communication latency. Because each process has two downstream partners, LU’s peak ingress volume counts two messages.
- **Alltoall:** This MPI collective is important for applications that perform parallel Fast Fourier Transforms (FFT) such as pF3D [30], NAMD [31], and VASP [32]. SST implements MPI_Alltoall through a multi-step ring exchange such that process N receives data from process $N - i$ and sends data to process $N + i$ in round i . Since each process only sends one message in each round, Alltoall operation’s peak ingress volume only counts one message. When the communicator is large, many rounds of send and receive are required hence leading to long communication latency.
The studied *FFT3D* application maps the problem to a two-dimensional processes array, where processes on the same row and column perform Alltoall operations.
- **Stencil:** Stencil computation is an important class of algorithms in scientific computing and a representative one-to-many communication pattern [33]. Under a stencil pattern, processes are arranged into a multi-dimensional grid, and each process talks with its nearest neighbors in both directions along all the dimensions. As a result, stencil’s peak ingress volume must consider both the message size and process grid dimension. Because at each communication step, processes send data to multiple destinations, stencils are high-throughput and aggressive on occupying network resources.
Halo3D is a 3D-stencil application that processes communicate with its six neighbors along three dimensions. Lattice Quantum ChromoDynamics (*LQCD*) is a QCD simulation application to study strong force theory featured with a 4D-stencil pattern of eight neighbors [34]. *Stencil5D* is a synthetic application used to study the impact of peak ingress volume that processes have up to ten neighbors.
- **Allreduce:** Distributed machine learning applications pe-

riodically aggregate and distribute the learned information among all processes to update model parameters [35]. As a result, collective communication such as Allreduce becomes heavily used and is presented as an important message movement pattern on the HPC systems. In SST, MPI_Allreduce is implemented by arranging processes into a binary tree where the desired data is firstly aggregated from leaf to root and then distributed in the opposite direction. Allreduce peak ingress volume counts two messages considering each tree node has two children.

CosmoFlow is a fully synchronous data parallel distributed deep learning (DL) cosmology application [36]. CosmoFlow shows strong scaling performance with the main communication pattern of 28.15 MB Allreduce messages every 129 ms. To make the application have similar execution time as other applications, we proportionally decrease both message size and communication interval by 25x and keep its intrinsic communication intensity. To represent large-scale distributed DL applications with a massive training dataset, a heavier Allreduce application named *DL* is studied. DL has similar message size but shorter communication interval, such that its message injection rate is around 4.7x higher than that of CosmoFlow.

- **Hybrid:** The Livermore Unstructured Lagrangian Explicit Shock Hydrodynamics (*LULESH*) is a widely studied proxy application designed for exascale hardware/software co-design effort [37]. LULESH represents a typical hydrocode such as ALE3D using MPI non-blocking (MPI_Isend, MPI_Irecv) and collective operations [38]. LULESH communication pattern is dominated by a 26-point 3D-stencil followed by a sweep3D data exchange [39], [40]. LULESH shows good weak scaling performance on production systems [41]. We implement LULESH based on the communication analysis listed in the literature [39], [40].

V. PAIRWISE WORKLOAD ANALYSIS

In this section, we study pairwise workload interference by co-running *a target application* with *a background application*. The goal is to analyze how the communication performance of the target application is impacted by the background application with different message injection rates and peak ingress volumes. Specifically, FFT3D, LU, LQCD, CosmoFlow, Stencil5D and LULESH are selected as the target applications covering different communication patterns and intensities. Each target application is co-run with a background application from different communication pattern categories. For pairwise analysis, the 1056-node system is equally divided to host a target application and a background application. To maintain a perfect process cube, LULESH takes 512 ($=8^3$) nodes leaving 16 nodes being idle. Random job placement is used in our experiments.

For each target application, its process-to-node mapping is kept unchanged across different runs (i.e., standalone or co-running with a background application). Therefore, a variation in communication time of the same application under different

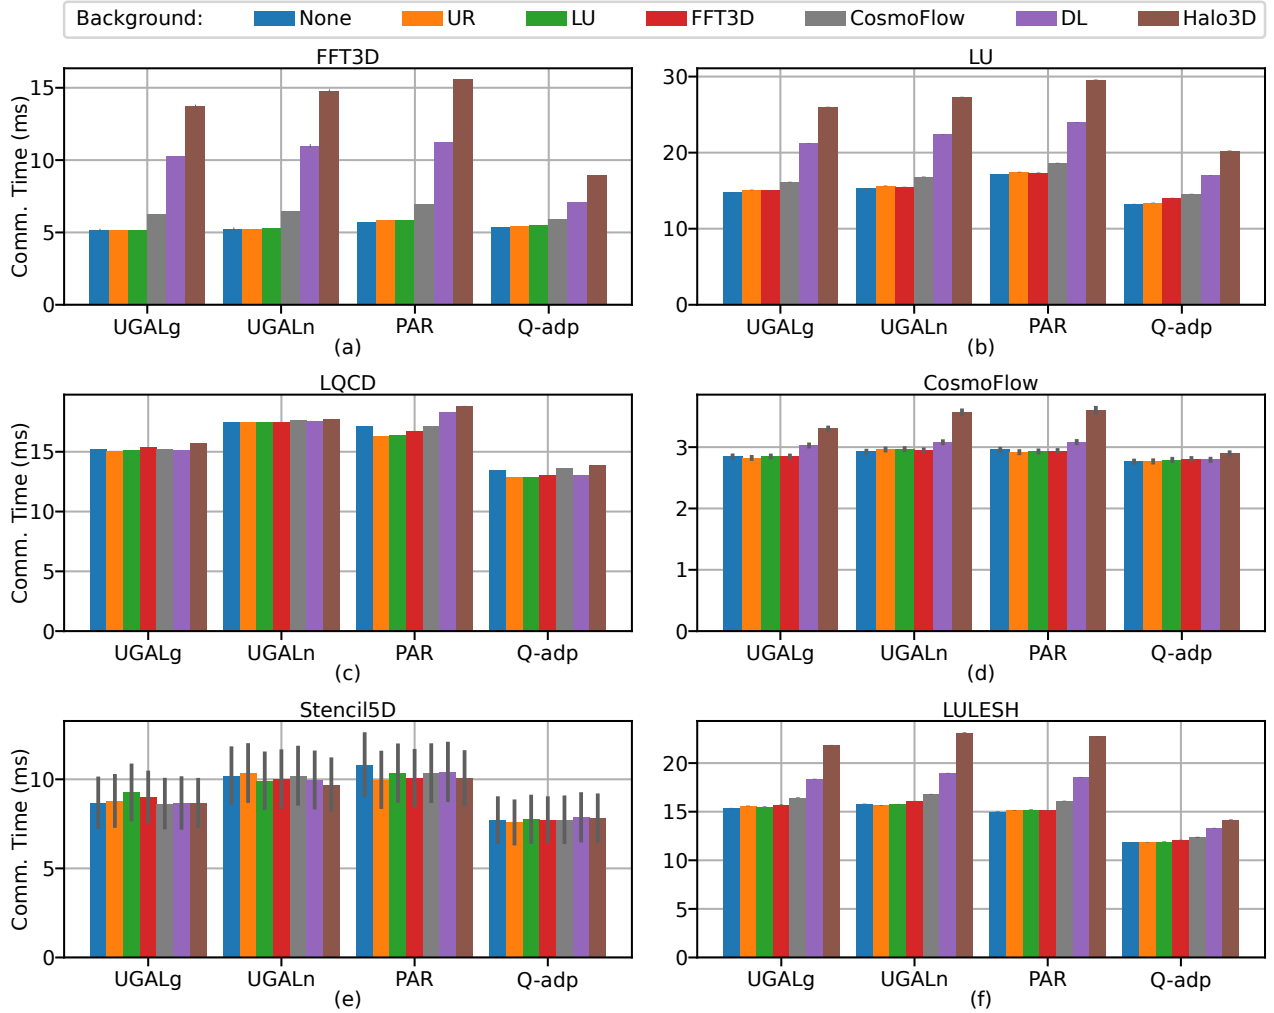


Fig. 4. Average communication time (bar) and standard deviation (vertical line) of a target application over all processes. The results for six target applications are presented in (a)-(f). For a target application, each colored bar indicates its communication time under a background application (or none).

runs unveils the effect of workload interference. *To make a fair comparison, Q-adaptive starts an application under the same condition as adaptive routing algorithms without any pre-trained information.* Thus, the time spent on training Q-adaptive routing is included in its communication time.

Note that we use both *application-level metrics* and *network-level metrics* for interference analysis. We mainly focus on the communication time and its variation across application processes for the application-level analysis. We leverage packet latency distribution, tail latency dispersion, and application communication throughput for the network-level analysis.

In order to clearly identify workload interference, we first conduct standalone experiments to collect target application performance in an interference-free environment. Figure 4 presents the application communication time under different pairwise analyses. The blue bars show the communication time of applications when it is the only job on the system. For the standalone cases, FFT3D and CosmoFlow perform similarly

among different routing algorithms, whereas LU, LQCD, Stencil5D, and LULESH perform the best under Q-adaptive routing with an average of 23.46% smaller communication time compared with PAR.

In short, *compared with adaptive routing, Q-adaptive achieves equal or better performance when each target application runs alone with random placement.*

A. Impact of Message Injection Rate

Applications with a higher injection rate tend to delay those with a lower rate. Q-adaptive routing reduces workload interference with shorter communication time and less variation.

Figure 4(a) shows that UR or LU causes negligible interference effect on FFT3D. The average communication time of FFT3D in both cases increases by less than 3% with all the studied adaptive and Q-adaptive routing algorithms. As summarized in Table I, FFT3D is hardly affected by UR and




Fig. 5. Network throughput of FFT3D and Halo3D along simulated time. Q-adaptive protects FFT3D’s performance from Halo3D’s interference with 2.58x higher throughput compared with that of PAR shown in green lines.

LU, both having lower injection rates. CosmoFlow causes a mild interference impact on FFT3D with an average of 22.51% communication delay under adaptive routing and a smaller 10.81% delay under Q-adaptive routing. Although the injection rate of CosmoFlow is smaller than that of FFT3D, CosmoFlow has a considerably larger peak ingress volume. The effect of peak ingress volume is discussed in §V-C. Compared with CosmoFlow as background application, DL — a higher injection rate application — causes more severe interference impact on target applications. Interference from DL makes FFT3D spend 2.00x, 2.11x, 1.97x, and 1.32x more time on communication with UGALg, UGALn, PAR, and Q-adaptive routing, respectively. Among adaptive routing algorithms, PAR’s 11.24ms average communication time is slightly longer than UGALg’s 10.23ms and UGALn’s 10.98ms. However, PAR has a smaller communication time variation of 97.14% compared with UGALg’s 99.53% and UGALn’s 110.77% variation. Although network interference is inevitable, Q-adaptive has the shortest communication time with the most stable performance. Q-adaptive’s 7.08ms communication time is 37.05% smaller than that of PAR. Compared with the standalone case, Q-adaptive’s communication time variation is 32.34%, which is 70.80% smaller than UGALn’s 110.77% variation. The highest injection rate application of Halo3D causes the most interference. Halo3D delays FFT3D’s communication time by 2.67x, 2.83x, and 2.73x using UGALg, UGALn, and PAR respectively. However, Q-adaptive reduces the interference effect by completing the communication within 8.93ms, which is 42.63% smaller than PAR’s 15.57ms.

To investigate the effect of injection rate, we examine the FFT3D-Halo3D co-running case. Figure 5 shows the application communication throughput along the simulated time under PAR and Q-adaptive routing. Both standalone and interfered cases are plotted in different colors. The standalone results in Figure 5(a) and (b) show that Halo3D (orange line) is a communication intensive application with continuous high throughput. In contrast, FFT3D (blue line) features burden network throughput on which the valleys correspond to the forward and backward FFT computation and the peaks for

the Alltoall operation between computations. When FFT3D and Halo3D are co-run together, Halo3D (red line), as the application with a higher injection rate, is hardly affected by FFT3D, with less than 2% variation with respect to both average throughput and application execution time under PAR and Q-adaptive routing. However, Halo3D maintains its high network throughput at the cost of degraded FFT3D performance (green line). In Figure 5(a), due to Halo3D’s aggressive network requirement, FFT3D’s average throughput (green line) using PAR is only 0.44GB/ms, which is 82.62% smaller than the 2.51GB/ms throughput of standalone case. Once Halo3D finishes at 14.66ms, FFT3D regains network bandwidth with an average of 2.46GB/ms throughput. In contrast, Q-adaptive is capable of protecting FFT3D’s performance while keeping Halo3D’s high throughput. In Figure 5(b), with Q-adaptive routing, the average throughput of FFT3D under Halo3D’s interference is 1.13GB/ms, which is 2.58x higher than PAR’s 0.44GB/ms. Once Halo3D finishes at 11.05ms, FFT3D regains network resources with an average of 2.50GB/ms throughput.

A similar effect is also observed for the LU shown in Figure 4(b). UR and FFT3D cause no more than 6.06% variation on average communication time for all the studied routing algorithms. CosmoFlow delays LU’s communication by 9.34% on average under adaptive routing and by 9.99% under Q-adaptive routing. In the Allreduce pattern category, the higher injection rate DL application causes more impact on LU. When co-running with DL, LU spends 43.49%, 46.63%, and 40.16% more time on communication under UGALg, UGALn and PAR, respectively. PAR has a more stable performance with a slightly smaller communication time variation compared with other adaptive routing algorithms. Nevertheless, Q-adaptive’s 16.96ms communication time is the smallest among all the studied routing algorithms, with only a 28.15% variation. The highest injection rate application, Halo3D, leads to significant interference on LU with the largest communication time and the greatest variation. The three adaptive routing algorithms spend an average of 27.63ms on communication with an average of 75.77% variation compared with the standalone cases. Although Q-adaptive’s average communication time is

increased to 20.24ms from 13.24ms, the interference effect is still the smallest compared with adaptive routing algorithms.

Figure 4(f) confirms the effect of message injection rate on workload interference. Significant interference effect is only observed with Halo3D background application, whose injection rate is the highest among all the studied applications. Halo3D causes LULESH to spend an average of 46.81% more time on communication under adaptive routing and only 19.62% more time under Q-adaptive routing. DL affects LULESH's communication in a much smaller magnitude with an average of 21.08% delay under adaptive routing and only 12.30% under Q-adaptive routing. Other background applications cause no more than 7.35% and 4.35% communication time increase under adaptive and Q-adaptive routing respectively.

B. Impact of Tail Latency

MPI collective performance degradation caused by network interference is manifested by long tail latency.

In Figure 4(a), Q-adaptive reduces up to 42.63% communication time for FFT3D under the network interference from Halo3D. Q-adaptive has better performance thanks to its better packet tail latency control capability. Figure 6 depicts FFT3D's packet latency distribution with 95th and 99th percentile latency for both standalone and interfered by Halo3D cases. The box ranges from the first quartile (Q1) to the third quartile (Q3) of the data, with a yellow line at the median. When FFT3D is the only job on the system, Q-adaptive has a better average packet latency with a mean of 1.27 μ s and a median of 1.16 μ s, which are 10.32% and 25.15% smaller than those of PAR. Because Q-adaptive routing starts without any pre-trained information, the learning and exploration process of reinforcement learning algorithm makes a few packets being routed with longer latency, shown as the larger tail latency under the standalone case. However, by comparing the standalone communication time in Figure 4(a), this small difference in tail latency is negligible for overall communication performance.




Fig. 6. FFT3D packet latency distribution under both standalone and interfered by Halo3D cases with the 95th and 99th percentile latencies.

When FFT3D is interfered by Halo3D, Q-adaptive routing provides much better tail latency control. Although PAR and Q-adaptive have similar median packet latency (1.55 μ s, 1.28 μ s), the 95th and 99th percentile latency of PAR are significantly delayed to 6.50 μ s and 16.92 μ s, which are 1.59x and 2.01x longer than Q-adaptive's 4.08 μ s and 8.40 μ s. For the Alltoall operation, it is only considered completed by a process when all messages are received. Therefore, the tail latency has a more important influence on the application's communication time than the mean and median. As a result, we observe that compared with PAR, although Q-adaptive has similar median packet latency, its 50.35% smaller 99th percentile latency helps FFT3D save 42.63% time spent on communication as shown in Figure 4(a) brown bars.

C. Impact of Peak Ingress Volume

Applications with larger peak ingress volume can better tolerate other application's interference by aggressively occupying network resources.

According to Table I, the peak ingress volume of LQCD is larger than all the background applications. As shown in Figure 4(c), LQCD is nearly immune to the network interference. With PAR, even the highest injection rate Halo3D application only makes LQCD spend up to 10.10% more time on communication. Similarly, the largest peak ingress volume application, Stencil5D, experiences no more than 8% communication time variation under all the background applications as shown in Figure 4(e). Stencil5D has higher communication time variance among its processes due to the imperfect multidimensional process cube. It is difficult to create a balanced 5D-grid with similar size along five dimensions. As a result, many processes are on the grid edges or surfaces with fewer neighbors, hence completing communication faster than those in the center of the grid.

LQCD has a 4D-stencil communication pattern with maximum of 8 neighbors, and Stencil5D has up to 10 neighbors. Both are considered communication intense applications with large peak ingress volume. Since router buffers and links process packets on a first-come, first-served basis, applications with a higher peak ingress volume have a greater chance to insert their packets in front of other applications. As a result, packets of applications with a smaller peak ingress volume, such as Halo3D, have to be queued and delayed.

To verify this effect, LQCD is co-run with Stencil5D, whose peak ingress volume is the largest among all the studied applications. Figure 7 illustrates the effect of packet delay by plotting packet latency along simulated time. In Figure 7(a), when LQCD is the only job in the system, its mean and 99th percentile latency with PAR are 13.59 μ s and 42.94 μ s in the first 15ms. However, introduction of Stencil5D into the system delays those values to 21.37 μ s and 77.44 μ s, a 57.32% and 80.36% increase respectively. As expected, the significantly delayed LQCD packets force the application to spend more time on communication as shown in Figure




Fig. 7. Packet latency of LQCD and Stencil5D. Stencil5D with larger peak ingress volume delays LQCD’s packets significantly under PAR.

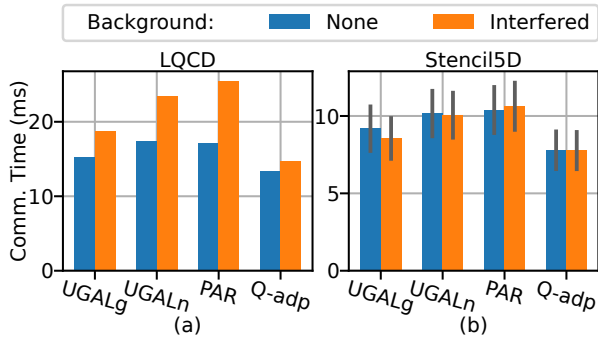


Fig. 8. Communication time of LQCD and Stencil5D. “None” means the application is the only job on the system and “interfered” denotes the pairwise running of LQCD and Stencil5D. LQCD is heavily interfered by Stencil5D, whose peak ingress volume is the largest.

8(a). Under PAR, LQCD’s communication time surges from 17.10ms to 25.51ms, a 49.14% increase. In contrast, Stencil5D with a larger peak ingress volume is barely affected by LQCD with less than 3% communication time variation under all the studied routing algorithms as shown in Figure 8(b). Q-adaptive outperforms adaptive routing algorithms by providing the smallest communication time for both applications in all scenarios. LQCD only suffers a 9.27% communication time increase from 13.43ms to 14.67ms. As a result, Q-adaptive provides better performance for both applications and reduces interference at the same time.

D. Impact of Computation

Application with long compute time and short communication time can mask communication interference.

In Figure 4(d), CosmoFlow exhibits a stable performance between standalone and interfered cases with less communication time variation. Compared with Halo3D, background applications such as UR, LU and FFT3D have lower injection rate and smaller peak ingress volume. These background applications introduce negligible network interference on CosmoFlow with less than 1.5% communication time increase

under all the studied routing algorithms. A larger interference effect on CosmoFlow is observed when co-running with DL, which has a similar peak ingress volume but a higher injection rate. CosmoFlow has to spend up to 6.27% more time on communication with UGALg, but only 1.10% with Q-adaptive routing. Halo3D as the highest injection rate application causes 15.93%, 22.03%, 21.88% communication time variation under UGALg, UGALn, and PAR respectively. However, Q-adaptive can still handle this intensive case very well with 2.90ms communication time, or a 4.88% variation. Figure 9 depicts CosmoFlow’s network throughput under the interference of Halo3D. Unlike the other applications, CosmoFlow is special because of its long computation interval between communications. As a result, Halo3D seems to be the only job on the system most of the time with a similar average throughput between the standalone and interfered cases. In Figure 9(b), CosmoFlow reaches its peak throughput of 3.15GB/ms when the first Allreduce operation is issued at 5.51ms. This high throughput pulse temporally brings down Halo3D’s throughput to 2.95GB/ms, which then quickly recovers to its average of 4.51GB/ms. However, the intermittent decrease in throughput does not affect the overall communication time.

VI. MIXED WORKLOAD ANALYSIS

In this set of experiments, we go beyond pairwise analysis by co-running multiple applications of different communication patterns, denoted as *mixed workload analysis*. The goal is to analyze workload interference under a mixture of applications. Our analysis focuses on capturing application communication performance under workload interference as well as system-wide network congestion and/or hot spots. Table II summarizes the mixed workload, in which each application takes around one-sixth of the system nodes. LQCD and Stencil5D occupy more compute nodes to construct the high-dimensional process grids.

TABLE II
APPLICATIONS IN MIXED WORKLOAD ANALYSIS

| Application | FFT3D | CosmoFlow | LU | UR | LQCD | Stencil5D |
|-------------|-------|-----------|-----|-----|------|-----------|
| Job size | 140 | 138 | 140 | 139 | 256 | 243 |




Fig. 9. CosmoFlow and Halo3D network throughput along simulated time. The peak of CosmoFlow throughput temporally drags down Halo3D’s throughput, but the over interference effect is negligible.

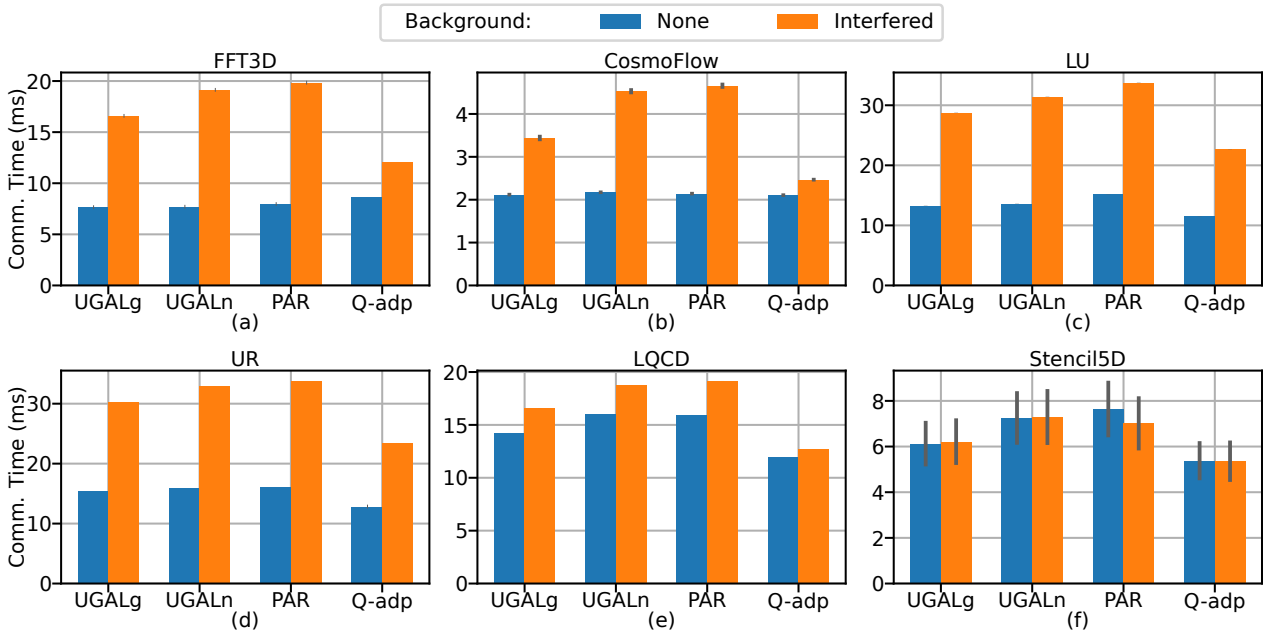


Fig. 10. Application communication time comparison: “none” denoting there is no other co-running application and “interfered” denoting multiple applications in the mixed workloads are co-running in the system.

A. Application Performance Under Workload Interference

Compared with adaptive routing, Q-adaptive routing reduces workload interference by 49.23% on average. Applications with large peak ingress volume can resist interference from other workloads.

Figure 10 depicts the network interference from the mixed workload on each individual application. Among the mixed workload applications, Stencil5D has the largest peak ingress volume and resists network interference with less than 2% communication delay caused by the other applications. The second largest peak ingress volume application, LQCD, also has a relative stable performance with an average of 17.87% communication delay under adaptive routing and only a 6.50% delay under Q-adaptive routing. On average, FFT3D,

CosmoFlow, LU, UR and LQCD spend 96.01% more time on communication under adaptive routing. However, these applications experience less communication delay under Q-adaptive routing. Specifically, compared with adaptive routing, Q-adaptive results in an average of 49.23% interference reduction for these applications.

B. System-wide Network Analysis

Q-adaptive routing achieves better system-wide performance by providing more balanced traffic distribution, less congestion, and less hot spots.

Adaptive routing algorithms are sensitive to local congestion, whereas Q-adaptive routing sends packets based on the overall network condition. As a result, in the case of a fully

loaded system, packets are more likely to be unnecessarily non-minimally forwarded by adaptive routing methods. This inefficiency consumes more network resources to deliver the same amount of traffic and leads to unbalanced network usage with hot spots.




Fig. 11. Network stall time analysis. Circles denote local links within a group, and a larger circle means higher stall time. Edges between groups denote global links, and a darker color means longer stall time on the global link. For visibility purpose, only global links from Group 0 are plotted. These plots clearly indicate that Q-adaptive routing outperforms PAR in terms of reducing network stall time on both local and global links.

Figure 11 depicts *network stall time* caused by congestion. Compared with PAR, Q-adaptive forwards packets more efficiently with less congestion shown as the smaller average stall time both within groups (31.42ms vs. 59.15ms) and between groups (0.52ms vs. 1.33ms). Q-adaptive routing avoids network congestion by having a more balanced system-wide traffic distribution than PAR. In Figure 11(b), obvious hot spots are observed at Group 1, 2, 10 and 21. These hot spots also congest the traffic from Group 0 causing a longer global link stall time. Moreover, among the overly loaded global links, PAR extremely underutilizes G0-G28 global link with only 0.59ms stall time compared with the average of 1.33ms.




Fig. 12. Heat map of global link congestion index (non-diagonal dots) and local link congestion index (diagonal dots). PAR results in unbalanced traffic distribution on global and local links with clear diagonal and vertical/horizontal lines.

We further analyze system-wide traffic distribution by adapting a metric called *congestion index* from [42]. It is defined as the ratio between average link throughput and its maximum capacity. As shown in Figure 12, compared with Q-adaptive routing, PAR generally has a darker color, indicating overall system congestion. The clear diagonal line in Figure

12(b) demonstrates the unbalanced utilization between local and global links. Additionally, the pattern of darker horizontal and vertical lines at group 8, 14, 16, 21 and lighter lines at group 2,3,5,6,27,28 unveils the unbalanced global traffic distribution.

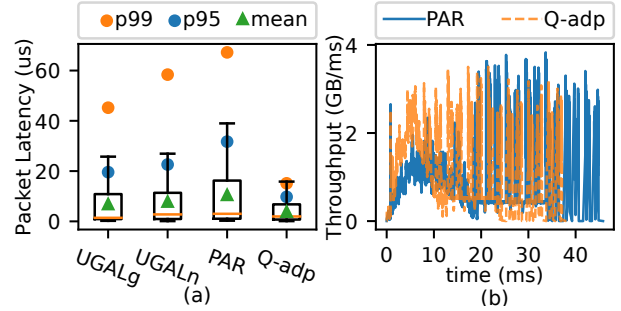


Fig. 13. System-wide packet latency distribution and aggregated network throughput under the mixed workload. Q-adaptive routing achieves significantly smaller tail latency with higher throughput.

Figure 13 depicts *system-wide packet latency distribution* and *the aggregated network throughput*. In Figure 13(a), Q-adaptive has the most concentrated packet latency distribution with the smallest average and tail latency. Q-adaptive’s 3.87 μ s average packet latency and 15.13 μ s 99th percentile latency are more than 63% smaller than those of PAR. Being able to deliver packets more quickly also leads to higher system throughput. As shown in Figure 13(b), Q-adaptive achieves an average throughput of 1.27GB/ms, which is 35.11% higher than PAR’s 0.94GB/ms average throughput.

To conclude, the mixed workload analyses demonstrate that Q-adaptive routing can effectively balance system traffic to avoid network hot spots and congestion. As a result, when Q-adaptive routing is used, the system can deliver packets more quickly, and therefore yields a higher network throughput, which in turn leads to better application performance.

VII. CONCLUSION

High-radix, low-diameter Dragonfly interconnect topology is a crucial component for exascale computing. A major problem on Dragonfly system is the network competition among co-existing applications for shared network resources (aka workload interference) such that the applications may experience a huge communication delay. In this study, through pairwise workload analysis and mixed workload network interference study, we have shown that without using any additional complicated techniques, intelligent routing such as Q-adaptive routing can greatly reduce network congestion and save communication time by up to 42.63%. We have also presented two metrics to formally quantify an application’s communication intensity: message injection rate and peak ingress volume. Several key insights have been presented in our pairwise and mixed workload analyses.

ACKNOWLEDGMENT

We thank the reviewers for their valuable feedback and insightful comments. This work is supported in part by

REFERENCES

- [1] J. Kim, W. J. Dally, S. Scott, and D. Abts, "Technology-driven, highly-scalable dragonfly topology," in *2008 International Symposium on Computer Architecture*. IEEE, 2008, pp. 77–88.
- [2] top500.org. (2021) Top500 list. [Online]. Available: <https://www.top500.org/lists/top500/2021/11/>
- [3] D. Sensi, S. Girolamo, K. McMahon, D. Roweth, and T. Hoefler, "An in-depth analysis of the slingshot interconnect," in *2020 SC20: International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. IEEE Computer Society, 2020, pp. 481–494.
- [4] B. Alverson, E. Froese, L. Kaplan, and D. Roweth, "Cray xc series network," *Cray Inc., White Paper WP-Aries01-1112*, 2012.
- [5] G. Faanes, A. Batatineh, D. Roweth, T. Court, E. Froese, B. Alverson, T. Johnson, J. Kopnick, M. Higgins, and J. Reinhard, "Cray cascade: a scalable hpc system based on a dragonfly network," in *SC'12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. IEEE, 2012, pp. 1–9.
- [6] N. Jiang, J. Kim, and W. J. Dally, "Indirect adaptive routing on large scale interconnection networks," in *Proceedings of the 36th annual international symposium on Computer architecture*, 2009, pp. 220–231.
- [7] S. Chunduri, K. Harms, S. Parker, V. Morozov, S. Oshin, N. Cherukuri, and K. Kumaran, "Run-to-run variability on xeon phi based cray xc systems," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2017, pp. 1–13.
- [8] B. Li, S. Chunduri, K. Harms, Y. Fan, and Z. Lan, "The effect of system utilization on application performance variability," in *Proceedings of the 9th International Workshop on Runtime and Operating Systems for Supercomputers*, 2019, pp. 11–18.
- [9] X. Wang, M. Mubarak, Y. Kang, R. B. Ross, and Z. Lan, "Union: An automatic workload manager for accelerating network simulation," in *2020 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 2020, pp. 821–830.
- [10] X. Yang, J. Jenkins, M. Mubarak, R. B. Ross, and Z. Lan, "Watch out for the bully! job interference study on dragonfly network," in *SC'16: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2016, pp. 750–760.
- [11] N. Jain, A. Bhatele, X. Ni, N. J. Wright, and L. V. Kale, "Maximizing throughput on a dragonfly network," in *SC'14: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2014, pp. 336–347.
- [12] D. De Sensi, S. Di Girolamo, and T. Hoefler, "Mitigating network noise on dragonfly networks through application-aware routing," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019, pp. 1–32.
- [13] K. A. Brown, N. McGlohon, S. Chunduri, E. Borch, R. B. Ross, C. D. Carothers, and K. Harms, "A tunable implementation of quality-of-service classes for hpc networks," in *International Conference on High Performance Computing*. Springer, 2021, pp. 137–156.
- [14] Y. Kang, X. Wang, and Z. Lan, "Q-adaptive: A multi-agent reinforcement learning based routing on dragonfly network," in *Proceedings of the 30th International Symposium on High-Performance Parallel and Distributed Computing*, 2021, pp. 189–200.
- [15] A. Bhatele, N. Jain, W. D. Gropp, and L. V. Kale, "Avoiding hot-spots on two-level direct networks," in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, 2011, pp. 1–11.
- [16] N. McGlohon, C. D. Carothers, K. Hemmert, M. Levenhagen, K. A. Brown, S. Chunduri, and R. B. Ross, "Exploration of congestion control techniques on dragonfly-class hpc networks through simulation," in *2021 International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*. Los Alamitos, CA, USA: IEEE Computer Society, nov 2021, pp. 40–50. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/PMBS54543.2021.00010>
- [17] M. Mubarak, P. Carns, J. Jenkins, J. K. Li, N. Jain, S. Snyder, R. Ross, C. D. Carothers, A. Bhatele, and K.-L. Ma, "Quantifying i/o and communication traffic interference on dragonfly networks equipped with burst buffers," in *2017 IEEE International Conference on Cluster Computing (CLUSTER)*, 2017, pp. 204–215.
- [18] Y. Kang, X. Wang, and Z. Lan, "Study of Workload Interference with Intelligent Routing on Dragonfly," Aug. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.6975878>
- [19] M. Flajslik, E. Borch, and M. A. Parker, "Megafly: A topology for exascale systems," in *International Conference on High Performance Computing*. Springer, 2018, pp. 289–310.
- [20] A. Shpiner, Z. Haramaty, S. Eliad, V. Zdornov, B. Gafni, and E. Zahavi, "Dragonfly+: Low cost topology for scaling datacenters," in *2017 IEEE 3rd International Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era (HiPINEB)*. IEEE, 2017, pp. 1–8.
- [21] J. Won, G. Kim, J. Kim, T. Jiang, M. Parker, and S. Scott, "Overcoming far-end congestion in large-scale networks," in *2015 IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2015, pp. 415–427.
- [22] Y. Kang, X. Wang, N. McGlohon, M. Mubarak, S. Chunduri, and Z. Lan, "Modeling and analysis of application interference on dragonfly+," in *Proceedings of the 2019 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, 2019, pp. 161–172.
- [23] M. Mubarak, N. McGlohon, M. Musleh, E. Borch, R. B. Ross, R. Huggahalli, S. Chunduri, S. Parker, C. D. Carothers, and K. Kumaran, "Evaluating quality of service traffic classes on the megafly network," in *International Conference on High Performance Computing*. Springer, 2019, pp. 3–20.
- [24] J. J. Wilke and J. P. Kenny, "Opportunities and limitations of quality-of-service in message passing applications on adaptively routed dragonfly and fat tree networks," in *2020 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 2020, pp. 109–118.
- [25] S. A. Smith, C. E. Cromey, D. K. Lowenthal, J. Domke, N. Jain, J. J. Thiagarajan, and A. Bhatele, "Mitigating inter-job interference using adaptive flow-aware routing," in *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2018, pp. 346–360.
- [26] N. McGlohon, C. D. Carothers, K. S. Hemmert, M. Levenhagen, K. A. Brown, S. Chunduri, and R. B. Ross, "Exploration of congestion control techniques on dragonfly-class hpc networks through simulation," in *2021 International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, 2021, pp. 40–50.
- [27] 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, 2016.
- [28] A. F. Rodrigues, K. S. Hemmert, B. W. Barrett, C. Kersey, R. Oldfield, M. Weston, R. Risen, J. Cook, P. Rosenfeld, E. Cooper-Balis *et al.*, "The structural simulation toolkit," *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 4, pp. 37–42, 2011.
- [29] S. J. Pennycook, S. D. Hammond, S. A. Jarvis, and G. R. Mudalige, "Performance analysis of a hybrid mpi/cuda implementation of the naslu benchmark," *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 4, pp. 23–29, 2011.
- [30] C. H. Still, R. Berger, A. Langdon, D. Hinkel, L. Suter, and E. Williams, "Filamentation and forward brillouin scatter of entire smoothed and aberrated laser beams," *Physics of Plasmas*, vol. 7, no. 5, pp. 2023–2032, 2000.
- [31] J. C. Phillips, G. Zheng, S. Kumar, and L. V. Kalé, "Namd: Biomolecular simulation on thousands of processors," in *SC'02: Proceedings of the 2002 ACM/IEEE Conference on Supercomputing*. IEEE, 2002, pp. 36–36.
- [32] G. Kresse and J. Hafner, "Ab initio molecular dynamics for liquid metals," *Physical Review B*, vol. 47, no. 1, p. 558, 1993.
- [33] D. Unat, X. Cai, and S. B. Baden, "Mint: realizing cuda performance in 3d stencil methods with annotated c," in *Proceedings of the international conference on Supercomputing*, 2011, pp. 214–224.
- [34] R. Babich, M. A. Clark, and B. Joó, "Parallelizing the quda library for multi-gpu calculations in lattice quantum chromodynamics," in *SC'10: Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2010, pp. 1–11.
- [35] A. Sergeev and M. Del Balso, "Horovod: fast and easy distributed deep learning in tensorflow," *arXiv preprint arXiv:1802.05799*, 2018.
- [36] A. Mathuriya, D. Bard, P. Mendygral, L. Meadows, J. Arnemann, L. Shao, S. He, T. Kärnä, D. Moise, S. J. Pennycook *et al.*, "Cos-

- moflow: Using deep learning to learn the universe at scale,” in *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2018, pp. 819–829.
- [37] I. Karlin, A. Bhatele, J. Keasler, B. L. Chamberlain, J. Cohen, Z. DeVito, R. Haque, D. Laney, E. Luke, F. Wang *et al.*, “Exploring traditional and emerging parallel programming models using a proxy application,” in *2013 IEEE 27th International Symposium on Parallel and Distributed Processing*. IEEE, 2013, pp. 919–932.
- [38] I. Karlin, J. Keasler, and R. Neely, “Lulesh 2.0 updates and changes,” Tech. Rep. LLNL-TR-641973, August 2013.
- [39] C. D. Carothers, J. S. Meredith, M. P. Blanco, J. S. Vetter, M. Mubarak, J. LaPre, and S. Moore, “Durango: Scalable synthetic workload generation for extreme-scale application performance modeling and simulation,” in *Proceedings of the 2017 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, 2017, pp. 97–108.
- [40] P. C. Roth, J. S. Meredith, and J. S. Vetter, “Automated characterization of parallel application communication patterns,” in *Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing*, 2015, pp. 73–84.
- [41] Y. Zheng, A. Kamil, M. B. Driscoll, H. Shan, and K. Yelick, “Upc++: a pgas extension for c++,” in *2014 IEEE 28th International Parallel and Distributed Processing Symposium*. IEEE, 2014, pp. 1105–1114.
- [42] F. He, X. Yan, Y. Liu, and L. Ma, “A traffic congestion assessment method for urban road networks based on speed performance index,” *Procedia engineering*, vol. 137, pp. 425–433, 2016.

Appendix: Artifact Description/Artifact Evaluation

SUMMARY OF THE EXPERIMENTS REPORTED

The experiments reported in the paper are simulated with the enhanced SST simulation toolkit. The enhancements include the addition of adaptive routing mechanisms (UGAL, PAR), intelligent routing (Q-adaptive), the extension of SST-Elements "Merlin" and "Ember", and several HPC and ML applications. The enhanced toolkit is open-sourced with the GitHub link provided below. The experiments setup configuration file is also provided along with the published code. The experiments include pairwise workload simulations between a target application and a background application. The target applications include FFT3D, LU, LQCD, and Cosmoflow. The background applications include UR, LU, FFT3D, DL, and Halo3D. All possible pairwise combinations are studied with an additional case of LQCD-Stencil5D simulation. A mixed workload consisting of FFT3D, CosmoFlow, LU, UR, and LQCD is also studied in the paper. For all the simulated workloads, we evaluated the application and system performance under UGALg, UGALn, PAR, and Q-adaptive routing algorithms. We used the Argonne JLSE Skylake cluster where each node contains two Intel Xeon Platinum 8180M CPUs with 768GB RAM. The simulation toolkit is compiled with openMPI/4.1.1, gcc/11.1.0, Python/3.6

AUTHOR-CREATED OR MODIFIED ARTIFACTS:

Artifact 1

Persistent ID: https://github.com/SPEAR-IIT/q-adaptive_sst/tree/df_interference_rl

Artifact name: Enhanced SST toolkit for Dragonfly network interference study

Citation of artifact: Arun F Rodrigues, K Scott Hemmert, Brian W Barrett, Chad Kersey, Ron Oldfield, Marlo Weston, Rolf Risen, Jeanine Cook, Paul Rosenfeld, Elliot Cooper-Balis, et al. 2011. The structural simulation toolkit. ACM SIGMETRICS Performance Evaluation Review 38, 4 (2011), 37–42. Y. Kang, X. Wang, and Z. Lan, "Q-adaptive: A multi-agent reinforcement learning based routing on dragonfly network," in Proceedings of the 30th International Symposium on High-Performance Parallel and Distributed Computing, 2021, pp. 189–200.

Artifact 2

Persistent ID: [10.5281/zenodo.6975878](https://zenodo.org/record/6975878)

Artifact name: DOI for the source code

Citation of artifact: Y. Kang, X. Wang, and Z. Lan, "Study of Workload Interference with Intelligent Routing on Dragonfly" SC'22

Reproduction of the artifact with container: The experiments reported in the paper are network system simulations. The simulation configuration and setup parameters are provided along with the published source code. The provided container contains both the source code and the compiled executables. A running script is also provided to reproduce the reported experiments