

PIAS: Practical Information-Agnostic Flow Scheduling for Commodity Data Centers

Wei Bai, Li Chen, Kai Chen, Dongsu Han, Chen Tian, and Hao Wang

Abstract—Many existing data center network (DCN) flow scheduling schemes, that minimize flow completion times (FCT) assume prior knowledge of flows and custom switch functions, making them superior in performance but hard to implement in practice. By contrast, we seek to minimize FCT with no prior knowledge and existing commodity switch hardware. To this end, we present PIAS, a DCN flow scheduling mechanism that aims to minimize FCT by mimicking shortest job first (SJF) on the premise that flow size is not known *a priori*. At its heart, PIAS leverages multiple priority queues available in existing commodity switches to implement a multiple level feedback queue, in which a PIAS flow is gradually demoted from higher-priority queues to lower-priority queues based on the number of bytes it has sent. As a result, short flows are likely to be finished in the first few high-priority queues and thus be prioritized over long flows in general, which enables PIAS to emulate SJF without knowing flow sizes beforehand. We have implemented a PIAS prototype and evaluated PIAS through both testbed experiments and ns-2 simulations. We show that PIAS is readily deployable with commodity switches and backward compatible with legacy TCP/IP stacks. Our evaluation results show that PIAS significantly outperforms existing information-agnostic schemes, for example, it reduces FCT by up to 50% compared to DCTCP [11] and L2DCT [32]; and it only has a 1.1% performance gap to an ideal information-aware scheme, pFabric [13], for short flows under a production DCN workload.

Index Terms—Data center networks, flow scheduling.

I. INTRODUCTION

IT HAS been a virtually unanimous consensus in the community that one of the most important objectives for data center network (DCN) transport designs is to minimize the flow completion times (FCT) [11]–[13], [24], [31], [32].

Manuscript received May 9, 2016; revised December 31, 2016; accepted January 29, 2017; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor A. Wierman. This work was supported in part by the Hong Kong RGC under Grant ECS-26200014, Grant GRF-16203715, Grant GRF-613113, and Grant CRF-C703615G; in part by the China 973 under Grant 2014CB340303; in part by the China NSFC under Grant 61602194 and Grant 61321491; in part by the Collaborative Innovation Center of Novel Software Technology and Industrialization, and the Jiangsu Innovation and Entrepreneurship (Shuangchuang) Program, the ICT R&D Program of MSIP/IITP, Republic of Korea, under Grant B0101-15-1368 and Grant B0126-15-1078; and in part by the Basic Science Research Program of NRF within the MSIP, Republic of Korea, under Grant 2013R1A1A1076024.

W. Bai, L. Chen, and K. Chen are with The Hong Kong University of Science and Technology, Hong Kong (e-mail: wbaia@cse.ust.hk; lchenad@cse.ust.hk; kaichen@cse.ust.hk).

D. Han is with the Korea Advanced Institute of Science and Technology, Daejeon 34141, South Korea (e-mail: dongsuh@ee.kaist.ac.kr).

C. Tian is with the State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210000, China (email: tianchen@nju.edu.cn).

H. Wang is with the University of Toronto, Toronto, ON M5S 3G4, Canada (e-mail: wh.sjtu@gmail.com).

Digital Object Identifier 10.1109/TNET.2017.2669216

This is because many of today’s cloud applications, such as web search, social networking, and retail recommendation, etc., have very demanding latency requirements, and even a small delay can directly affect application performance and degrade user experience [11], [32].

To minimize FCT, most previous proposals [13], [24], [31], [43] assume prior knowledge of accurate flow information, *e.g.*, flow sizes or deadlines, to achieve superior performance. For example, PDQ, pFabric and PASE [13], [24], [31] all assume flow size is known *a priori*, and attempt to approximate Shortest Job First (SJF, preemptive), which is the optimal scheduling discipline for minimizing the average FCT over a single link. In this paper, we question the validity of this assumption, and point out that, for many applications, such information is difficult to obtain, and may even be unavailable (§II). Existing transport layer solutions with this assumption are therefore very hard to implement in practice.

We take one step back and ask: without prior knowledge of flow size information, what is the best scheme that minimizes FCT with existing commodity switches?

Motivated by the above question, we set up our key design goals as follows:

- **Information-agnostic:** Our design must not assume *a priori* knowledge of flow size information being available from the applications.
- **FCT minimization:** The solution must be able to enforce a near-optimal information-agnostic flow schedule. It should minimize the average and tail FCT for latency-sensitive short flows, while not adversely affecting the FCT of long flows.
- **Readily-deployable:** The solution must work with existing commodity switches in DCNs and be backward compatible with legacy TCP/IP stacks.

When exploring possible solution spaces, we note that some existing approaches such as DCTCP, Timely, HULL, etc. [11], [12], [30], [32], [47] reduce FCT without relying on flow size information. They generally improve FCT by maintaining low queue occupancy through mechanisms like adaptive congestion control, ECN, pacing, etc. However, they do not provide a full answer to our question, because they mainly perform end-host based rate control which is ineffective for flow scheduling [13], [31].

In this paper, we answer the question with PIAS, a practical information-agnostic flow scheduling that minimizes the FCT in DCNs. In contrast to previous FCT minimization schemes [13], [24], [31] that emulate SJF by using prior knowledge of flow sizes, PIAS manages to mimic SJF with

no prior information. At its heart, PIAS leverages multiple priority queues available in existing commodity switches to implement a Multiple Level Feedback Queue (MLFQ), in which a PIAS flow is gradually demoted from higher-priority queues to lower-priority queues based on the bytes it has sent during its lifetime. In this way, PIAS ensures in general that short flows are prioritized over long flows, effectively emulating SJF without knowing flow sizes beforehand.

Generally speaking, there are three components in the PIAS framework. A central entity collects traffic information reported from each end host periodically, calculates the demotion thresholds based on the aggregated traffic statistics from the entire network, and distributes the thresholds to all end hosts. Then each host uses such thresholds to tag the traffic. And finally, the PIAS switches simply perform strict priority queueing based on the packet tags.

However, we face several concrete challenges before making PIAS truly effective. First, how to determine the demotion threshold for each queue of MLFQ? Second, as the traffic varies across both time and space, how to keep PIAS's performance in such a dynamic environment? Third, how to ensure PIAS's compatibility with legacy TCP/IP stacks in production DCNs?

For the first challenge, we derive a set of optimal demotion thresholds for MLFQ through solving a FCT minimization problem. We further show that the derived threshold set is robust to a reasonable range of traffic distributions. This enables PIAS to distribute the demotion thresholds to the end hosts for packet tagging while only performing strict priority queueing, a built-in function, in the PIAS switches.

For the second challenge, since one set of demotion thresholds only works the best for a certain range of traffic distributions, PIAS adjusts the thresholds to keep up with the traffic dynamics. However, the key problem is that a mismatch between thresholds and underlying traffic is inevitable. Once happens, short flows may be adversely affected by large ones in a queue, impacting on their latency. Inspired by ideas from ECN-based rate control [11], PIAS employs ECN to mitigate the mismatch problem. Our insight is that, by maintaining low queue occupancy, short flows always see small queues and thus will not be seriously delayed even if they are mistakenly placed in a queue with a long flow due to mismatched thresholds.

For the third challenge, we employ DCTCP-like transport at the PIAS end hosts and find that PIAS interacts favorably with DCTCP or other legacy TCP protocols with ECN enabled. A potential problem is that many concurrent short flows may starve a coexisting long flow, triggering TCP timeouts and degrading application performance. We measure the extent of starvation on our testbed with a realistic workload and analyze possible solutions. We ensure that all mechanisms in PIAS can be implemented by a shim layer over NIC without touching the TCP stack.

We have implemented a PIAS prototype (§IV). On the end host, we implement PIAS as a kernel module in Linux, which resides between the Network Interface Card (NIC) driver and the TCP/IP stack as a shim layer. It does not touch any TCP/IP implementation that natively supports various OS versions.

In virtualized environments, PIAS can also support virtual machines well by residing in hypervisor (or Dom 0). At the switch, PIAS only needs to enable priority queues and ECN which are both built-in functions readily supported by existing commodity switch hardware.

We evaluate PIAS on a small-scale testbed with 16 Dell servers and a commodity Pronto-3295 Gigabit Ethernet switch (Broadcom BCM#56538). In our experiments, we find that PIAS reduces the average FCT for short flows by $\sim 29\text{-}49\%$ and $\sim 18\text{-}34\%$ compared to DCTCP under two realistic DCN traffic patterns. It also improves the query performance by $\sim 28\text{-}30\%$ in a Memcached [6] application (§V-A). We further dig into different design components of PIAS such as queues, optimal demotion threshold setting, ECN, and demonstrate the effectiveness of each of their contributions to PIAS's performance (§V-B).

To complement our small-scale testbed experiments, we further conduct large-scale simulations in a simulated 10/40G network using ns-2 [9]. In our simulations, we show that PIAS outperforms all existing information-agnostic solutions under realistic DCN workloads, reducing the average FCT for short flows by up to 50% compared to DCTCP and L2DCT. In addition, our results show that PIAS, as a readily-deployable information-agnostic scheme, also delivers a comparable performance to a clean-slate information-aware design, pFabric [13], in certain scenarios. For example, there is only a 1.1% gap to pFabric for short flows in a data mining workload [22] (§V-C).

To make our work easy to reproduce, we made our implementation and evaluation scripts available online at: <http://sing.cse.ust.hk/projects/PIAS>.

II. MOTIVATION

To motivate our design, we introduce a few cases in which the accurate flow size information is hard to obtain, or simply not available.

HTTP Chunked Transfer: Chunked transfer has been supported since HTTP 1.1 [20], where dynamically generated content is transferred during the generation process. This mode is widely used by data center applications. For example, applications can use chunked transfer to dump database content into OpenStack Object Storage [8]. In chunked transfer, a flow generally consists of multiple chunks, and the total flow size is not available at the start of the transmission.

Database Query Response: Query response in database systems, such as Microsoft SQL Server [7], is another example. Typically, SQL servers send partial query results as they are created, instead of buffering the result until the end of the query execution process [7]. The flow size again is not available at the start of a flow.

Stream Processing: Stream processing systems are currently gaining popularity. In Apache Storm [2], after the master node distributes tasks to worker nodes, workers will analyze the tasks and pre-establish persistent connections with related worker nodes. During the data processing, data tuples completed in one node are continuously delivered to the next node in the stream processing chain. The amount of data to be processed is unknown until the stream finishes.

Practical Limitations: We note that there are certain cases where the flow size information can be obtained or inferred. For example, in Hadoop [4] the mapper will first write the intermediate data to disk before the corresponding reducer starts to fetch the data, thus the flow size can be obtained in advance [33]. Even so, practical implementation issues are still prohibitive. First, we need to patch all modules in every application that generate network traffic, which is a burden for applications programmers and/or network operators. Second, current operating systems lack appropriate interface for delivering the flow size information to the transport layer. Thus, kernel modifications are also required.

III. THE PIAS DESIGN

A. Design Rationale

Compared to previous solutions [13], [24], [31] that emulate SJF based on prior knowledge of flow sizes, PIAS distinguishes itself by emulating SJF with no prior information. At its core, PIAS exploits multiple priority queues available in commodity switches and implements a MLFQ, in which a PIAS flow is demoted from higher-priority queues to lower-priority queues dynamically according to its bytes sent. Through this way, PIAS enables short flows to finish in the first few priority queues, and thus in general prioritizes them over long flows, effectively mimicking SJF without knowing the flow sizes.

We note that scheduling with no prior knowledge is known as non-clairvoyant scheduling [28]. Least Attained Service (LAS) is one of the best known algorithms that minimize the average FCT in this case [37]. LAS tries to approximate SJF by guessing the remaining service time of a job based on the service it has attained so far. LAS is especially effective in DCN environments where traffic usually exhibits long-tail distribution—most flows are short and a small percent are very large [11], [22], and LAS favors short flows.

PIAS is partially inspired by LAS. However, we note that directly enabling LAS at the switch requires us to compare the amount of bytes transferred for each flow, which is not supported in existing commodity switches. Furthermore, although DCN traffic is generally long-tailed, it varies across both time and space, and on some switch ports the distribution may temporarily not be so. Blindly using LAS will exacerbate the problem when multiple long flows coexist. This is because pure LAS needs to compare and prioritize between any two flows, when some long flows meet a longer flow, it is possible that they can collectively cause the longer one to starve.

To this end, PIAS leverages multiple priority queues available in existing commodity switches (typically 4–8 queues per port [13]) to implement a MLFQ (see Figure 1). Packets in different queues of MLFQ are scheduled with strict priority, while packets in the same queue are scheduled based on FIFO. In a flow’s lifetime, it is demoted dynamically from i th queue down to the $(i+1)$ th queue after transmitting more bytes than queue i ’s demotion threshold, until it enters the last queue. To further prevent switches from maintaining the per-flow state, PIAS distributes packet priority tagging (indicating a flow’s sent size) to end hosts, allowing the PIAS switches to

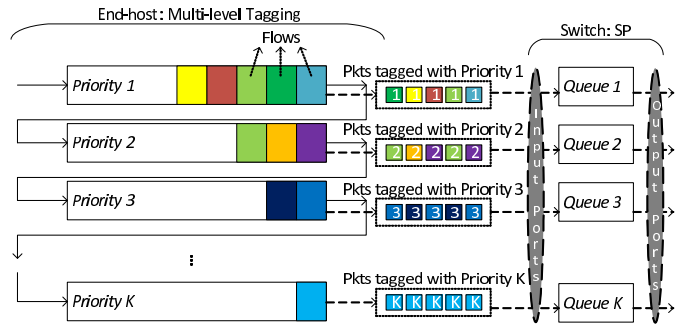


Fig. 1. PIAS overview.

perform strict priority queueing only, which is already a built-in function in today’s commodity switches.

By implementing MLFQ, PIAS gains two benefits. First, it prioritizes short flows over large ones because short flows are more likely to finish in the first few higher priority queues while large flows are eventually demoted to lower priority queues. This effectively enables PIAS to approximate SJF scheduling that optimizes average FCT while being readily implementable with existing switch hardware. Second, it allows large flows that are demoted to the same low priority queues to share the link fairly. This helps to minimize the response time of long flows, mitigating the starvation problem.

However, there are several concrete challenges to consider in order to make PIAS truly effective. First, how to determine the demotion threshold for each queue of MLFQ to minimize the FCT? Second, as DCN traffic varies across both time and space, how to make PIAS perform efficiently and stably in such a dynamic environment? Third, how to ensure PIAS’s compatibility with legacy TCP/IP stacks in production DCNs? Next, we explain the details of the mechanism we design to address all these challenges.

B. Detailed Mechanisms

At a high level, PIAS’s main mechanisms include packet tagging, switch design, and rate control.

1) *Packet Tagging at End-Hosts:* PIAS performs packet tagging at each end host as shown in Figure 1. There are K priorities $P_i, 1 \leq i \leq K$ and $(K-1)$ demotion thresholds $\alpha_j, 1 \leq j \leq K-1$. We assume $P_1 > P_2 \dots > P_K$ and $\alpha_1 \leq \alpha_2 \dots \leq \alpha_{K-1}$.

At the end host, when a new flow is initialized, its packets will be tagged with the highest priority P_1 , giving it the highest priority in the network. As more bytes are sent, the packets of this flow will be tagged with decreasing priorities P_j ($2 \leq j \leq K$) and receive decreasing priorities in the network. The threshold to demote priority from P_{j-1} to P_j is α_{j-1} .

One challenge is to determine the demotion threshold for each priority to minimize the average FCT. By solving a FCT minimization problem, we derive a set of analytical solutions for optimal demotion thresholds (details in §III-C). Note that in PIAS we calculate the thresholds based on traffic information from the entire DCN and distribute the same threshold setting to all the end hosts. Our experiments and analysis show that such threshold setting is effective and also robust to a certain range of traffic variations (§V). This is a key

reason we can decouple packet tagging from switches to end hosts while still maintaining good performance, which relieves the PIAS switches of having to keep the per-flow state.

As traffic changes over time, PIAS needs to adjust the demotion thresholds accordingly. To keep track of traffic variations, each end host can periodically report its local traffic information to a central entity for statistics, and there are many existing techniques available for this purpose [46]. However, historical traffic cannot predict the future perfectly, and traffic distribution also differs across different links. Mismatches between threshold setting and underlying traffic are inevitable, which can hurt latency sensitive short flows. Therefore, mitigating the impact of the mismatch is a must for PIAS to operate in the highly dynamic DCNs. Our solution, as shown subsequently, is to employ ECN.

2) *Switch Design*: The PIAS switches enable the following two basic mechanisms, which are built-in functions for existing commodity switches [31].

- Priority scheduling: Packets are dequeued based on their priorities strictly when a fabric port is idle.
- ECN marking: The arriving packet is marked with Congestion Experienced (CE) if the instant buffer occupancy is larger than the marking threshold.

With priority scheduling at the switches and packet tagging at the end hosts, PIAS performs MLFQ-based flow scheduling using stateless switches in the network. Packets with different priority tags are classified into different priority queues. When the link is idle, the head packet from the highest non-empty priority queue is transmitted.

We choose priority scheduling over weighted fair queueing (WFQ) for two reasons. First, strict priority queueing provides better in-network prioritization than WFQ and thus can achieve lower average FCT. Second, packets from a PIAS flow will be placed and scheduled in multiple queues, with WFQ, it may cause packet out-of-order problem as a latter packet has a chance to dequeue before an earlier packet, which degrades TCP performance. However, with strict priority queueing, PIAS is free of packet reordering, because a flow is always demoted from a higher priority queue to a lower priority queue, and an earlier packet is guaranteed to dequeue before a latter packet at each hop. Compared with priority scheduling, one advantage of WFQ is that it can potentially avoid starvation for long-lived flows. We quantify the starvation problem using testbed experiments subsequently and find that, even with priority scheduling, PIAS is not adversely affected by starvation.

Our intention to employ ECN is to mitigate the effect of the mismatch between the demotion thresholds and the traffic distribution. We use a simple example to illustrate the problem and the effectiveness of our solution. We connect 4 servers to a Gigabit switch as in Figure 2. One server is receiver (R) and the other three are senders (S1, S2 and S3). In our experiment, the receiver R continuously fetches 10MB data from S1 and S2, and 20KB data from S3. We configure the strict priority queueing with 2 queues on the switch egress port to R. Since there are two priorities, we only need to have one demotion threshold from the high priority queue to the




Fig. 2. Illustration example: (a) threshold right; (b) threshold too small, packets of short flow get delayed by long flow after prematurely demoted to the low priority queue; (c) threshold too large, packets of large flow stay too long in the high priority queue, affecting short flow.

TABLE I
COMPLETION TIMES OF 20KB SHORT FLOWS

	10KB threshold		20KB threshold		1MB threshold	
	Avg	99th	Avg	99th	Avg	99th
w/o ECN	1660 μ s	2257 μ s	599 μ s	676 μ s	685 μ s	3345 μ s
w/ ECN	799 μ s	916 μ s	603 μ s	660 μ s	622 μ s	1384 μ s

low priority queue. Clearly, in this case the optimal demotion threshold should be 20KB.

We intentionally apply three different demotion thresholds 10KB, 20KB and 1MB respectively, and measure the FCT of the 20KB short flows. Table I shows the results of PIAS and PIAS without ECN. When the threshold is 20KB (right), both PIAS and PIAS without ECN achieve an ideal FCT. However, with a larger threshold (1MB) or a smaller threshold (10KB), PIAS exhibits pronounced advantages over PIAS without ECN at both the average and 99th percentile. This is because, 1) if the threshold is too small, packets of short flows prematurely enter the low priority queue and experience queueing delay behind the long flows (see the scenario in Figure 2 (b)); 2) if the threshold is too large, packets of long flows over-stay in the high priority queue, thus also affecting the latency of short flows (see scenario in Figure 2 (c)).

By employing ECN, we can keep low buffer occupancy and minimize the impact of long flows on short flows, which makes PIAS more robust to the mismatch between the demotion thresholds and traffic distribution.

3) *Rate Control*: PIAS employs DCTCP [11] as end host transport, and other legacy TCP protocols with ECN enabled (e.g., ECN* [44] and DCQCN [47]) can also be integrated into PIAS. We require PIAS to interact smoothly with the legacy TCP stack. One key issue is to handle flow starvation: when packets of a large flow get starved in a low priority queue for long time, this may trigger TCP timeouts and retransmissions. The frequent flow starvation may disturb the transport layer and degrade application performance. For example, a TCP connection which is starved for long time may be terminated unexpectedly.

To address the problem, we first note that PIAS can well mitigate the starvation between long flows, because two long flows in the same low priority queue will fairly share the link in a FIFO manner. In this way, PIAS minimizes the response time of each long flow, effectively eliminating TCP timeouts.

However, it is still possible that many concurrent short flows will starve a long flow, triggering its TCP timeouts. To quantify

this, we run the web search benchmark traffic [11] (as shown in Figure 4) at 0.8 load in our 1G testbed, which has 16 servers connected to a Gigabit switch. We set RTO_{min} to 10ms and allocate 8 priority queues for PIAS. We enable both ECN and dynamic buffer management in our switch. In such setting, TCP timeouts, if happen, will be mainly caused by starvation rather than packet drops. This experiment consists of 5,000 flows, around 5.7 million MTU-sized packets. We measure the number of TCP timeouts to quantify the extent of starvation. We find that there are only 200 timeout events and 31 two consecutive timeout events in total. No TCP connection is terminated unexpectedly. The result indicates that, even at a high load, starvation is not common and will not degrade application performance adversely. We believe one possible reason is that the per-port ECN we used (see §IV-A2) may mitigate starvation by pushing back high priority flows when many packets from low priority long flows get starved. Another possible solution for handling flow starvation is treating a long-term starved flow as a new flow. For example, if a flow experiences two consecutive timeouts, we set its bytes sent back to zero. This ensures that a long flow can always make progress after timeouts. Note that the implementation of the above mechanism can be integrated to our packet tagging module without any changes to the networking stack.

4) *Discussion*: We further discuss two potential issues.

Local Decision: The key idea of PIAS is to emulate SJF which is optimal to minimize average FCT over a single link. However, there does not exist an optimal scheduling policy to schedule flows over an entire DCN with multiple links [13]. In this sense, similar to pFabric [13], PIAS also makes switch local decisions. This approach in theory may lead to some performance loss over the fabric [31]. For example, when a flow traverses multiple hops and gets dropped at the last hop, it causes bandwidth to be wasted on the upstream links that could otherwise have been used to schedule other flows. We note that some existing solutions [24], [31] leverage arbitration, where a common network entity allocates rates to each flow based on global network visibility, to address this problem. However, they are hard to implement because they require either non-trivial switch changes [24] or a complex control plane [31], which is against our design goal. Furthermore, some arbitration mechanisms [31] implicitly assume that each flow uses a single path. This assumption does not hold when network operators use advanced per-packet/flowlet/flowcell load balancing schemes [10], [17], [23]. Fortunately, local-decision based solutions maintain very good performance for most scenarios [13] and only experience performance loss at extremely high loads, *e.g.*, over 90% [31]. However, most DCNs operate at moderate loads, *e.g.*, 30% [15]. Our ns-2 simulation (§V-C) with production DCN traffic further confirms that PIAS works well in practice.

Demotion Threshold Updating: By employing ECN, PIAS can effectively handle the mismatch between the demotion thresholds and traffic distribution (see §V-B). This suggests that we do not need to frequently change our demotion thresholds which may be an overhead. In this paper, we simply assume the demotion thresholds are updated periodically according to the network scale (which decides the time for

information collection and distribution) and leave dynamic threshold updates as future work.

C. Optimizing Demotion Thresholds

In this section, we describe our formulation to derive the optimal demotion thresholds for minimizing the average FCT. We leverage existing optimization software to derive the optimal thresholds numerically for *any* given load and flow size distribution. We find that the demotion thresholds depend on both load and flow size distribution. As flow size distribution and load change across both time and space, ideally, one should use different thresholds for different links at different times. However, in practice, it is quite challenging to obtain such fine-grained link level traffic information across the entire DCN. Hence, we use the overall flow size distribution and load measured in the entire DCN as an estimate to derive a common set of demotion thresholds for all end hosts. We note that this approach is not optimal and there is room for improvement. However, it is more practical and provides considerable gains, as shown in our evaluation (§V).

Problem Formulation: We assume there are K priority queues P_i , $1 \leq i \leq K$, where P_1 has the highest priority. We denote the threshold for demoting the priority from $i-1$ to i as α_{i-1} , $2 \leq i \leq K$. We define $\alpha_K = \infty$, so that the largest flows are all in this queue, and $\alpha_0 = 0$.

Denote the cumulative density function of flow size distribution as $F(x)$, thus $F(x)$ is the probability that a flow size is no larger than x . Let L_i denote the number of packets a given flow brings in queue P_i for $i = 1, \dots, K$. Thus, $E[L_i] \leq (\alpha_i - \alpha_{i-1})(1 - F(\alpha_{i-1}))$. Denote flow arrival rate as λ , then the packet arrival rate to queue i is $\lambda_i = \lambda E[L_i]$.

The service rate for a queue depends on whether the queues with higher priorities are all empty. Thus, P_1 (highest priority) has capacity $\mu_1 = \mu$ where μ is the service rate of the link. The fraction of time that queue 1 is idle is $(1 - \rho_1)$ where $\rho_i = \lambda_i / \mu_i$ is simply the utilization of queue P_i . Thus, the service rate of P_2 is $\mu_2 = (1 - \rho_1)\mu$ since its service rate is μ (the full link capacity) given that P_1 is empty. Thus, we have $\mu_i = \prod_{j=0}^{i-1} (1 - \rho_j)\mu$, with $\rho_0 = 0$. Thus, $T_i = 1/(\mu_i - \lambda_i)$ which is the average delay of queue i assuming M/M/1 queues.¹

For a flow with size in $[\alpha_{i-1}, \alpha_i)$, it experiences the delays in different priority queues up to the i -th queue. Denote T_i as the average time spent in the i -th queue. Let $i_{max}(x)$ be the index of the smallest demotion threshold larger than x . So the average FCT for a flow with size x , $T(x)$, is upper-bounded by: $\sum_{i=1}^{i_{max}(x)} T_i$.

Thus we have an optimization problem of choosing an optimal set of thresholds $\{\alpha_i\}$ to minimize the average FCT of flows on this bottleneck link. Let $\theta_i = F(\alpha_i) - F(\alpha_{i-1})$ denote the percentage of flows with sizes in $[\alpha_{i-1}, \alpha_i)$. Using θ_i to equivalently express α_i , we can formulate the problem as:

$$\begin{aligned} \min_{\{\theta_i\}} \mathcal{T} &= \sum_{l=1}^K (\theta_l \sum_{m=1}^l T_m) = \sum_{l=1}^K (T_l \sum_{m=l}^K \theta_m) \\ \text{subject to } \theta_i &\geq 0, \quad i = 1, \dots, K-1 \end{aligned} \quad (1)$$

¹We use M/M/1 queues to simplify the analysis, and the robustness of this formulation is verified in §V experimentally.

The solution of this problem is $\{\theta_i^{opt}\}$, from which we can retrieve the thresholds $\{\alpha_i^{opt}\}$ with $\alpha_i^{opt} = F^{-1}(\sum_{j=1}^i \theta_j)$, where $F^{-1}(\cdot)$ is the inverse of $F(\cdot)$.

Solution Method: We have products of the variables (θ_i) both in numerators and denominators of the objective. Since the number of variables is small (the number of queues ≤ 8), the optimization problem can be solved numerically and quickly with existing solvers. We use global optimization toolbox available in SciPy [27], which searches for the optimal solution using basin-hopping algorithm [41], [42] to overcome the non-convexity of the problem. For 2, 4, 8-queue cases, the problem is on average solved in 59.61s, 86.87s, 160.11s respectively on a server in our testbed (described in §V). Since the flow size distribution is collected over a long period of time,² the time for calculating the thresholds is negligible in comparison.

In summary, given flow arrival rate λ^3 and flow size distribution $F(\cdot)$, we can compute the thresholds relatively quickly, and enforce them at the end-hosts.

IV. IMPLEMENTATION AND TESTBED SETUP

A. PIAS Implementation

We have implemented a prototype of PIAS. We now describe each component of our prototype in detail.

1) *Packet Tagging:* The packet tagging module is responsible for maintaining per-flow state and marking packets with priority at end hosts. We implement it as a kernel module in Linux. The packet tagging module resides between the TCP/IP stack and Linux TC, which consists of three components: a NETFILTER [5] hook, a hash based flow table, and a packet modifier.

The operations are as follows: 1) the NETFILTER hook intercepts all outgoing packets using the LOCAL_OUT hook and directs them to the flow table. 2) Each flow in the flow table is identified by the 5-tuple: src/dst IPs, src/dst ports and protocol. When a packet comes in, we identify the flow it belongs to (or create a new entry) and increment the amount of bytes sent. 3) Based on the flow information, the packet modifier sets the packet's priority by modifying the DSCP field in the IP header to the corresponding value.

In certain cases, applications may build persistent TCP connections to keep delivering request-response short messages for a long time. In such scenarios, PIAS should provide message-level rather than connection-level prioritization. Otherwise, these persistent connections will eventually be assigned to the lowest priority due to the large cumulative size of bytes sent. To address this challenge, we monitor TCP send buffer instead of modifying applications [14] to get messages semantics. In our implementation, we add a JPROBE hook to TCP_SENDMSG function and record its per-flow called times. If the time gap since the last function call is larger than a threshold (500us in our experiments) and the TCP send

buffer is empty now, we regard this as the beginning of a new message and reset bytes sent back to 0. We find that this simple solution works well in the testbed experiments.

Offloading techniques like large segmentation offload (LSO) may degrade the accuracy of packet tagging. With LSO, the packet tagging module may not be able to set the right DSCP value for each individual MTU-sized packet within a large segment. To quantify this, we sample more than 230,000 TCP segments with payload data in our 1G testbed and find that the average segment size is only 7,220 Bytes. This introduces little impact on packet tagging. We attribute this to the small window size in DCN environment which has small bandwidth-delay product and large number of concurrent connections. We expect that the final implementation solution for packet tagging should be in NIC hardware to permanently avoid this interference.

To quantify system overhead introduced by the PIAS packet tagging module, we installed it on a Dell PowerEdge R320 server with a Intel 82599EB 10GbE NIC and measured CPU usage. LSO is enabled in this experiment. We started 8 TCP long flows and achieved ~ 9.4 Gbps goodput. The extra CPU usage introduced by PIAS is $< 1\%$ compared with the case where the PIAS packet tagging module is not enabled.

In the near future, link speed in data centers may reach 100Gbps. In such high speed networks, people may offload network stacks to hardware to reduce CPU overhead and processing latency. Despite the above software implementation, the PIAS packet tagging module can also be implemented in hardware, such as programmable NICs and FPGA.

2) *Switch Configuration:* We enforce strict priority queues at the switches and classify packets based on the DSCP field. Similar to previous work [11], [44], we use ECN marking based on the instant queue lengths with a single marking threshold. In addition to the switch queueing delay in the network, the sender's NIC also introduces latency because it is actually the first contention point of the fabric [13], [26]. Hardware and software queues at the end hosts can introduce large queueing delay, which might severely degrade the application performance [25], [45]. To solve this problem, our software solution hooks into the TX datapath at POST_ROUTING and rate-limits outgoing traffic at the line rate. Then, we perform ECN marking and priority queueing at the end host as well as the switches.

Per-Queue vs Per-Port ECN Marking: We observe that some of today's commodity switching chips offer multiple ways to configure ECN marking when configured to use multiple queues per port.

For example, our Broadcom BCM#56538-based switch allows us to enable either per-queue ECN marking or per-port ECN marking. In per-queue ECN marking, each queue has its own marking threshold and performs ECN marking independently to other queues. In per-port ECN marking, each port is assigned a single marking threshold and marks packets when the sum of all queue sizes belong to the port exceeds the marking threshold.

Per-queue ECN is widely used in many DCN transport protocols [11], [31], [39], however, we find it has limitation when supporting multiple queues. Each queue requires a

²The shortest period to collect enough flow information to form an accurate and reliable flow size distribution is an interesting research problem, which we leave as future work. Previous reported distributions [11], [16], [29] are all collected over periods of at least days.

³The traffic load used in §V can be calculated by $\lambda \cdot E[M]$, where $E[M]$ is the average flow size given $F(\cdot)$.




Fig. 3. Testbed Topology.

moderate ECN marking threshold h to fully utilize the link independently (e.g., $h = 20$ packets for 1G and 65 packets for 10G in DCTCP [11]). Thus, supporting multiple queues may require the shared memory be at least multiple times (e.g., 8) the marking threshold, which is not affordable for most shallow buffered commodity switches. For example, our Pronto-3295 switch has 4MB (≈ 2667 packets) memory shared by 384 queues (48x 1G ports with 8 queues per port). If we set $h = 20$ packets as suggested above, we need over 11MB memory in the worst case, otherwise when the traffic is bursty, the shallow buffer may overflow before ECN takes effect.

Per-port ECN, to the best of our knowledge, has rarely been exploited in recent DCN transport designs. Although per-port ECN marking cannot provide ideal isolation among queues as per-queue ECN marking, it can provide much better burst tolerance and support a larger number of queues in shallow buffered switches. Moreover, per-port ECN marking can potentially mitigate the starvation problem. It can push back high priority flows when many packets of low priority flows get queued in the switch. Therefore, we use per-port ECN marking.

3) *Rate Control*: We use Linux 3.18.11 kernel that supports DCTCP congestion control algorithm. We further observe an undesirable interaction between the DCTCP implementation and our switch. The DCTCP implementation does not set the ECN-capable (ECT) codepoint on TCP SYN packets and retransmitted packets, following the ECN standard [38]. However, our switch drops any non-ECT packets from ECN-enabled queues, when the instant queue length is larger than the ECN marking threshold. This problem severely degrades the TCP performance [44]. To address this problem, we set ECT on every TCP packet at the packet modifier.

B. Testbed Setup

We built a small testbed that consists of 16 servers connected to a Pronto 3295 48-port Gigabit Ethernet switch with 4MB shared memory, as shown in Figure 3. Our switch supports ECN and strict priority queuing with at most 8 class of service queues [1]. Each server is a Dell PowerEdge R320 with a 4-core Intel E5-1410 2.8GHz CPU, 8G memory, a 500GB hard disk, and a Broadcom BCM5719 NetXtreme Gigabit Ethernet NIC. Each server runs Debian 7.8-64bit with Linux 3.18.11 kernel. By default, advanced NIC offload mechanisms are enabled to reduce the CPU overhead. The base round-trip time (RTT) of our testbed is around 135 μ s.

In addition, we have also built a smaller 10G testbed for measuring the end host queuing delay in high speed network. We connect three servers to the same switch (Pronto 3295 has




Fig. 4. Traffic distributions used for evaluation.

four 10GbE ports). Each server is equipped with an Intel 82599EB 10GbE NIC.

V. EVALUATION

We evaluate PIAS using a combination of testbed experiments and large-scale ns-2 simulations. Our evaluation centers around four key questions:

- **How does PIAS perform in practice?** Using realistic workloads in our testbed experiments, we show that PIAS reduces the average FCT of short flows by ~ 29 -49% with the web search workload [11] and ~ 18 -34% with the data mining workload [22] compared to DCTCP. In an application benchmark with Memcached [6], we show that PIAS achieves ~ 28 -30% lower average query completion time than DCTCP.
- **How effective are individual design components of PIAS**, and how sensitive is PIAS to parameter settings? We show that PIAS achieves reasonable performance even with two queues. We also demonstrate that ECN is effective in mitigating the harmful effect of a mismatch between the demotion thresholds and traffic, but PIAS performs the best with the optimal threshold setting.
- **Does PIAS work well even in large data centers?** Using large-scale ns-2 simulations, we show that PIAS scales to multi-hop topologies and performs best among all information-agnostic schemes (DCTCP [11], L2DCT [32], and LAS [37]). PIAS shows a 1.1% performance (measured in average FCT) gap from pFabric [13], an idealized information-aware scheme, for short flows in the data mining workload.

A. Testbed Experiments

Setting: PIAS uses 8 priority queues by default and enables per-port ECN marking as discussed in §IV. Given the base RTT is only 135 μ s, we set the ECN marking threshold to 20KB. As previous work [11], [40] suggests, we set TCP RTOMin to 10ms. The TCP initial window is 10 packets. We choose the demotion thresholds according to the guidelines in §III-C.

We use two realistic workloads, a web search workload [11] and a data mining workload [22], based on measurement from production data centers. Their overall flow size distributions are shown in Figure 4. We also evaluate PIAS using an application benchmark with Memcached [6].

Results With Realistic Workloads: For this experiment, we develop a client/server model⁴ to generate realistic

⁴<https://github.com/HKUST-SING/TrafficGenerator>




Fig. 5. Overall average flow completion times for both workloads. (a) Web search workload. (b) Data mining workload.

workloads and measure the FCT on application layer. The client application, running on 1 machine, initially opens 5 persistent TCP connections to each of the rest 15 machines. During the experiment, the client application generates requests through available connections (if no connection is available, the client application will establish a new persistent connection) to fetch data based on a Poisson process. The request sizes are drawn from Figure 4. The server applications, running on 15 other machines, respond with requested data. Hence, a TCP connection can carry multiple flows (messages). We evaluate the performance of PIAS, DCTCP and TCP, while varying the network loads from 0.5 to 0.8. We run 10000 flows for each setting.

Figure 5 gives the overall average FCT for the web search workload and the data mining workload at various loads. In general, PIAS delivers the best performance. In the web search workload, the overall average FCT with PIAS is up to 6.5% lower compared to DCTCP and 10.2% lower compared to TCP. In the data mining workload, PIAS reduces the overall average FCT by up to 1.2% and 11.2% compared to DCTCP and TCP, respectively.

We further break down FCT across different flow sizes. Figure 6 and Figure 7 show the FCT across small (0,100KB] (a, b), medium (100KB,10MB] (c), and large (10MB, ∞) (d) flows, respectively; for the web search and data mining workloads, respectively.

We make the following three observations. First, PIAS achieves the best performance for small flows. Compared to DCTCP, PIAS reduces the average FCT of small flows by $\sim 29\text{--}49\%$ for web search workload and $\sim 18\text{--}34\%$ for data mining workload. The improvement of PIAS over DCTCP in the 99th percentile FCT of short flows is also obvious: $\sim 37\text{--}62\%$ for the web search workload and $\sim 15\text{--}31\%$ for the data mining workload. Second, PIAS also provides the best performance for medium flows. It achieves up to 20% lower average FCT of medium flows than DCTCP for the web search workload. Third, PIAS does not severely penalize large flows. For example, from Figure 6 (d) and Figure 7 (d), we can see that for the data mining workload PIAS is comparable or slightly better than TCP and DCTCP, while for the web search workload it is worse than DCTCP by up to 21%. This is expected because PIAS prioritizes short flows over long flows and $\sim 60\%$ of all bytes in the web search workload are from flows smaller than 10MB. Note that the performance degradation of large flows would not affect the overall average

FCT since data center workloads are dominated by small and medium flows.

Results With the Memcached Application: To assess how PIAS improves the performance of latency-sensitive applications, we build a Memcached [6] cluster with 16 machines. One machine is used as a client and the other 15 are used as servers to emulate a partition/aggregate soft-real-time service [11], [43]. We pre-populate server instances with 4B-key, 1KB-value pairs. The client sends a GET query to all 15 servers and each server responds with a 1KB value. A query is completed only when the client receives all the responses from the servers. We measure the query completion time as the application performance metric. Since a 1KB response can be transmitted within one RTT, the query completion time is mainly determined by the tail queueing delay. The base query completion time is around $650\mu\text{s}$ in our testbed. We also generate background traffic, a mix of mice flows and elephant flows following the distribution of the web search workload [11]. We use queries per second, or *qps*, to denote the application load. We vary the load of the background traffic from 0.5 to 0.8 and compare the performance of PIAS with that of DCTCP.

Figure 8 and Figure 9 show the results of the query completion time at 20 and 40 qps loads respectively. Since we enable both dynamic buffer management and ECN at the switch, none of queries suffers from TCP timeout. With the increase in background traffic load, the average query completion time of DCTCP also increases ($1016\text{--}1189\mu\text{s}$ at 40qps and $1014\text{--}1198\mu\text{s}$ at 20qps). By contrast, PIAS maintains a relatively stable performance. At 0.8 load, PIAS can achieve $\sim 28\text{--}30\%$ lower average query completion times than those of DCTCP. Moreover, PIAS also reduces the 99th percentile query completion time by $\sim 20\text{--}27\%$. In summary, PIAS can effectively improve the performance of the Memcached application by reducing the queueing delay of short flows.

End Host Queueing Delay: The above experiments mainly focus on network switching nodes. PIAS extends its switch design to the end hosts as the sender's NIC is actually the first contention point of the fabric [13], [26].

To quantify this, we conduct an experiment in our 10G setting with three servers (one sender and two receivers). We start several (1 to 8) long-lived TCP flows from the sender to a receiver. Then we measure RTT from the sender to the other receiver by sending ICMP ping packets. Without PIAS, ping packets could experience up to $6748\mu\text{s}$ queueing delay with 8 background flows. Then we deploy a 2-queue PIAS end host scheduling module (as described in §IV-A2) with a threshold of 100KB. Each ICMP packet is identified as a new flow by PIAS. We measure the RTTs with PIAS and compare them with the results without PIAS in Figure 10. In general, PIAS can significantly reduce the average RTT to $\sim 200\mu\text{s}$ and ensure that the 99th percentile RTT is smaller than $450\mu\text{s}$. Note that the PIAS scheduling module does not affect network utilization and large flows still maintain more than 9Gbps goodput during the experiment. Since we enable LSO to reduce CPU overhead, it is difficult for us to achieve fine-grained transmission control and some delay may still exist in NIC's transmission queues. We believe




Fig. 6. Web search workload: FCT across different flow sizes. TCP's performance is outside the plotted range of (b). (a) (0,100KB]: Avg. (b) (0,100KB]: 99th Percentile. (c) (100KB,10MB]: Avg. (d) (10MB,∞): Avg.




Fig. 7. Data mining workload: FCT across different flow sizes. TCP's performance is outside the plotted range of (a) and (b). (a) (0,100KB]: Avg. (b) (0,100KB]: 99th Percentile. (c) (100KB,10MB]: Avg. (d) (10MB,∞): Avg.




Fig. 8. Query completion times at 20 qps. (a) Mean. (b) 99th Percentile.




Fig. 10. RTT with background flows. (a) Mean. (b) 99th Percentile.




Fig. 9. Query completion times at 40 qps. (a) Mean. (b) 99th Percentile.

there is still room to improve by offloading the scheduling to NIC hardware [36].

B. PIAS Deep Dive

In this section, we conduct a series of targeted experiments to answer the following three questions:

- **How sensitive is PIAS to the number of queues available?** Network operators may reserve some queues

for other usage while some commodity switches [3] only support 2 priority queues. We find that, even with only 2 queues, PIAS still effectively reduces the FCT of short flows. However, in general, more queues can further improve PIAS's overall performance.

- **How effective is ECN in mitigating the mismatch?** ECN is integrated into PIAS to mitigate the mismatch between the demotion thresholds and traffic distribution. In an extreme mismatch scenario, we find that without ECN, PIAS's performance suffers with medium flows and is worse than DCTCP. However, with ECN, PIAS effectively mitigates this problem, and is better than, or at least comparable to DCTCP.
- **What is the effect of the optimal demotion thresholds?** Compared to PIAS with thresholds derived from simple heuristics, PIAS with the optimal demotion thresholds achieves up for $\sim 8.3\%$ improvement in medium flows, which improves the overall performance.

Impact of Number of Queues: In general, the more queues we use, the better we can segregate different flows, thus




Fig. 11. Web search workload with different numbers of queues. (a) (0,100KB]: Avg. (b) (100KB,10MB]: Avg.




Fig. 12. Web search workload with mismatch thresholds derived from data mining workload. (a) (0,100KB]: Avg. (b) (100KB,10MB]: Avg.

improving overall performance. For this experiment, we generate traffic using the web search workload and do the evaluation with 2, 4 and 8 priority queues. The results are shown in Figure 11. We observe that three schemes achieve the similar average FCT for short flows. As expected, the average FCT of medium flows improves with the increasing number of queues. For example, at 80% load, PIAS with 4 queues and 8 queues provide similar performance, but improve the FCT by 14.5% compared to 2 queues. The takeaway is that PIAS can effectively reduce FCT of short flows even with 2 queues and more queues can further improve PIAS's overall performance.

Effect of ECN Under Thresholds–Traffic Mismatch: We use the web search workload for evaluation while choosing the demotion thresholds derived from the data mining workload. We consider PIAS, PIAS without ECN, and DCTCP. Figure 12 shows the FCT results of short and medium flows. Both PIAS and PIAS without ECN greatly outperforms DCTCP for short flows. PIAS is slightly better than PIAS without ECN since it can efficiently reduce the switch buffer occupancy when many small flows collide in the highest priority queue. Furthermore, PIAS without ECN generally delivers the worst performance for medium flows. This is because, due to the threshold mismatch, medium and large flows coexist in low priority queues for long time. Without ECN, medium flows would experience large queuing delay and packet drops due to the impact of large flows. With ECN, PIAS effectively mitigates this side-effect by keeping low buffer occupancy as explained in §III-B2.

Impact of Demotion Thresholds: To explore the effectiveness of our optimal demotion threshold setting, we compare the optimized PIAS with the PIAS using thresholds derived

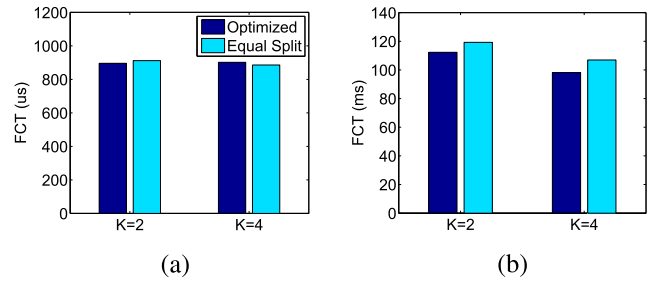


Fig. 13. Web search workload with different thresholds. (a) (0,100KB]: Avg. (b) (100KB,10MB]: Avg.

from the equal split heuristic as [13]. More specifically, given a flow size distribution and K queues, we set the first threshold to the size of $100/K$ th percentile flow, the second threshold to the size of $200/K$ th percentile flow, and so on. We run the web search workload at 80% load and summarize results in the Figure 13. We test PIAS with 2 and 4 queues. We observe that there is an obvious improvement in the average FCT of medium flows with the optimized thresholds. Specifically, PIAS (4-queue) with the optimized thresholds can achieve $\sim 8.3\%$ lower FCT for medium flows than that of equal split, and a $\sim 5.8\%$ improvement for the 2-queue PIAS. This partially validates the effectiveness of our optimal threshold setting.

C. Large-Scale NS-2 Simulations

We use ns-2 [9] simulations to answer four questions.

- **How does PIAS perform compared to information-agnostic schemes?** PIAS outperforms DCTCP [11] and L2DCT [32] in general, and significantly improves their average FCTs for short flows by up to 50%. Furthermore, PIAS is close to LAS for short flows and greatly outperforms LAS for long flows, reducing its average FCT by up to 55% for the web search workload.
- **How does PIAS perform compared to information-aware schemes?** As a practical information-agnostic scheme, PIAS can also deliver comparable performance to a clean-slate information-aware design, pFabric [13], in certain scenarios. For example, it only has a 1.1% gap to pFabric for short flows in the data mining workload.
- **How does PIAS perform in the oversubscribed network?** In a 3:1 oversubscribed network with ECMP load balancing, PIAS still delivers very good performance. Compared to DCTCP, the average FCT for the short flows with PIAS is up to $\sim 57\%$ lower for the web search workload.
- **How does PIAS perform in the heterogeneous workload?** We create a heterogeneous traffic pattern where different links have different flow size distributions. Under such challenging traffic pattern, PIAS still obviously outperforms DCTCP.

Setting: We use a leaf-spine topology with 9 leaf (ToR) switches, 4 spine (Core) switches and 144 hosts. Each leaf switch is connected to 16 hosts using 10Gbps links and 4 spine switches using 40Gbps links, thus forming a non-oversubscribed network. The base end-to-end round-trip time




Fig. 14. Overall average flow completion times. (a) Web search workload. (b) Data mining workload.




Fig. 15. Web search workload: FCT. (a) (0,100KB]: Avg. (b) (10MB,∞): Avg.

across the spine (4 hops) is $85.2\mu\text{s}$. We use packet spraying [19] for load balancing and disable dupACKs to avoid packet reordering. The TCP initial window is 70 packets, which is approximately equal to the bandwidth-delay product (BDP). The ECN marking threshold is 65 packets. For pFabric, each switch port has 210KB ($\sim 2 \times \text{BDP}$) buffer. For the other schemes, each switch port has 360KB buffer that is completely shared by all the queues. The TCP RTomin is $250\mu\text{s}$ ($\sim 3 \times \text{RTT}$) for pFabric and 2ms for the others. The traffic is generated similarly to our testbed experiments (§V – A). We open persistent connections between each pair of hosts. Flows arrive according to a Poisson process and their sources and destinations are chosen randomly among 144 hosts. Same as §V – A, the flow sizes are drawn from the web search and data mining distributions in Figure 4. We run 100000 flows for each simulation.

1) *Comparison With Information-Agnostic Schemes:* We mainly compare PIAS with three other information-agnostic schemes: DCTCP, L2DCT [32] and LAS [37] (pFabric using bytes sent as the priority).

Overall Performance: Figure 14 shows the average FCT of information-agnostic schemes under different workloads and load levels. From the figure, we see that PIAS delivers the best overall performance. By contrast, the performance of LAS is varied. PIAS significantly outperforms LAS by 37% (at 90% load) for the web search workload. This is because PIAS effectively mitigates the starvation between long flows unlike LAS. In the data mining workload, there are not so many large flows on the same link concurrently. As a result, LAS does not suffer from starvation as significantly.

Breakdown by Flow Size: We now breakdown the average FCT across different flow sizes: (0, 100KB] and (10MB, ∞) (Figure 15 and 16).




Fig. 16. Data mining workload: FCT. (a) (0,100KB]: Avg. (b) (10MB,∞): Avg.




Fig. 17. Average FCT for (0,100KB]. (a) Web search workload. (b) Data mining workload.

For short flows in (0,100KB], we find that PIAS significantly outperforms both DCTCP and L2DCT, improving the average FCT by up to $\sim 50\%$. This is because DCTCP and L2DCT use reactive rate control at the end host, which is not so effective as in-network prioritization used by PIAS. We further observe that PIAS achieves similar performance as LAS for short flows. PIAS only performs slightly worse than LAS for the web search workload when some packets from short flows get dropped.

For long flows in (10MB,∞), we find that PIAS significantly outperforms LAS for the web search workload (55% reduction in FCT at 90% load). This is because, in the web search workload, multiple large flows are likely to collide in the same link. In such scenarios, LAS always stops older flows to send new flows. This causes a serious starvation problem as large flows usually take very long times to finish. However, with PIAS, large flows receive their fair shares in the lowest priority queue, thus mitigating this problem. Furthermore, PIAS slightly outperforms DCTCP and L2DCT for both workloads. We note that L2DCT is marginally worse than DCTCP for large flows. We think this is due to L2DCT’s slower window increase in the congestion avoidance phase.

2) *Comparison With Ideal Information-Aware Schemes:* We compare PIAS to an ideal information-aware approach for DCN transport, pFabric [13], on small flows of the two workloads. We note that the most recent work PASE [31] can achieve better performance than pFabric in particular scenarios (*e.g.*, very high load and single rack). However in our topology setting with realistic workloads, pFabric is better than PASE and PDQ [24], and achieves near-optimal performance. Thus, we directly compare PIAS with pFabric.

The result is shown in Figure 17. In general, PIAS delivers comparable average FCT for short flows as pFabric,




Fig. 18. Web search workload in a 3:1 oversubscribed network with ECMP load balancing. (a) Overall: Avg. (b) (0,100KB]: Avg.

particularly within 1.1% in the data mining workload. We find that the gap between PIAS and pFabric is smaller in the data mining workload than that in the web search workload. This is mainly due to the fact that the data mining workload is more skewed than the web search workload. Around 82% flows in data mining are smaller than 100KB, while only 54% of flows in web search are smaller than 100KB. For the web search workload, it is more likely that large flows coexist with short flows in the high priority queues temporarily, increasing the queuing delay for short flows. pFabric, by assuming prior knowledge of flow sizes, is immune to such problem.

3) *Performance in Oversubscribed Networks*: We evaluate PIAS in a 3:1 oversubscribed network and employ ECMP for fabric load balancing. In this network, there are 12 leaf switches, 4 spine switches and 144 hosts. Each leaf switch is connected to 12 hosts and 4 spine switches with 10Gbps links. Given the source and destination of flows are chosen randomly, $\frac{11}{12}$ of all traffic traverses the spine. Hence, the load at the fabric's core is $\frac{11}{12} \times 3 = 2.75$ larger than the load at the edge. Figure 18 gives the results of PIAS and DCTCP for the web search workload. Note that the load in the figure is at the fabric's core. Compared to DCTCP, PIAS achieves up to $\sim 57\%$ and $\sim 12\%$ lower average FCT for short flows and all the flows, respectively.

4) *Performance With Heterogeneous Traffic*: In previous simulations, sizes of all flows are drawn from the same distribution. As a result, all the links should have the same flow size distribution. In this simulation, we create a heterogeneous traffic pattern instead. As we have 144 hosts, we have 144×143 host pairs in total. The sources and destinations of flows are chosen randomly. For flows from host i to j ,⁵ their sizes are drawn from either the web search workload if $i < j$ or the data mining workload. Consequently, the flow size distribution varies across different links. As PIAS uses a common set of demotion thresholds (based on traffic of entire DCN) for all end hosts, it cannot capture traffic variations in different links under such heterogeneous traffic pattern. We evaluate the performance of PIAS (using demotion thresholds derived for the web search workload) and DCTCP under such heterogeneous traffic pattern. As Figure 19 shows, compared to DCTCP, PIAS achieves up to $\sim 48\%$ and $\sim 8.1\%$ lower average FCT for short flows and all the flows, respectively. This demonstrates the robustness of PIAS in large-scale production data center networks.

⁵ i and j are host IDs.




Fig. 19. Heterogeneous traffic pattern: FCT. (a) Overall: Avg. (b) (0,100KB]: Avg.

VI. RELATED WORK

We classify prior work on minimizing FCT in DCNs into two categories: information-agnostic solutions (*e.g.*, [11], [12], [30], [32], [47]) and information-aware solutions (*e.g.*, [13], [24], [31]).

Information-agnostic solutions [11], [12], [30], [32], [47] generally improve the FCT for short flows by keeping low queue occupancy. For example, DCTCP [11] and DCQCN [47] try to keep the fabric queues small by employing ECN-based adaptive congestion control algorithms to throttle elephant flows. HULL [12] further improves the latency of DCTCP at the cost of trading network bandwidth. In summary, these solutions mainly perform end-host based rate control which is ineffective for flow scheduling. By contrast, PIAS leverages in-network priority queues to emulate SJF for flow scheduling, which is more efficient in terms of FCT minimization.

Information-aware solutions [13], [24], [31] attempt to approximate ideal Shortest Remaining Processing Time (SRPT) scheduling. For example, PDQ [24] employs switch arbitration and uses explicit rate control for flow scheduling. pFabric [13] decouples flow scheduling from rate control and achieves near-optimal FCT with decentralized in-network prioritization. PASE [31] synthesizes the strengths of previous solutions to provide good performance. In general, these solutions can potentially provide ideal performance, but they require non-trivial modifications on switch hardware or a complex control plane for arbitration. By contrast, PIAS does not touch the switch hardware or require any arbitration in the control plane, while still minimizing FCT.

There are also some other efforts [18], [39], [43] focusing on meeting flow deadlines. D3 [43] assigns rates to flows according to their sizes and deadlines explicitly, whereas D2TCP [39] and MCP [18] add deadline-awareness to ECN-based congestion window adjustment implicitly. They all require prior knowledge of flow information and do not directly minimize FCT, unlike PIAS.

To achieve diverse scheduling goals, some generic scheduling architectures [21], [34] have been proposed recently. In theory, we can also leverage these architectures to enforce MLFQ scheduling policy. In practice, it is challenging to deploy these architectures in large-scale production data centers. For example, Fastpass [34] uses a centralized scheduler, which suffers from failures and centralized scheduling overheads. pHost [21] requires the core of the underlying network


is congestion-free, which may not hold in practice [35]. By contrast, PIAS is easy to deploy in production data centers since it just leverages features of commodity switches without extra requirements.

VII. CONCLUSION

In this paper, we introduced PIAS, a solution that exploits existing commodity switches in DCNs to minimize the average FCT for flows, especially the smaller ones, without assuming any prior knowledge of flow sizes. We have implemented a PIAS prototype using all commodity hardware and evaluated PIAS through a series of small-scale testbed experiments as well as large-scale packet-level ns-2 simulations. Both our implementation and evaluation results demonstrate that PIAS is a viable solution that achieves all our design goals.

REFERENCES

- [1] *Pica8 OS*. [Online]. Available: <http://www.pica8.com/documents/pica8-datasheet-picos.pdf>
- [2] *Apache Storm*. [Online]. Available: <https://storm.incubator.apache.org/>
- [3] *Cisco Nexus 5500 Series NX-OS Quality of Service Configuration Guide*. [Online]. Available: <http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5500/sw/qos/7x/b5500QoSConfig7x.pdf>
- [4] *Hadoop*. [Online]. Available: <http://hadoop.apache.org/>
- [5] *Linux Netfilter*. [Online]. Available: <http://www.netfilter.org>
- [6] *Memcached*. [Online]. Available: <http://memcached.org/>
- [7] *Microsoft SQL Server*. [Online]. Available: <http://www.microsoft.com/en-us/server-cloud/products/sql-server/>
- [8] *OpenStack Object Storage*. [Online]. Available: <http://docs.openstack.org/api/openstack-object-storage/1.0/content/chunked-transfer-encoding.html>
- [9] *The Network Simulator NS-2*. [Online]. Available: <http://www.isi.edu/nsnam/ns/>
- [10] M. Alizadeh *et al.*, “CONGA: Distributed congestion-aware load balancing for datacenters,” in *Proc. SIGCOMM*, 2014, pp. 503–514.
- [11] M. Alizadeh *et al.*, “Data center TCP (DCTCP),” in *Proc. SIGCOMM*, 2010, pp. 63–74.
- [12] M. Alizadeh *et al.*, “Less is more: Trading a little bandwidth for ultra-low latency in the data center,” in *Proc. NSDI*, 2012, p. 19.
- [13] M. Alizadeh *et al.*, “pFabric: Minimal near-optimal datacenter transport,” in *Proc. SIGCOMM*, 2013, pp. 435–446.
- [14] H. Ballani *et al.*, “Enabling end-host network functions,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 45, no. 4, pp. 493–507, 2015.
- [15] T. Benson, A. Akella, and D. A. Maltz, “Network traffic characteristics of data centers in the wild,” in *Proc. IMC*, 2010, pp. 267–280.
- [16] T. Benson, A. Anand, A. Akella, and M. Zhang, “Understanding data center traffic characteristics,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 1, pp. 92–99, 2010.
- [17] J. Cao *et al.*, “Per-packet load-balanced, low-latency routing for close-based data center networks,” in *Proc. CoNEXT*, 2013, pp. 49–60.
- [18] L. Chen, S. Hu, K. Chen, H. Wu, and D. H. K. Tsang, “Towards minimal-delay deadline-driven data center TCP,” in *Proc. HotNets*, 2013, Art. no. 21.
- [19] A. Dixit, P. Prakash, Y. C. Hu, and R. R. Kompella, “On the impact of packet spraying in data center networks,” in *Proc. IEEE INFOCOM*, 2013, pp. 2130–2138.
- [20] R. Fielding *et al.*, *Hypertext Transfer Protocol-HTTP/1.1*, document RFC 2616, 1999.
- [21] P. X. Gao *et al.*, “pHost: Distributed near-optimal datacenter transport over commodity network fabric,” in *Proc. CoNEXT*, 2015, Art. no. 1.
- [22] A. Greenberg *et al.*, “VL2: A scalable and flexible data center network,” in *Proc. SIGCOMM*, 2009, pp. 51–62.
- [23] K. He *et al.*, “Presto: Edge-based load balancing for fast datacenter networks,” in *Proc. SIGCOMM*, 2015, pp. 465–478.
- [24] C.-Y. Hong, M. Caesar, and P. B. Godfrey, “Finishing flows quickly with preemptive scheduling,” in *Proc. SIGCOMM*, 2012, pp. 127–138.
- [25] K. Jang, J. Sherry, H. Ballani, and T. Moncaster, “Silo: Predictable message completion time in the cloud,” Univ. California, Berkeley, CA, USA, Tech. Rep. MSR-TR-2013-95, 2013.
- [26] V. Jeyakumar *et al.*, “EyeQ: Practical network performance isolation at the edge,” in *Proc. NSDI*, 2013, pp. 297–312.
- [27] E. Jones, T. Oliphant, and P. Peterson, “SciPy: Open source scientific tools for Python,” 2014.
- [28] B. Kalyanasundaram and K. R. Pruhs, “Minimizing flow time nonclairvoyantly,” *J. ACM*, vol. 50, no. 4, pp. 551–567, 2003.
- [29] S. Kandula, S. Sengupta, A. Greenberg, P. Patel, and R. Chaiken, “The nature of data center traffic: Measurements & analysis,” in *Proc. 9th ACM SIGCOMM Conf. Internet Meas. Conf.*, 2009, pp. 202–208.
- [30] R. Mittal *et al.*, “TIMELY: RTT-based congestion control for the datacenter,” in *Proc. SIGCOMM*, 2015, pp. 537–550.
- [31] A. Munir *et al.*, “Friends, not foes: Synthesizing existing transport strategies for data center networks,” in *Proc. SIGCOMM*, 2014, pp. 491–502.
- [32] A. Munir *et al.*, “Minimizing flow completion times in data centers,” in *Proc. IEEE INFOCOM*, 2013, pp. 2157–2165.
- [33] Y. Peng *et al.*, “Hadoopwatch: A first step towards comprehensive traffic forecasting in cloud computing,” in *Proc. IEEE INFOCOM*, 2014, pp. 19–27.
- [34] J. Perry, A. Ousterhout, H. Balakrishnan, D. Shah, and H. Fugal, “Fastpass: A centralized ‘zero-queue’ datacenter network,” in *Proc. SIGCOMM*, 2014, pp. 307–318.
- [35] L. Popa *et al.*, “ElasticSwitch: Practical work-conserving bandwidth guarantees for cloud computing,” in *Proc. SIGCOMM*, 2013, pp. 351–362.
- [36] S. Radhakrishnan *et al.*, “SENIC: Scalable NIC for end-host rate limiting,” in *Proc. NSDI*, vol. 14, 2014, pp. 475–488.
- [37] I. A. Rai, G. Urvoy-Keller, M. K. Vernon, and E. W. Biersack, “Performance analysis of LAS-based scheduling disciplines in a packet switched network,” in *Proc. SIGMETRICS*, 2004, pp. 106–117.
- [38] K. Ramakrishnan, S. Floyd, and D. Black, *The Addition of Explicit Congestion Notification (ECN) to IP*, document RFC 3168, 2001.
- [39] B. Vamanan, J. Hasan, and T. Vijaykumar, “Deadline-aware datacenter TCP (D2TCP),” in *Proc. SIGCOMM*, 2012, pp. 115–126.
- [40] V. Vasudevan *et al.*, “Safe and effective fine-grained TCP retransmissions for datacenter communication,” in *Proc. SIGCOMM*, 2009, pp. 303–314.
- [41] D. J. Wales and J. P. Doye, “Global optimization by basin-hopping and the lowest energy structures of Lennard-Jones clusters containing up to 110 atoms,” *J. Phys. Chem. A*, vol. 101, no. 28, pp. 5111–5116, 1997.
- [42] D. J. Wales and H. A. Scheraga, “Global optimization of clusters, crystals, and biomolecules,” *Science*, vol. 285, no. 5432, pp. 1368–1372, 1999.
- [43] C. Wilson, H. Ballani, T. Karagiannis, and A. Rowtron, “Better never than late: Meeting deadlines in datacenter networks,” in *Proc. SIGCOMM*, 2011, pp. 50–61.
- [44] H. Wu *et al.*, “Tuning ECN for data center networks,” in *Proc. CoNEXT*, 2012, pp. 25–36.
- [45] Y. Xu, M. Bailey, B. Noble, and F. Jahanian, “Small is better: Avoiding latency traps in virtualized data centers,” in *Proc. SOCC*, 2013, Art. no. 7.
- [46] M. Yu *et al.*, “Profiling network performance for multi-tier data center applications,” in *Proc. NSDI*, 2011, pp. 57–70.
- [47] Y. Zhu *et al.*, “Congestion control for large-scale RDMA deployments,” in *Proc. SIGCOMM*, 2015, pp. 523–536.



Wei Bai received the B.E. degree in information security from Shanghai Jiao Tong University, China, in 2013. He is currently pursuing the Ph.D. degree in computer science with The Hong Kong University of Science and Technology. His current research interests are in the area of data center networks.



Li Chen received the B.E. degree (Hons.) in electronic and computer engineering with a minor in mathematics and the M.Phil. degree from The Hong Kong University of Science and Technology (HKUST) in 2011 and 2013, respectively. He is currently pursuing the Ph.D. degree with the Department of Computer Science and Engineering, HKUST. He is involved in topics in data center networking, under the supervision of Prof. K. Chen. He is currently a Microsoft Research Asia Ph.D. Fellow.



Kai Chen received the Ph.D. degree in computer science from Northwestern University, Evanston, IL, USA, in 2012. He is currently an Assistant Professor with the Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Hong Kong. His research interest includes networked systems design and implementation, data center networks, and cloud computing.



Chen Tian received the B.S., M.S., and Ph.D. degrees from the Department of Electronics and Information Engineering, Huazhong University of Science and Technology, China, in 2000, 2003, and 2008, respectively. He was an Associate Professor with the School of Electronics Information and Communications, Huazhong University of Science and Technology. From 2012 to 2013, he was a Post-Doctoral Researcher with the Department of Computer Science, Yale University. He is currently an Associate Professor with the State Key Laboratory for Novel Software Technology, Nanjing University, China. His research interests include data center networks, network function virtualization, distributed systems, Internet streaming and urban computing.



Dongsu Han received the B.S. degree in computer science from the Korea Advanced Institute of Science (KAIST) in 2003, and the Ph.D. degree in computer science from Carnegie Mellon University in 2012. He is currently an Assistant Professor with the School of Electrical Engineering, Graduate School of Information Security, KAIST. He is interested in networking, distributed systems, and network/system security.



Hao Wang received the B.E. and M.E. degrees from Shanghai Jiao Tong University, in 2012 and 2015, respectively. He is currently pursuing the Ph.D. degree with the University of Toronto. He was a Research Assistant with The Hong Kong University of Science and Technology for 14 months. His research interests include datacenter networking, distributed computing, and software defined networking.