



# **Your Coflow has Many Flows: Sampling them for Fun and Speed**

Akshay Jajoo, Y. Charlie Hu, and Xiaojun Lin, *Purdue University*

<https://www.usenix.org/conference/atc19/presentation/jajoo>

**This paper is included in the Proceedings of the  
2019 USENIX Annual Technical Conference.**

**July 10–12, 2019 • Renton, WA, USA**

ISBN 978-1-939133-03-8

**Open access to the Proceedings of the  
2019 USENIX Annual Technical Conference  
is sponsored by USENIX.**

# Your Coflow Has Many Flows: Sampling Them for Fun and Speed

Akshay Jajoo  
ajajoo@purdue.edu

Y. Charlie Hu  
ychu@purdue.edu  
Purdue University

Xiaojun Lin  
linx@purdue.edu

## Abstract

Coflow scheduling improves data-intensive application performance by improving their networking performance. State-of-the-art online coflow schedulers in essence approximate the classic Shortest-Job-First (SJF) scheduling by learning the coflow size online. In particular, they use multiple priority queues to simultaneously accomplish two goals: to sieve long coflows from short coflows, and to schedule short coflows with high priorities. Such a mechanism pays high overhead in learning the coflow size: moving a large coflow across the queues delays small and other large coflows, and moving similar-sized coflows across the queues results in inadvertent round-robin scheduling.

We propose PHILAE, a new online coflow scheduler that exploits the spatial dimension of coflows, *i.e.*, a coflow has many flows, to drastically reduce the overhead of coflow size learning. PHILAE pre-schedules sampled flows of each coflow and uses their sizes to estimate the average flow size of the coflow. It then resorts to Shortest Coflow First, where the notion of shortest is determined using the learned coflow sizes and coflow contention. We show that the sampling-based learning is robust to flow size skew and has the added benefit of much improved scalability from reduced coordinator-local agent interactions. Our evaluation using an Azure testbed, a publicly available production cluster trace from Facebook shows that compared to the prior art Aalo, PHILAE reduces the coflow completion time (CCT) in average (P90) cases by  $1.50\times$  ( $8.00\times$ ) on a 150-node testbed and  $2.72\times$  ( $9.78\times$ ) on a 900-node testbed. Evaluation using additional traces further demonstrates PHILAE's robustness to flow size skew.

## 1 Introduction

### 1.1 Motivation

In big data analytics jobs, speeding up the communication stage where the data is transferred between compute nodes is important to speed up the jobs. However, improving network level metrics such as flow completion time may not translate into improvements at the application level metrics such as job completion time. The coflow abstraction [18] was proposed to bridge such a gap. The abstraction captures the collective network requirements of applications, as reduced coflow completion time (CCT) can directly lead to faster job completion time [20, 24].

There have been a number of efforts on network designs for coflows [7, 21, 27] that assume complete prior knowledge of coflow sizes (The coflow size is defined as the total size of its constituent flows.). However, in many practical settings, coflow characteristics are not known a priori. For example, multi-stage jobs pipeline data from one stage to the next as soon as the data is generated, which makes it difficult to know the size of each flow [22, 40]. A recent study [40] shows various other reasons why it is not very plausible to learn flow sizes from applications, for example, learning flow sizes from applications requires changing either the network stack or the applications.

Scheduling coflows in such *non-clairvoyant* settings, however, is challenging. The major challenge in developing an effective non-clairvoyant coflow scheduling scheme has centered around how to learn the coflow sizes online quickly and accurately, as once the coflow sizes (bytes to be transferred) can be estimated, one can apply variations of the classic Shortest-Job-First (SJF) algorithm such as Shortest Coflow First [21] or apply an LP solver (*e.g.*, [7]).

State-of-the-art online non-clairvoyant schedulers such as Saath [30], Gravtion [29] and Aalo [19] in essence learn coflow sizes and approximate SJF using discrete priority queues, where all newly arriving coflows start from the highest priority queue, and move to lower priority queue as they send more data (without finishing), *i.e.*, cross the per-queue thresholds. In this way, the smaller coflows finish in high priority queues, while the larger coflows gradually move to the lower priority queues where they finish after smaller coflows.

To realize the above idea in scheduling coflows which have flows at many network ports, *i.e.*, in a distributed setting, Aalo uses a global coordinator to assign coflows to logical priority queues, and uses the total bytes sent by all flows of a coflow as its logical "length" in moving coflows across the queues. The logical priority queues are mapped to local priority queues at each port, and the individual local ports then schedule the flows in its local priority queues, *e.g.*, by enumerating flows from the highest to lowest priority queues and using FIFO to order the flows within each queue.

In essence, Aalo learns coflow sizes by actually scheduling the coflow, a "try and miss" approach to approximate SJF. As coflow sizes are not known, in each queue, Aalo schedules each coflow for a fixed amount of data (try). If the coflow does not finish (miss), it is demoted to a lower priority queue.

Afterwards, such a coflow will no longer block coflows in higher priority queues.

Using multiple priority queues to learn the relative coflow sizes of coflows this way, however, negatively affects the average CCT and the scalability of the coordinator:

**(1) Intrinsic queue-transit overhead:** Every coflow that Aalo transits through the queues before reaching its final queue worsens the average CCT because during transitions, such a coflow effectively *blocks other shorter* coflows in the earlier queues it went through, which would have been scheduled before this coflow starts in a perfect SJF.

**(2) Overhead due to inadvertent round-robin:** Although Aalo attempts to approximate SJF, it inadvertently ends up doing *round-robin* for coflows of similar sizes as it moves them across queues. Aalo assigns a *fixed threshold of data transfer* for each coflow in each queue. Assume there are “ $N$ ” coflows in a queue that do not finish in that queue. Aalo schedules one coflow (chosen using FIFO) and demotes it to a lower priority queue when the coflow reaches the data threshold. At that point, the next coflow from the same queue is scheduled, which joins the previous coflow at a lower priority queue after exhausting its quantum, and this cycle continues as coflows of similar sizes move through the queues. Effectively, these coflows experience the round-robin scheduling which is known to have the worst average CCT [39], when jobs are of similar sizes.

**(3) Limited scalability from frequent updates from local ports:** To support the try-and-error style learning, the coordinator requires frequent updates from all local ports of the bytes sent for each coflow in order to move coflows across multiple queues timely. This results in high load on the central coordinator from receiving frequent updates and calculating and sending new rate allocations, which limits the scalability of the overall approach.

**Empirical measurement** We quantify the coflow size *learning overhead* of Aalo, defined as the portion of the bytes of a coflow that has been transferred (or the fraction of its CCT spent in doing so) before reaching its correct queue, using a trace from Facebook clusters [4] (see detailed methodology in §8). Figure 1 shows that 40% of the coflows that moved beyond the initial queue reached the correct priority queue after spending more than 20% of their CCT moving across early queues.

## 1.2 Our Contribution

We propose PHILAE, a new non-clairvoyant coflow scheduler with a dramatically different approach to learning coflow sizes to enable online SJF. To leverage optimal scheduling SJF in coflow scheduling, it is vital to learn the coflow sizes quickly and accurately. PHILAE achieves this objective by exploiting the *spatial dimension* of coflows, *i.e.*, a coflow typically consists of many flows, via *sampling*, a highly effective technique used in large-scale surveys [34]. In particular, PHILAE pre-schedules sampled flows, called *pilot flows*, of each




Figure 1: CDF of learning overhead per coflow, *i.e.*, the time to reach the correct priority queue as a fraction of CCT, excluding coflows directly scheduled by PHILAE or finish in Aalo’s first queue.

coflow and uses their measured size to estimate the coflow size. It then resorts to SCF using the estimated job size.

Intuitively, such a sampling scheme avoids all three sources of overhead in Aalo – Once the coflow sizes are learned, the coflows are assigned to the correct queues, which avoids the intrinsic queue-transit and round-robin effects. Further, a sampling-based design has an important benefit – it offers much higher scalability than priority-queue-based learning in Aalo. This is because unlike Aalo, after estimating the coflow size, PHILAE clients do not need to send periodic updates of bytes sent-so-far to the centralized coordinator.

Developing a complete non-clairvoyant coflow scheduler based on the simple sampling idea raises three questions: (1) Why is sampling more efficient than the priority-queue-based coflow size learning? (2) Will sampling be effective in the presence of skew of flow sizes? (3) How to design the complete scheduler architecture? We systematically address these questions with design rational, theoretical analysis, system design, prototyping, and extensive evaluation.

In summary, this paper makes the following contributions:

- (1) Using a production datacenter trace from Facebook, we show that the prior art scheduler Aalo spends substantial amount of time and network bandwidth in learning coflow sizes, which negatively affects the CCT of coflows.
- (2) We propose the novel idea of applying sampling in the spatial dimension of coflow to significantly reduce the overhead of online learning coflow sizes.
- (3) We present theoretical underpinning explaining why sampling remains effective in the presence of flow size skew.
- (4) We present the design and implementation of PHILAE.
- (5) We extensively evaluate PHILAE via simulations and testbed experiments, and show that compared to the prior art, the new design reduces the average CCT by  $1.51\times$  for the Facebook coflow trace and by  $1.36\times$  for a trace with properties similar to a Microsoft production cluster.
- (6) The CCT improvement mainly stems from reduced coflow size learning overhead. PHILAE reduces the median latency and data sent in finding the right queue for coflows in Aalo by  $19.0\times$  and  $20.0\times$ , respectively (§8.2).

## 2 Background and Problem Statement

We start with a brief review of the coflow abstraction and the need for non-clairvoyant coflow scheduling. We then state the network model and problem formulation.

**Coflow abstraction** In data-parallel applications such as Hadoop [1] and Spark [2], the job completion time heavily depends on the completion time of the communication stage [12, 20]. The coflow abstraction [18] was proposed to speed up the communication stage to improve application performance. A coflow is defined as a set of flows between several nodes that accomplish a common task. For example, in map-reduce jobs, the set of all flows from all map to all reduce tasks in a single job forms a typical coflow. The coflow completion time (CCT) is defined as the time duration between when the first flow arrives and the last flow completes. In such applications, improving CCT is more important than improving individual flows' completion time (FCT) for improving the application performance [19, 21, 24, 29, 30].

**Non-clairvoyant coflows** Data-parallel directed acyclic graphs (DAGs) typically have multiple stages which are represented as multiple coflows with dependencies between them. Recent systems (e.g., [3, 22, 28, 36]) employ optimizations that pipeline the consecutive computation stages which removes the barrier at the end of each coflow, making knowing flow sizes of each coflow beforehand difficult. Thus in this paper, we focus on *non-clairvoyant* coflow scheduling which do not assume knowledge about coflow characteristics such as flow sizes upon coflow arrival.

**Non-blocking network fabric** We assume the same non-blocking network fabric model in recent network designs for coflows [7, 19, 21, 29, 30], where the datacenter network fabric is abstracted as a single non-blocking switch that interconnects all the servers, and each server (computing node) is abstracted as a network port that sends and receives flows. In such a model, the ports, i.e., server uplinks and downlinks, are the only source of contention as the network core is assumed to be able to sustain all traffic injected into the network. We note that the abstraction is to simplify our description and analysis, and is not required or enforced in our evaluation.

**Problem statement** Our goal is to *develop an efficient non-clairvoyant coflow scheduler that optimizes the communication performance, in particular the average CCT, of data-intensive applications without prior knowledge, while guaranteeing starvation freedom and work conservation and being resilient to the network dynamics.* The problem of non-clairvoyant coflow scheduling is NP-hard because coflow scheduling even assuming all coflows arrive at time 0 and their size are known in advance is already NP-hard [21]. Thus practical non-clairvoyant coflow schedulers are approximation algorithms. Our approach is to dynamically prioritize coflows by efficiently learning their flow sizes online.

## 3 Key Idea

Our new non-clairvoyant coflow scheduler design, PHILAE, is based on a key observation about coflows that a coflow has a *spatial dimension*, i.e., it typically consists of many flows. We thus propose to explicitly learn coflow sizes online by using *sampling*, a highly effective technique used in large-scale surveys [34]. In particular, PHILAE preschedules sampled flows, called *pilot flows*, of each coflow and uses their measured sizes to estimate the coflow size. It then resorts to SJF or variations using the estimated coflow sizes.

Developing a complete non-clairvoyant coflow scheduler based on the simple sampling idea raises three questions:

(1) *Why is sampling more efficient than the priority-queue-based coflow size learning? Would scheduling the remaining flows after sampled pilot flows are completed adversely affect the coflow completion time?*

(2) *Will sampling be effective in the presence of skew of flow sizes?*

(3) *How to design the complete scheduler architecture?*

We answer the first two questions below, and present the complete architecture design in §4.

### 3.1 Why is sampling more efficient?

Scheduling pilot flows first before the rest of the flows can potentially incur two sources of overhead. First, scheduling pilot flows of a newly arriving coflow consumes port bandwidth which can delay other coflows (with already estimated sizes). However, compared to the multi-queue based approach, the overhead is much smaller for two reasons: (1) PHILAE schedules only a small subset of the flows (e.g., fewer than 1% for coflows with many flows). (2) Since the CCT of a coflow depends on the completion of its last flow, some of its earlier finishing flows could be delayed without affecting the CCT. PHILAE exploits this observation and schedules pilot flows on the least-busy ports to increase the odds that it only affects earlier finishing flows of other coflows.

Second, scheduling pilot flows first may elongate the CCT of the newly arriving coflow itself whose other flows cannot start until the pilot flows finish. This is again typically insignificant for two reasons: (1) A coflow (e.g., from a MapReduce job) typically consists of flows from all sending ports to all receiving ports. Conceptually, pre-scheduling one out of multiple flows from each sender may not delay the coflow progress at that port, because all flows at that port have to be sent anyway. (2) Coflow scheduling is of high relevance in a busy cluster (when there is a backlog of coflows in the network), in which case the CCT of coflow is expected to be much higher than if it were the only coflow in the network, and hence the piloting overhead is further dwarfed by a coflow's actual CCT.

### 3.2 Why is sampling effective in the presence of skew?

The flow sizes within a coflow may vary (*skew*). Intuitively, if the skew across flow sizes is small, sampling even a small number of pilot flows will be sufficient to yield an accurate estimate. Interestingly, even if the skew across flow sizes is large, our experiment indicates that sampling is still highly effective. In the following, we give both the intuition and theoretical underpinning for why sampling is effective.

Consider, for example, two coflows and the simple setting where both coflows share the same set of ports. In order to improve the average CCT, we wish to schedule the shorter coflow ahead of the longer coflow. If the total sizes of the two coflows are very different, then even a moderate amount of estimation error of the coflow sizes will not alter their ordering. On the other hand, if the total sizes of the two coflows are close to each other, then indeed the estimation errors will likely alter their ordering. However, in this case since their sizes are not very different anyway, switching the order of these two coflows will not significantly affect the average CCT.

**Analytic results.** To illustrate the above effect, we show that the gap between the CCT based on sampling and assuming perfect knowledge is small, even under general flow size distributions. Specifically, coflows  $C_1$  and  $C_2$  have  $cn_1$  and  $cn_2$  flows, respectively. Here, we assume that  $n_1$  and  $n_2$  are fixed constants. Thus, by taking  $c$  to be larger, we will be able to consider wider coflows. Assume that each flow of  $C_1$  (correspondingly,  $C_2$ ) has a size that is distributed within a bounded interval  $[a_1, b_1]$  ( $[a_2, b_2]$ ) with mean  $\mu_1$  ( $\mu_2$ ), *i.i.d.* across flows. However, the exact distributions can be arbitrary. Let  $T^c$  be the total completion time when the exact flow sizes are known in advance. Let  $\tilde{T}^c$  be the average CCT by sampling  $m_1$  and  $m_2$  flows from  $C_1$  and  $C_2$ , respectively. Without loss of generality, we assume that  $n_2\mu_2 \geq n_1\mu_1$ . Then, using Hoeffding's Inequality, we can show that,

$$\lim_{c \rightarrow \infty} \frac{\tilde{T}^c - T^c}{T^c} \leq 4 \exp \left[ - \frac{2(n_2\mu_2 - n_1\mu_1)^2}{\left( \frac{n_2(b_2 - a_2)}{\sqrt{m_2}} + \frac{n_1(b_1 - a_1)}{\sqrt{m_1}} \right)^2} \right] \frac{n_2\mu_2 - n_1\mu_1}{n_2\mu_2 + 2n_1\mu_1} \quad (1)$$

(Note that here we have used the fact that, since both coflows share the same set of ports and  $c$  is large, the CCT is asymptotically proportional to the coflow size.)

Equation (1) can be interpreted as follows. First, due to the first exponential term, the relative gap between  $\tilde{T}^c$  and  $T^c$  decreases as  $b_1 - a_1$  and  $b_2 - a_2$  decrease. In other words, as the skew of each coflow decreases, sampling becomes more effective. Second, when  $b_1 - a_1$  and  $b_2 - a_2$  are fixed, if  $n_2\mu_2 - n_1\mu_1$  is large (i.e., the two coflow sizes are very different), the value of the exponential function will be small. On the other hand, if  $n_2\mu_2 - n_1\mu_1$  is close to zero (i.e., the two coflow sizes are close to each other), the numerator on

the second term on the right hand side will be small. In both cases, the relative gap between  $\tilde{T}^c$  and  $T^c$  will also be small, which is consistent with the intuition explained earlier. The largest gap occurs when  $n_2\mu_2 - n_1\mu_1$  is on the same order as  $\frac{n_2(b_2 - a_2)}{\sqrt{m_2}} + \frac{n_1(b_1 - a_1)}{\sqrt{m_1}}$ . Finally, although these analytical results assume that both coflows share the same set of ports, similar conclusions on the impact of estimation errors due to sampling also apply under more general settings.

The above analytical results suggest that, when  $c$  is large, the relative performance gap for CCT is a function of the number of pilot flows sampled for each coflow, but is independent of the total number of flows in each coflow. In practice, large coflows will dominate the total CCT in the system. Thus, these results partly explain that, while in our experiments the number of pilot flows is never larger than 1% of the total number of flows, the performance of our proposed approach is already very good.

Finally, the above analytical results do not directly tell us how to choose the number of pilot flows, which likely depends on the probability distribution of the flow size. In practice, we do not know such distribution ahead of time. Further, while choosing a larger number of pilot flows reduces the estimation errors, it also incurs higher overhead and delay. Therefore, our design (§4) needs to have practical solutions that carefully address these issues.

## 4 PHILAE Design

In this section, we present the detailed design of PHILAE, which addresses three design challenges: (1) *Coflow size estimation*: How to choose and schedule the pilot flows for each newly arriving coflow? (2) *Starvation avoidance*: How to schedule coflows after size estimation using variations of SJF that avoid starvation? (3) *Coflow scheduling*: How to schedule among all the coflows with estimated sizes?

### 4.1 PHILAE architecture

Fig. 2 shows the PHILAE architecture. PHILAE models the entire datacenter as a single big-switch with each computing node as an individual port. The scheduling task in PHILAE is divided among (1) a central coordinator, and (2) local agents that run on individual ports. A computing framework such as Spark [42] first registers (removes) a coflow when a job arrives (finishes). Upon a new coflow arrival, old coflow completion, or pilot flow completion, the coordinator calculates a new coflow schedule, which includes (1) coflows that are to be scheduled in the next time slot, and (2) flow rates for the individual flows of a coflow, and pushes this information to the local agents which use this information to allocate their bandwidth. The local agents will follow the current schedule until they receive a new schedule.




Figure 2: PHILAE architecture.

## 4.2 Sampling pilot flows

As discussed in §3, PHILAE estimates the size of a coflow online by actually scheduling a subset of its flows (*pilot flows*) at their ports. We do not schedule the flows of a coflow other than the pilot flows until the completion of the pilot flows in order to avoid unnecessary extra blocking of other potentially shorter coflows.

**How many pilot flows?** When a new coflow arrives, PHILAE first needs to determine the number of pilot flows. As discussed at the end of §3, the number of pilot flows affects the trade-off between the coflow size estimation accuracy and scheduling overhead. For coflows with skewed flow sizes, accurately estimating the total coflow size potentially requires sampling the sizes of many pilot flows. However, scheduling pilot flows has associated overhead, *i.e.*, if the coflow turns out to be a large coflow and should have been scheduled to run later under SJF.

We explore several design options for choosing the number of pilot flow. Two natural design choices are using a constant number of pilot flows or a fixed fraction of the total number of flows of a coflow. In addition, we observe that typical coflows consist of flows between a set of senders (*e.g.*, mappers) and a set of receivers (*e.g.*, reducers) [23]. We thus include a third design choice of a fixed fraction of sending ports. This design also spreads the pilot flows to avoid having multiple pilot flows contending for the same sending ports. We empirically found that (§8.2) limiting the pilot flows to 5% to 10% of the number of its sending ports (*e.g.*, mappers in a MapReduce coflow) strikes a good balance between estimation accuracy and overhead. We note the total number of flows sampled in this case is still under 1%.

Finally, we estimate the total coflow size as  $S = f_i \cdot N$ , where  $N$  is the number of flows in a coflow, and  $f_i$  is the average size of the sampled pilot flows.

**Which flows to probe?** Second, PHILAE needs to decide which ports to schedule the chosen number of probe flows for a coflow. For this, we use a simple heuristic where, upon the arrival of a new coflow, we select the ports for its pilot

flows that are least busy, *i.e.*, having pilot flows from the least number of other coflows. PHILAE starts with the least busy sending port and iterates over receiving ports starting with the least busy receiving port and assigns the flow if it exists. It then updates the statistics for the number of pilot flows scheduled at each port and repeats the above process. Such a choice will likely delay fewer coflows when the pilot flows are scheduled and hence reduce the elongation on their CCT. We note that such an online heuristic may not be optimal; more sophisticated algorithms can be derived by picking ports for multiple coflows together. However, we make this design choice for its simplicity and low time complexity to ensure that the coordinator makes fast decisions.

**How to schedule pilot flows?** In PHILAE, we prioritize the pilot flows of a new coflow over existing flows to accelerate learning the size of the new coflow. In particular, at each port, pilot flows have high priority over non-pilot flows. If there are multiple outstanding pilot flows (of different coflows) at a port, PHILAE schedules them in the FIFO order.

## 4.3 Coflow scheduling with starvation avoidance

Once the sizes of coflows are learned, we can apply variations of the SJF policy to schedule them. However, it is well known that such policies can lead to starvation.

There are many ways to mitigate the starvation issue. However, a subtlety arises where *even slight difference in how starvation is addressed can result in different performance*. For example, the multiple priority queues in Aalo has the benefit of ensuring progress of all coflows, but assigning different time-quanta to different priority queues can result in different average CCT for the same workload. To ensure the fairness of performance comparison with Aalo, we need to ensure that both PHILAE and Aalo provide the same level of starvation freedom (or progress measure).

For this reason, in this paper, we inherit the multiple priority queue structure from Aalo for coflow scheduling. As in Aalo, PHILAE sorts the coflows among multiple priority queues. In particular, PHILAE uses  $N$  queues,  $Q_0$  to  $Q_{N-1}$ , with each queue having lower queue threshold  $Q_q^{lo}$  and higher threshold  $Q_q^{hi}$ , where  $Q_0^{lo} = 0$ ,  $Q_{N-1}^{hi} = \infty$ ,  $Q_{q+1}^{lo} = Q_q^{hi}$ , and the queue thresholds grow exponentially, *i.e.*,  $Q_{q+1}^{hi} = E \cdot Q_q^{hi}$ .

The overall coflow scheduling in PHILAE works as follows. After the coflow size is estimated using pilot flows, PHILAE assigns the coflow to the priority queue using inter-coflow policies discussed in §4.4. Within a queue, we use FIFO to schedule coflows. Lastly, we use weighted sharing of network bandwidth among the queues, where a priority queue receives a network bandwidth based on its priority. As in Aalo, the weights decrease exponentially with decrease in the priority of the queues.

Using FIFO within the priority queue and weighted fair sharing among the queues together ensure the same starva-

tion freedom and thus meaningful performance comparison between PHILAE and Aalo [19].

#### 4.4 Inter-coflow scheduling policies

In PHILAE, we explore four different scheduling policies based on different combinations of coflow size and contention, two size-based policies (*A*, *B*) as in Aalo, a contention-based, similar to the intra-queue policy used in Saath [30] (*C*), and a new contention-and-length-based policy (*D*):

**(A) Smallest job first:** Coflows are sorted based on coflow size ( $l \cdot n$ ).

**(B) Smallest remaining data first:** Coflows are sorted based on remaining data ( $l \cdot n - d$ ).

**(C) Least contention first:** Coflows are sorted based on their contention ( $c$ ).

**(D) Least length-weighted total-port contention first:** Coflows are sorted based on the sum of port-wise contention times estimated flow length  $\sum_p c^p \cdot l$ .

We use the following parameters of a coflow to define the metrics in scheduling algorithms: (1) average flow length ( $l$ ) from piloting, (2) number of flows ( $n$ ), (3) number of sender and receiver ports ( $s, r$ ), (4) total amount of data sent so far ( $d$ ), (5) contention ( $c$ ), defined as the number of other coflows sharing any ports with the given coflow, and (6) port-wise contention ( $c^p$ ), defined as the number of other coflows blocked at a given port  $p$ .

PHILAE uses Policy D by default, as it results in the least average CCT (§8). For all policies, we continue to use the priority-queue based scheduling, and the algorithms only differ in what metric they use in assigning coflows to the priority queues. In contrast, Aalo does not handle inter-coflow contention, and uses the total bytes sent so far ( $d$ ) to move coflows across multiple priority queues.

#### 4.5 Rate allocation

Once the scheduling order of the coflows is determined, we need to determine the rates for the individual flows at each port. First, since we want to quickly finish the pilot flow, at any port that has pilot flows, PHILAE assigns the *entire* port bandwidth to the pilot flows. For the remaining ports, as discussed in §4.3, across multiple queues, PHILAE assigns weighted shares of the port bandwidth, by assigning them varying numbers of scheduling intervals according to the weights assigned to each priority queues.

Second, at each scheduling interval, PHILAE assigns rates for the flows of the coflow in the head of the FIFO queue as follows. It assigns equal rates at all the ports containing its flows as there is no benefit in speeding-up its flows at certain ports when its CCT depends on the slowest flow. At each port, we could use max-min fairness to schedule the individual flows of the coflow (to different receivers), and then assign the rate of the slowest flow to all the flows in the coflow. Afterwards, the port-allocated bandwidths are incremented

accordingly at the coordinator, which then allocates rates for the next coflow in the same FIFO queue, and so on.

Though the above max-min approach has the advantage of minimizing bandwidth wastage, it slows down the coordinator which has to iterate over many flows. In our experiments, we used a simple scheme where we assign the *entire* bandwidth at the sender and receiver ports to one flow of the coflow at the head of the FIFO queue at a time. We found that this simple scheme has very marginal effect on CCTs but makes the rate assignment process considerably faster.

#### 4.6 Additional design issues

**Thin coflow bypass** Recall that, in PHILAE, when a new coflow arrives, PHILAE only schedules its pilot flows. All other flows of that coflow are delayed until the pilot flows finish and coflow size is known. However, such a design choice can inadvertently lead to higher CCTs for coflows, particularly for thin coflows, e.g., a two-flow coflow would end up serializing scheduling its two flows, one for the piloting purpose.

To avoid CCT degradations for thin coflows, we schedule all flows of a coflow if its width is under a threshold (set to 7 in PHILAE; §8.6 provides sensitivity analysis for thresholds).

**Failure tolerance and recovery** Cluster dynamics such as stragglers or node failure can delay some of the flows of a coflow or start new flows, increasing their CCT. The PHILAE design automatically self-adjusts to speed up coflows that are affected by cluster dynamics using the following mechanisms: (1) It adjusts the coflow size as the amount of data left by the coflow, which is essentially the difference between the size calculated using pilot flows and amount of data already sent. (2) It calculates contention only on the ports that have unfinished flows.

**Work Conservation** By default, PHILAE schedules non-pilot flows of a coflow only after all its pilot flows are over. This can lead to some ports being idle where the non-pilot flows are waiting for the pilot flows to finish. In such cases, PHILAE schedules non-pilot flows of coflows which are still in the sampling phase at those ports. In work conservation, the coflows are scheduled in the FIFO order of arrival of coflows.

### 5 Scalability Analysis

Compared to learning coflow sizes using priority queues (PQ-based) [19,30], learning coflow sizes by sampling PHILAE not only reduces the learning overhead as discussed in §3.1 and shown in §8.2, but also significantly reduces the amount of interactions between the coordinator and local agents and thus makes the coordinator highly scalable, as summarized in Table 1.

First, *PQ-based learning requires much more frequent update from local agents*. PQ-based learning estimates coflow sizes by incrementally moving coflows across priority queues

Table 1: Comparison of frequency of interactions between the coordinator and local agents.

	Update of data sent	Update of flow completion	Rate calculation
PHILAE	No	Yes	Event triggered
Aalo	Periodic ( $\delta$ )	Yes	Periodic ( $\delta$ )

according to the data sent by them so far. As such, the scheduler needs frequent updates (every  $\delta$  ms) of data sent per coflow from the local agents. In contrast, PHILAE directly estimates a coflow’s size upon the completion of all its pilot flows. The only updates PHILAE needs from the local agents are about the flow completion which is needed for updating contentions and removing flows from active consideration..

Second, *PQ-based learning results in much more frequent rate allocation*. In sampling-based approach, since coflow sizes are estimated only once, coflows are re-ordered only upon coflow completion or arrival events or in the case of contention based policies only when contention changes, which is triggered by completion of all the flows of a coflow at a port. In contrast, in PQ-based learning, at every  $\delta$  interval, coflow data sent are updated and coflow priority may get updated, which will trigger new rate assignment.

Our scalability experiments in §9.3 confirms that PHILAE achieves much higher scalability than Aalo.

## 6 Implementation

We implemented both PHILAE and Aalo scheduling policies in the same framework consisting of the global coordinator and local agents (Fig. 2), in 5.2 KLoC in C++.

**Coordinator:** The coordinator schedules the coflows based on the operations received from the registering framework. The key implementation challenge for the coordinator is that it needs to be fast in computing and updating the schedules. The PHILAE coordinator is optimized for speed using a variety of techniques including pipelining, process affinity, and concurrency whenever possible.

**Local agents:** The local agents update the global coordinator only upon completion of a flow, along with its length if it is a pilot flow. Local agents schedule the coflows based on the last schedule received from the coordinator. They comply to the last schedule until a new schedule is received. To intercept the packets from the flows, local agents require the compute framework to replace `datasend()`, `datarecv()` APIs with the corresponding PHILAE APIs, which incurs very small overhead.

**Coflow operations:** The global coordinator runs independently from, and is not coupled to, any compute framework, which makes it general enough to be used with any framework. It provides RESTful APIs to the frameworks for coflow operations: (a) `register()` for registering a new coflow when it enters, (b) `deregister()` for removing a

coflow when it exits, and (c) `update()` for updating coflow status whenever there is a change in the coflow structure, e.g., during task migration and restarts after node failures.

## 7 Evaluation Highlights

We evaluated PHILAE using a 150-node and a 900-node testbed cluster in Azure and using large scale simulations by utilizing a publicly available Hive/MapReduce trace collected from a 3000-machine, 150-rack Facebook production cluster [4] and multiple derived traces with varying degrees of flow size skew to measure PHILAE’s robustness to skew.

- **Facebook (FB) trace:** The trace contains 150 ports and 526 ( $> 7 \times 10^5$  flows) coflows, that are extracted from Hive/MapReduce jobs from a Facebook production cluster. Each coflow consists of pair-wise flows between a set of senders and a set of receivers.

Due to the lack of other publicly available coflow trace<sup>1</sup>, we derived three additional traces using the original Facebook trace in order to more thoroughly evaluate PHILAE under varying coflow size skew:

- **Low-skew-filtered:** Starting with the FB trace, we filtered out coflows that have skew (*max flow length/min flow length*) less than a constant  $k$ . We generated five traces in this class with  $k = 1, 2, 3, 4, 5$ . The filtered traces have 142, 100, 65, 51 and 43 coflows, respectively.
- **Mantri-like:** Starting with the FB trace, we adjusted the sizes of the flows sent by the mappers, keeping the total reducer data the same as given in the original trace, to match the skew of a large Microsoft production cluster trace as described in Mantri [12]. In particular, the sizes are adjusted so that the coefficients of variation across mapper data are about 0.34 in the 50<sup>th</sup> percentile case and 3.1 in the 90<sup>th</sup> percentile case. This trace has the same numbers of coflows and ports as the FB trace.
- **Wide-coflows-only:** We filtered out all the coflows in the FB trace with the total number of flows  $\leq 7$ , the default thin coflow bypass threshold (`thinLimit`) in PHILAE. The filtered trace has 269 coflows spreading over 150 ports.

The primary performance metrics used in the evaluation are CCT or CCT speedup, defined as the ratio of a CCT under other baseline algorithms and under PHILAE, piloting overhead, and coflow size estimation accuracy.

The highlights of our evaluation results are:

- (1) PHILAE significantly improves the CCTs. In simulation using the FB trace, the average CCT is improved by  $1.51 \times$  over the prior art, Aalo. Individual CCT speedups are  $1.78 \times$  in the median case ( $P90 = 9.58 \times$ ). For the Mantri-like trace,

<sup>1</sup>A challenge that has also been faced by previous work on coflow scheduling such as [19, 27, 29, 44].



Table 2: Performance improvement over Aalo for varying pilot flow selection schemes.

	Constant	Proportional to number of senders					Proportional to number of flows	
	2	5%	10%	20%	50%	100%	1%	10%
Avg. error	13.21%	6.14%	5.42%	4.94%	5.53%	4.25%	4.15%	2.90%
Avg. CCT	1.27x	1.51x	1.45x	1.50x	1.50x	1.50x	1.43x	0.49x
P50 speedup	1.75x	1.78x	1.76x	1.71x	1.52x	1.40x	1.33x	0.69x
P90 speedup	9.00x	9.58x	9.00x	9.15x	8.33x	8.45x	8.23x	8.23x

the average CCT is improved by  $1.36\times$  and individual CCT speedups are  $1.75\times$  in the median case ( $P90 = 12.0\times$ ).

(2) The CCT improvement mainly stems from the reduction in the learning overhead (in terms of latency and amount of data sent) in determining the right queue for the coflows. Compared to Aalo, median reduction in the absolute latency in finding the right queue for coflows in PHILAE is  $19.0\times$ , and in absolute amount of data sent is  $20.0\times$  (§8.2).

(3) PHILAE improvements are consistent when varying the skew among the flow sizes in a coflow (§8.5).

(4) PHILAE improvements are consistent when varying its parameters (§8.6).

(5) The PHILAE coordinator is much more scalable than that of Aalo (§9.3).

## 8 Simulation

We present detailed simulation results in this section, and the testbed evaluation of our prototype in §9.

**Experimental setup:** Our simulated cluster uses the same number of nodes (sending and receiving network ports) as in the trace. As in [19], we assume full bisection bandwidth is available, and congestion can happen only at network ports.

The *default parameters* for Aalo and PHILAE in the experiments are: starting queue threshold ( $Q_0^{hi}$ ) is 10MB, exponential threshold growth factor ( $E$ ) is 10, number of queues ( $K$ ) is set to 10, the weights assigned to individual priority queues decrease exponentially by a factor of 10, and the new schedule calculation interval  $\delta$  is set to 8ms for the 150-node cluster<sup>2</sup>, the default suggested in its publicly available simulator [19]. In PHILAE, a new schedule is calculated on demand, upon arrival of a new coflow, completion of a coflow, or completion of all pilot flows of a coflow. Finally, in PHILAE the threshold for thinLimit (T) is set to 7, the number of pilot flows assigned to wide coflows are  $\max(1, 0.05 \cdot S)$ , where  $S$  is the number of senders, and the default inter-coflow scheduling policy in PHILAE is Least length-weighted total-port contention.

### 8.1 Pilot flow selection policies

We start by evaluating the impact of different policies in choosing the pilot flows for a coflow in PHILAE. Table 2 summarizes the improvement in average CCT of PHILAE over

<sup>2</sup>8ms is the time to send 1MB of data.

Aalo and average error in size estimation normalized to the actual coflow size, when varying the pilot flow selection policy while keeping other parameters as the default in PHILAE, using the FB trace.

Unsurprisingly, the estimation accuracy increases when increasing the number of pilot flows across the three selection schemes: constant, fraction of senders, and fraction of total flows. However, as the number of pilot flows increases (over the range of parameter choices), the CCT speedup (P50 and P90 of individual coflow CCT speedups) decreases. This is because the benefit from size estimation accuracy improvement from using additional pilot flows does not offset the added overhead from completing the additional pilot flows and the delay they incur to other coflows.

We find sampling 5% of the number of senders per coflow strikes a good trade-off between piloting overhead and size estimation accuracy leading to the best CCT reduction. We thus set it ( $0.05 \cdot S$ ) as the default pilot flow selection policy.

### 8.2 Piloting overhead and accuracy




Figure 3: PHILAE coflow size learning accuracy. Coflows that did not go through the piloting phase (48%) are not shown.

Next, using the default pilot selection policy, we evaluate PHILAE’s effectiveness in estimating coflow sizes by sampling pilot flows. Fig. 3 shows a scatter plot of the actual coflow size vs. estimated size from running PHILAE under the default settings. We observe that PHILAE coflow’s size estimation is highly accurate except for a few outliers. Overall, the average and standard deviation of relative estimation error are 0.06 and 0.15, respectively, and for the top 99% and 95% coflows (in terms of estimation accuracy), the average (standard deviation) of relative error are only 0.05 (0.12) and 0.03 (0.07) respectively. Interestingly, a few coflows experience large estimation errors, and we found they all have very high skew

in their flow lengths; the mean standard deviation in flow lengths, normalized by the average length, of the bottom 1% (in terms of accuracy) ranges between 4.6 and 6.8.

Fig. 1 shows the cost of estimating the correct queue for each coflow in PHILAE and Aalo, measured as the time in learning the coflow size as a fraction of the coflow’s CCT in PHILAE and Aalo. We see that under PHILAE, about 63% of the coflows spent less than 1% of their CCT in the learning phase, while under Aalo, 63% coflows reached the correct priority queue after spending up to 22% of their CCT moving across other queues. Compared to Aalo, PHILAE in the median case sends  $20\times$  less data in determining the right queue and reduces the latency in determining the right queue by  $19\times$ .

### 8.3 Inter-coflow scheduling policies

PHILAE differs from Aalo in two ways: online size estimation and inter-flow scheduling policy. Here, we evaluate the effectiveness of the four inter-coflow scheduling policies of PHILAE discussed in §4.4, keeping the remaining parameters as the default. Such evaluation allows us to decouple the contribution of sampling-based learning from the effect of scheduling policy difference.

Table 3 shows the CCT improvement of PHILAE under the four inter-flow scheduling policies over Aalo. We make the following observations.

First, PHILAE with the purely sized-based policy, **Smallest job first (A)**, which uses the same inter-queue and intra-queue scheduling policy as Aalo and only differs from Aalo in coflow size estimation, reduces the average CCT (P50) of Aalo by  $1.40\times$  ( $1.48\times$ ).

In contrast, the default PHILAE uses **Least length-weighted total-port contention (D)**, which uses the sum of size-weighted port contention to assign coflows to priority queues, and slightly outperforms the size-based policy A; it reduces the average CCT (P50) of Aalo by  $1.51\times$  ( $1.78\times$ ). This is because it captures the diversity of contention at different ports, which happens often in real distributed settings, and at the same time accounts for the coflow size by using length-weighted sum of the port-wise contention. The above results for policy A and policy D indicate that the primary improvement in PHILAE comes from its sampling-based coflow size estimation scheme.

**Shortest remaining time first (B)** performs similarly as smallest job first. This is because the preemptive nature of SRTF will kick in only on arrival of new coflows. Also, although SRTF is advantageous for small coflows, since PHILAE already schedules thin coflows at high priority, many thin and thus small coflows are anyways being scheduled at high priority under both policies A and B, and as a result they perform similarly.

Finally, **Least contention first (C)** performs poorly. This is because contention for a coflow is defined as the unique number of other coflows that share ports, and as a result such a policy completely ignores the size (length) of the coflows.

Table 3: CCT speedup in PHILAE under different inter-coflow scheduling policies (§4.4) over Aalo.

Priority estimation metric	P50	P90	Avg. CCT
Estimated size (A)	1.48x	8.27x	1.40x
Remaining size (B)	1.54x	8.34x	1.37x
Global Contention (C)	0.75x	8.26x	0.13x
Length-weighted total-port contention (D) (PHILAE)	1.78x	9.58x	1.51x

### 8.4 Average CCT improvement

We now compare the CCT speedups of PHILAE against 5 well-known coflow scheduling policies: (1) Aalo, (2) Aalo-Oracle, which is an oracle version of Aalo where the scheduler knows the final queue of a coflow upon its arrival time and directly starts the coflow from that queue, (3) SEBF in Varys [21] which assumes the knowledge of coflow sizes apriori and uses the Shortest Effective Bottleneck First policy, where the coflow whose slowest flow will finish first is scheduled first. (4) FIFO, which is a single queue FIFO based coflow scheduler, and (5) FAIR, which uses per-flow fair sharing. We do not include Saath [30] in the comparison as it does not provide the same liveness guarantees as PHILAE which as discussed in §4.3 can obscure the comparison result. All experiments use the default parameters discussed in the setup, including  $K, E, S$ , unless otherwise stated. The results are shown in Fig. 4(a). We make the following observations.

First, we compare CCT under PHILAE against under Aalo-Oracle, where Aalo-Oracle starts all coflows at the correct priority queues (*i.e.*, no learning overhead). PHILAE improves the average CCT by  $1.18\times$  and P50 CCT by  $1.40\times$ , respectively. Since Aalo-Oracle pays no overhead for coflow size estimation, its worse performance suggests that using length-weighted total-port contention in assigning coflows to the priority queues in PHILAE outperforms Aalo’s size-based, contention-oblivious policy in assigning coflows to the queues.

Second, PHILAE improves the average CCT over Aalo by  $1.51\times$  (median) and P50 by  $1.78$ . The significant *additional* improvement on top of the gain over Aalo-Oracle comes from fast and accurate estimation of the right queues for the coflows (Fig. 1).

Third, PHILAE, which requires no coflow size knowledge a priori, achieves comparable performance as SEBF [21]; it reduces the average CCT by  $1.16\times$ . Again this is because its total-port contention policy outperforms the contention-oblivious SEBF.

Finally, PHILAE significantly outperforms the single-queue FIFO-based coflow scheduler, with a median (P90) CCT speedup of  $3.00$  ( $77.96\times$ ) and average CCT speedup of  $3.16\times$ , and the un-coordinated flow-level fair-share scheduler, with a median (P90) CCT speedup of  $70.82\times$  ( $1947\times$ ) and average CCT speedup of  $5.66\times$ .

To gain insight into how different coflows are affected by PHILAE over Aalo, we group the coflows in the trace into




Figure 4: CCT speedup using PHILAE compared to using other coflow schedulers on different traces. In Fig. 4(c), the x-axis denotes the minimum skew in the 5 Low-skew-filtered traces.




Figure 5: Performance breakdown into bins shown in Table 4.

Figure 6: [Testbed] Distribution of speedup in CCT and JCT in PHILAE using the FB trace.

Table 4: Bins based on total coflow size and width (number of flows). The numbers in brackets denote the fraction of coflows in that bin.

	width $\leq 7$ (thin)	width $> 7$ (wide)
size $\leq 100$ MB (small)	bin-1 (44.3%)	bin-2 (24.1%)
size $> 100$ MB (large)	bin-3 (4.5%)	bin-4 (27.1%)

four bins defined in Table 4, and show in Fig. 5 the CCT speedups for each bin. We see that PHILAE improves CCT for all coflows in bin 1 and 3 and for large fraction in bin-4. Most of the underperforming coflows fall in bin-2. Coflows in bin-2 have width  $> 7$  and size  $< 100$ MB, *i.e.*, the flows are short but wide. Because the width exceeds the thinLimit, PHILAE schedules the pilot flows to estimate the coflow size first (§4). Thus, although the remaining flows are short, they get delayed until the completion of the pilot flows, which results in CCT increase.

Finally, since thin coflows benefit from PHILAE’s scheme of bypassing probing for thin coflows, we also compare PHILAE with other schemes using the Wide-coflows-only trace which consists of all coflows wider than the default thinLimit (7) in PHILAE. Fig. 4(b) shows that PHILAE continues to perform well, reducing the average CCT by  $1.54\times$ ,  $1.15\times$ , and  $1.12\times$  over Aalo, Aalo-Oracle, and SEBF, respectively.

## 8.5 Robustness to coflow data skew

Next, we evaluate PHILAE’s robustness to flow size skew by comparing it against Aalo using traces with varying degrees of skew. First, we evaluate PHILAE using the Mantri-like trace. Fig. 4(d) shows that PHILAE consistently outperforms prior-art coflow schedulers. In particular, PHILAE reduces the average CCT by  $1.36\times$  compared to Aalo. Second, we evalu-

ate PHILAE using the Low-skew-filtered traces which have low skew coflows filtered out. Fig. 4(c) shows that PHILAE performs better than Aalo even with highly skewed traces and reduces the average CCT by  $1.45\times$ ,  $1.44\times$ ,  $1.44\times$ ,  $1.40\times$  and  $1.38\times$  for the five Low-skew-filtered traces containing coflows with skew of at least 1, 2, 3, 4 and 5, respectively.

## 8.6 Sensitivity analysis

Compared to Aalo, PHILAE has only two additional parameters: thinLimit and flow sampling rate. We already discussed the choice of sampling rate in §8.1. Below, we evaluate the sensitivity of PHILAE to thinLimit and other design parameters common to Aalo by varying one parameter at a time while keeping the rest as the default.

**Thin coflow bypassing limit ( $T$ )** In this experiment, we vary thinLimit ( $T$ ) in PHILAE for bypassing coflows from the probing phase. The result in Fig. 7(a) shows that the average CCT remains almost the same as  $T$  increases. This is because the average CCT is dominated by wide and large coflows, which are not affected by thinLimit. However, the P50 speedup increases till  $T = 7$  and tapers off after  $T = 7$ . The reason for the CCT improvement until  $T = 7$  is that all flows of thin coflows (with width  $\leq 7$ ) are scheduled immediately upon arrival which improves their CCT, and the number of thin coflows is significant.

**Start queue threshold ( $Q_0^{hi}$ )** We next vary the threshold for the first priority queue from 2 MB to 64 MB. Fig. 7(b) shows the average CCT of PHILAE over Aalo. Overall, PHILAE is not very sensitive to the threshold of first priority queue and the CCT speedup over Aalo is within 8% of the default PHILAE (10 MB). The speedup appears to oscillate with a periodicity of  $5\times$  to  $10\times$ . For example, the speedups for 2 MB and 64 MB are close to that of the default (10 MB), while for 4 MB and 32 MB are lower. This can be explained by the impact of the first queue threshold on job segregation; with the default queue threshold growth factor of 10, every time the first queue threshold changes by close to  $10\times$ , the distribution of jobs across the queues become similar.

**Multiplication factor ( $E$ )** In this experiment, we vary the queue threshold growth factor from 2 to 64. Recall that the queue thresholds are computed as  $Q_q^{hi} = Q_{q-1}^{hi} \cdot E$ . Thus, as  $E$  grows, the number of queues decreases. As shown in Fig. 7(c),




Figure 7: [Simulation] PHILAE sensitivity analysis. We vary one parameter of PHILAE keeping rest same as default and compare it with Aalo.

smaller queue threshold multiplication factor which leads to more queues performs better because of fine-grained priority segregation.

## 9 Testbed Evaluation

Next, we deployed PHILAE in a 150-machine Azure cluster and a 900-machine cluster to evaluate its performance and scalability.

**Testbed setup:** We rerun the FB trace on a Spark-like framework on a 150-node cluster in Microsoft Azure [5]. The coordinator runs on a Standard DS15 v2 server with 20-core 2.4 GHz Intel Xeon E5-2673 v3 (Haswell) processor and 140GB memory. The local agents run on D2v2 with the same processor as the coordinator with 2-core and 7GB memory. The machines on which local agents run have 1 Gbps network bandwidth. Similarly as in simulations, our testbed evaluation keeps the same flow lengths and flow ports in trace replay. All the experiments use default parameters  $K, E, S$  and the default pilot flow selection policy.

### 9.1 CCT Improvement

In this experiment, we measure CCT improvements of PHILAE compared to Aalo. Fig. 6 shows the CDF of the CCT speedup of individual coflows under PHILAE compared to under Aalo. The average CCT improvement is  $1.50\times$  which is similar to the results in the simulation experiments. We also observe  $1.63\times$  P50 speedup and  $8.00\times$  P90 speedup.

We also evaluated PHILAE using the Wide-coflow-only trace. Table 5 shows that PHILAE achieves  $1.52\times$  improvement in average CCT over Aalo, similar to that using the full FB trace. This is because the improvement in average CCT is dominated by large coflows, PHILAE is speeding up large coflows, and the Wide-coflow-only trace consists of mostly large coflows.

### 9.2 Job Completion Time

Next, we evaluate how the improvement in CCT affects the job completion time (JCT). In data clusters, different jobs

Table 5: [Testbed] CCT improvement in PHILAE as compared to Aalo.

	P50	P90	Avg. CCT
FB Trace	$1.63\times$	$8.00\times$	$1.50\times$
Wide-coflow-only	$1.05\times$	$2.14\times$	$1.49\times$

Table 6: [Testbed] Average (standard deviation) coordinator CPU time (ms) per scheduling interval in 900-port runs. PHILAE did not have to calculate and send new rates in 66% of intervals, which contributes to its low average.

	Rate Calc.	New Rate Send	Update Recv.	Total
PHILAE	2.99 (5.35)	4.90 (11.25)	6.89 (17.78)	14.80 (28.84)
Aalo	4.28 (4.14)	17.65 (20.9)	10.97 (19.98)	32.90 (34.09)

spend different fractions of their total job time in data shuffle. In this experiment, we used 526 jobs, each corresponding to one coflow in the FB trace. The fraction of time that the jobs spent in the shuffle phase follows the same distribution used in Aalo [19], *i.e.*, 61% jobs spent less than 25% of their total time in shuffle, 13% jobs spent 25-49%, another 14% jobs spent 50-74%, and the remaining spent over 75% of their total time in shuffle. Fig. 6 shows the CDF of individual speedups in JCT. Across all jobs, PHILAE reduces the job completion time by  $1.16\times$  in the median case and  $7.87\times$  in the 90<sup>th</sup> percentile. This shows that improved CCT translates into better job completion time. As expected, the improvement in job completion time is smaller than the improvement in CCT because job completion time depends on the time spent in both compute and shuffle (communication) stages, and PHILAE improves only the communication stage.

### 9.3 Scalability

Finally, we evaluate the scalability of PHILAE by comparing its performance with Aalo on a 900-node cluster. To drive the evaluation, we derive a 900-port trace by replicating the FB trace 6 times across ports, *i.e.*, we replicated each job 6 times, keeping the arrival time for each copy the same but assigning sending and receiving ports in increments of 150 (the cluster size for the original trace). We also increased the

Table 7: [Testbed] Percentage of scheduling intervals where synchronization and rate calculation took more than  $\delta$  for 150-port and  $\delta' (= 6 \times \delta)$  for 900-port runs.

	150 ports	900 ports
PHILAE	1%	10%
Aalo	16%	37%

scheduling interval  $\delta$  by 6 times to  $\delta' = 6 \times \delta$ .

PHILAE achieved  $2.72 \times$  ( $9.78 \times$ ) speedup in average (P90) CCT over Aalo. The higher speedup compared to the 150-node runs ( $1.50 \times$ ) comes from higher scalability of PHILAE. In 900-node runs, Aalo was not able to finish receiving updates, calculating new rates and updating local agents of new rates within  $\delta'$  in 37% of the intervals, whereas PHILAE only missed the deadline in 10% of the intervals. For 150-node runs these values are 16% for Aalo and 1% for PHILAE. The 21% increase in missed scheduling intervals in 900-node runs in Aalo resulted in local agents executing more frequently with outdated rates. As a result, PHILAE achieved even higher speedup in 900-node runs.

As discussed in §5, Aalo’s poorer coordinator scalability comes from more frequent updates from local agents and more frequent rate allocation, which result in longer coordinator CPU time in each scheduling interval. Table 6 shows the average coordinator CPU usage per interval and its breakdown. We see that (1) on average PHILAE spends much less time than Aalo in receiving updates from local agents, because PHILAE does not need updates from local agents at every interval – on average in every scheduling interval PHILAE receives updates from 49 local agents whereas Aalo receives from 429 local agents, and (2) on average PHILAE spends much less time calculating new rates and send new rates. This is because rate calculation in PHILAE is triggered by events and PHILAE did not have to flush rates in 66% of the intervals.

## 10 Related Work

**Coflow scheduling:** In this paper, we have shown PHILAE outperforms prior-art non-clairvoyant coflow scheduler Aalo from more efficient learning of coflow sizes online. Saath [30] and Graviton [29] also learn coflow sizes online using priority queues and hence suffers the same inefficiency as Aalo. Graviton [29] uses the number of ports a coflow is present at, as an additional indicator of its size. In [19], Aalo was shown to outperform previous non-clairvoyant coflow schedulers Baraat [24] by using global coordination, and Orchestra [20] by avoiding head-of-line blocking.

Clairvoyant coflow schedulers such as Varys [21] and Sincronia [7] assume prior knowledge of coflows upon arrival. Varys runs a shortest-effective-bottleneck-first heuristic for inter-coflow scheduling and performs per-flow rate allocation at the coordinator. Sincronia improves the scalability of the centralized coordinator of Varys by only calculating the coflow ordering at the coordinator (by solving an LP) and

offloading flow rate allocation to individual local agents. Sincronia is orthogonal to PHILAE; once coflow sizes are learned through sampling, ideas from Sincronia can be adopted in PHILAE to order coflows and offload rate allocation to local ports. CODA [44] tackles an orthogonal problem of identifying flows of individual coflows online.

However, recent studies [19, 40] have shown various reasons why it is not very plausible to learn flow sizes from applications beforehand. For example, many applications stream data as soon as data are generated and thus the application does not know the flow sizes until flow completion, and learning flow sizes from applications requires changing either the network stack or the applications.

**Flow scheduling:** There exist a rich body of prior work on flow scheduling. Efforts to minimize flow completion time (FCT), both with prior information (e.g., PDQ [26], pFabric [9]) and without prior information (e.g., Fastpass [35], PIAS [13], [14]), fall short in minimizing CCTs which depend on the completion of the last flow [21]. Similarly, Hedera [8] and MicroTE [15] schedule the flows with the goal of reducing the overall FCT, which again is different from reducing the overall CCT of coflows.

**Speculative scheduling** Recent works [16, 33] use the idea of online requirement estimation for scheduling in data-center. In [31], recurring big data analytics jobs are scheduled using their history.

**Job scheduling:** There have been much work on scheduling in analytic systems and storage at scale by improving speculative tasks [11, 12, 43], improving locality [10, 41], and end-point flexibility [17, 38]. The coflow abstraction is complementary to these work, and can benefit from them. Combining coflow with these approaches remains a future work.

**Scheduling in parallel processors:** Coflow scheduling by exploiting the spatial dimension bears similarity to scheduling processes on parallel processors and multi-cores, where many variations of FIFO [37], FIFO with backfilling [32] and gang scheduling [25] have been proposed.

## 11 Conclusion

State-of-the-art online coflow schedulers approximate the classic SJF by implicitly learning coflow sizes and pay a high penalty for large coflows. We propose the novel idea of sampling in the spatial dimension of coflows to explicitly and efficiently learn coflow sizes online to enable efficient online SJF scheduling. Our extensive simulation and testbed experiments show the new design offers significant performance improvement over prior art. Further, the sampling-in-spatial-dimension technique can be generalized to other distributed scheduling problems such as cluster job scheduling. We have made our simulator publicly available at <https://github.com/coflowPhilae/simulator> [6].

**Acknowledgement** We thank our shepherd Patrick Stuedi and the anonymous reviewers for their insightful comments.

## References

- [1] Apache hadoop. <http://hadoop.apache.org>.
- [2] Apache spark. <http://spark.apache.org>.
- [3] Apache tez. <http://tez.apache.org>.
- [4] Coflow trace from facebook datacenter. <https://github.com/coflow/coflow-benchmark>.
- [5] Microsoft azure. <http://azure.microsoft.com>.
- [6] Philae simulator. <https://github.com/coflowPhilae/simulator>.
- [7] Saksham Agarwal, Shijin Rajakrishnan, Akshay Narayan, Rachit Agarwal, David Shmoys, and Amin Vahdat. Sincronia: Near-optimal network design for coflows. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, SIGCOMM '18*, pages 16–29, New York, NY, USA, 2018. ACM.
- [8] Mohammad Al-Fares, Sivasankar Radhakrishnan, Barath Raghavan, Nelson Huang, and Amin Vahdat. Hedera: Dynamic flow scheduling for data center networks. In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation, NSDI'10*, pages 19–19, Berkeley, CA, USA, 2010. USENIX Association.
- [9] Mohammad Alizadeh, Shuang Yang, Milad Sharif, Sachin Katti, Nick McKeown, Balaji Prabhakar, and Scott Shenker. pfabric: Minimal near-optimal datacenter transport. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, SIGCOMM '13*, pages 435–446, New York, NY, USA, 2013. ACM.
- [10] Ganesh Ananthanarayanan, Sameer Agarwal, Srikanth Kandula, Albert Greenberg, Ion Stoica, Duke Harlan, and Ed Harris. Scarlett: Coping with skewed content popularity in mapreduce clusters. In *Proceedings of the Sixth Conference on Computer Systems, EuroSys '11*, pages 287–300, New York, NY, USA, 2011. ACM.
- [11] Ganesh Ananthanarayanan, Ali Ghodsi, Scott Shenker, and Ion Stoica. Effective straggler mitigation: Attack of the clones. In *Proceedings of the 10th USENIX Conference on Networked Systems Design and Implementation, nsdi'13*, pages 185–198, Berkeley, CA, USA, 2013. USENIX Association.
- [12] Ganesh Ananthanarayanan, Srikanth Kandula, Albert Greenberg, Ion Stoica, Yi Lu, Bikas Saha, and Edward Harris. Reining in the outliers in map-reduce clusters using mantri. In *Proceedings of the 9th USENIX Conference on Operating Systems Design and Implementation, OSDI'10*, pages 265–278, Berkeley, CA, USA, 2010. USENIX Association.
- [13] Wei Bai, Li Chen, Kai Chen, Dongsu Han, Chen Tian, and Weicheng Sun. Pias: Practical information-agnostic flow scheduling for data center networks. In *Proceedings of the 13th ACM Workshop on Hot Topics in Networks, HotNets-XIII*, pages 25:1–25:7, New York, NY, USA, 2014. ACM.
- [14] Wei Bai, Li Chen, Kai Chen, Dongsu Han, Chen Tian, and Hao Wang. Information-agnostic flow scheduling for commodity data centers. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, pages 455–468, Oakland, CA, 2015. USENIX Association.
- [15] Theophilus Benson, Ashok Anand, Aditya Akella, and Ming Zhang. Microte: Fine grained traffic engineering for data centers. In *Proceedings of the Seventh Conference on Emerging Networking EXperiments and Technologies, CoNEXT '11*, pages 8:1–8:12, New York, NY, USA, 2011. ACM.
- [16] Inho Cho, Keon Jang, and Dongsu Han. Credit-scheduled delay-bounded congestion control for datacenters. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication, SIGCOMM '17*, pages 239–252, New York, NY, USA, 2017. ACM.
- [17] Mosharaf Chowdhury, Srikanth Kandula, and Ion Stoica. Leveraging endpoint flexibility in data-intensive clusters. In *Proceedings of the ACM SIGCOMM 2013 Conference on SIGCOMM, SIGCOMM '13*, pages 231–242, New York, NY, USA, 2013. ACM.
- [18] Mosharaf Chowdhury and Ion Stoica. Coflow: A networking abstraction for cluster applications. In *Proceedings of the 11th ACM Workshop on Hot Topics in Networks, HotNets-XI*, pages 31–36, New York, NY, USA, 2012. ACM.
- [19] Mosharaf Chowdhury and Ion Stoica. Efficient coflow scheduling without prior knowledge. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM '15*, pages 393–406, New York, NY, USA, 2015. ACM.
- [20] Mosharaf Chowdhury, Matei Zaharia, Justin Ma, Michael I. Jordan, and Ion Stoica. Managing data transfers in computer clusters with orchestra. In *Proceedings of the ACM SIGCOMM 2011 Conference, SIGCOMM '11*, pages 98–109, New York, NY, USA, 2011. ACM.

- [21] Mosharaf Chowdhury, Yuan Zhong, and Ion Stoica. Efficient coflow scheduling with varies. In *Proceedings of the 2014 ACM Conference on SIGCOMM*, SIGCOMM '14, pages 443–454, New York, NY, USA, 2014. ACM.
- [22] Tyson Condie, Neil Conway, Peter Alvaro, Joseph M. Hellerstein, Khaled Elmeleegy, and Russell Sears. Mapreduce online. In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*, NSDI'10, pages 21–21, Berkeley, CA, USA, 2010. USENIX Association.
- [23] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: Simplified data processing on large clusters. volume 51, pages 107–113, New York, NY, USA, January 2008. ACM.
- [24] Fahad R. Dogar, Thomas Karagiannis, Hitesh Ballani, and Antony Rowstron. Decentralized task-aware scheduling for data center networks. In *Proceedings of the 2014 ACM Conference on SIGCOMM*, SIGCOMM '14, pages 431–442, New York, NY, USA, 2014. ACM.
- [25] Dror G. Feitelson and Morris A. Jette. Improved utilization and responsiveness with gang scheduling. In *Proceedings of the Job Scheduling Strategies for Parallel Processing*, IPSP '97, pages 238–261, London, UK, UK, 1997. Springer-Verlag.
- [26] Chi-Yao Hong, Matthew Caesar, and P. Brighten Godfrey. Finishing flows quickly with preemptive scheduling. In *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM '12, pages 127–138, New York, NY, USA, 2012. ACM.
- [27] Xin Sunny Huang, Xiaoye Steven Sun, and T.S. Eugene Ng. Sunflow: Efficient optical circuit scheduling for coflows. In *Proceedings of the 12th International on Conference on Emerging Networking EXperiments and Technologies*, CoNEXT '16, pages 297–311, New York, NY, USA, 2016. ACM.
- [28] Michael Isard, Mihai Budiu, Yuan Yu, Andrew Birrell, and Dennis Fetterly. Dryad: Distributed data-parallel programs from sequential building blocks. In *Proceedings of the 2Nd ACM SIGOPS/EuroSys European Conference on Computer Systems 2007*, EuroSys '07, pages 59–72, New York, NY, USA, 2007. ACM.
- [29] Akshay Jajoo, Rohan Gandhi, and Y. Charlie Hu. Graviton: Twisting space and time to speed-up coflows. In *8th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 16)*, Denver, CO, 2016. USENIX Association.
- [30] Akshay Jajoo, Rohan Gandhi, Y. Charlie Hu, and Cheng-Kok Koh. Saath: Speeding up coflows by exploiting the spatial dimension. In *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies*, CoNEXT '17, pages 439–450, New York, NY, USA, 2017. ACM.
- [31] Virajith Jalaparti, Peter Bodik, Ishai Menache, Sriram Rao, Konstantin Makarychev, and Matthew Caesar. Network-aware scheduling for data-parallel jobs: Plan when you can. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, SIGCOMM '15, pages 407–420, New York, NY, USA, 2015. ACM.
- [32] David A. Lifka. The anl/ibm sp scheduling system. In *Proceedings of the Workshop on Job Scheduling Strategies for Parallel Processing*, IPSP '95, pages 295–303, London, UK, UK, 1995. Springer-Verlag.
- [33] Masoud Moshref, Minlan Yu, Ramesh Govindan, and Amin Vahdat. Trumpet: Timely and precise triggers in data centers. In *Proceedings of the 2016 ACM SIGCOMM Conference*, SIGCOMM '16, pages 129–143, New York, NY, USA, 2016. ACM.
- [34] Stanley Lemeshow Paul S. Levy. *Sampling of Populations: Methods and Applications*. Wiley, 4 edition, Jun 2012.
- [35] Jonathan Perry, Amy Ousterhout, Hari Balakrishnan, Devavrat Shah, and Hans Fugal. Fastpass: A centralized "zero-queue" datacenter network. In *Proceedings of the 2014 ACM Conference on SIGCOMM*, SIGCOMM '14, pages 307–318, New York, NY, USA, 2014. ACM.
- [36] Christopher J. Rossbach, Yuan Yu, Jon Currey, Jean-Philippe Martin, and Dennis Fetterly. Dandelion: A compiler and runtime for heterogeneous systems. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, SOSP '13, pages 49–68, New York, NY, USA, 2013. ACM.
- [37] Uwe Schwiegelshohn and Ramin Yahyapour. Analysis of first-come-first-serve parallel job scheduling. In *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '98, pages 629–638, Philadelphia, PA, USA, 1998. Society for Industrial and Applied Mathematics.
- [38] David Shue, Michael J. Freedman, and Anees Shaikh. Performance isolation and fairness for multi-tenant cloud storage. In *Proceedings of the 10th USENIX Conference on Operating Systems Design and Implementation*, OSDI'12, pages 349–362, Berkeley, CA, USA, 2012. USENIX Association.
- [39] Abraham Silberschatz, Peter B. Galvin, and Greg Gagne. *Process Scheduling. Operating System Concepts*. John Wiley & Sons, 8 edition, 2010.

- [40] Vojislav Đukić, Sangeetha Abdu Jyothi, Bojan Karlas, Muhsen Owaida, Ce Zhang, and Ankit Singla. Is advance knowledge of flow sizes a plausible assumption? In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, pages 565–580, Boston, MA, 2019. USENIX Association.
- [41] Matei Zaharia, Dhruba Borthakur, Joydeep Sen Sarma, Khaled Elmeleegy, Scott Shenker, and Ion Stoica. Delay scheduling: A simple technique for achieving locality and fairness in cluster scheduling. In *Proceedings of the 5th European Conference on Computer Systems, EuroSys '10*, pages 265–278, New York, NY, USA, 2010. ACM.
- [42] Matei Zaharia, Mosharaf Chowdhury, Michael J. Franklin, Scott Shenker, and Ion Stoica. Spark: Cluster computing with working sets. In *Proceedings of the 2Nd USENIX Conference on Hot Topics in Cloud Computing, HotCloud'10*, pages 10–10, Berkeley, CA, USA, 2010. USENIX Association.
- [43] Matei Zaharia, Andy Konwinski, Anthony D. Joseph, Randy Katz, and Ion Stoica. Improving mapreduce performance in heterogeneous environments. In *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation, OSDI'08*, pages 29–42, Berkeley, CA, USA, 2008. USENIX Association.
- [44] Hong Zhang, Li Chen, Bairen Yi, Kai Chen, Mosharaf Chowdhury, and Yanhui Geng. Coda: Toward automatically identifying and scheduling coflows in the dark. In *Proceedings of the 2016 ACM SIGCOMM Conference, SIGCOMM '16*, pages 160–173, New York, NY, USA, 2016.