



HAL
open science

Performance bounds for priority-based stochastic coflow scheduling

Olivier Brun, Balakrishna Prabhu

► **To cite this version:**

Olivier Brun, Balakrishna Prabhu. Performance bounds for priority-based stochastic coflow scheduling. LAAS - CNRS. 2025. ⟨hal-05017753⟩

HAL Id: hal-05017753

<https://laas.hal.science/hal-05017753v1>

Submitted on 2 Apr 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Performance bounds for priority-based stochastic coflow scheduling

Olivier Brun and Balakrishna J. Prabhu,
LAAS-CNRS, University of Toulouse, CNRS, 31400 Toulouse, France.
E-mails: {brun, balakrishna.prabhu}@laas.fr.

April 2, 2025

Abstract

We consider the coflow scheduling problem in the non-clairvoyant setting, assuming that flow sizes are realized on-line according to given probability distributions. The goal is to minimize the weighted average completion time of coflows in expectation. We first obtain inequalities for this problem that are valid for all non-anticipative order-based rate-allocation policies and define a polyhedral relaxation of the performance space of such scheduling policies. This relaxation is used to analyze the performance of a simple priority policy in which the priority order is computed by *Sincronia* from expected flow sizes instead of their unknown actual values. We establish a bound on the approximation ratio of this priority policy with respect to the optimal priority policy for arbitrary probability distributions of flow sizes (with finite first and second moments). Tighter upper bounds are obtained for some specific distributions. Extensive numerical results suggest that performance of the proposed policy is much better than the upper bound.

1 Introduction

Modern data centres are increasingly used to run data-parallel computing applications such as MapReduce, Hadoop and Spark [13, 33]. These applications alternate between computation stages and communication stages, during which the application's tasks exchange intermediate results using the data centre network. Each application may generate hundreds or thousands of concurrent data transfers which have to share the resources of the datacenter network not only between themselves, but also with the data flows of the other applications. Chowdhury and Stoica showed in [10] that by allocating rates to all these data transfers in some appropriate way, it was possible to significantly reduce the average time spent in communication by data-parallel applications. They introduced the coflow scheduling problem, a coflow being defined as a collection of flows generated by the same parallel applicaton and its completion time being determined by the completion time of the last flow in the collection (that is, it corresponds to the time at which the next computation stage can start). The coflow scheduling problem is then to allocate rates to flows over time so that the average coflow completion time (CCT) is minimized.

Even in the clairvoyant offline setting, in which all coflows to be scheduled are initially present in the system and the precise volume of each and every flow is known, the problem is known to be NP-hard [12]. This follows from the fact that it can be seen as a generalization of the concurrent open-shop scheduling problem in which resources are coupled, as transmitting a flow requires capacity on both the input and output ports. Furthermore, the coflow scheduling problem was shown to be inapproximable below a factor of 2 [5, 25] in this setting. Several heuristics and approximation algorithms for clairvoyant coflow scheduling have been proposed in the last decade, see e.g. [12, 8, 31, 32, 9, 19, 28, 35]. At the present time, the state-of-the-art algorithm for coflow scheduling in the clairvoyant setting is *Sincronia* [3], which was proposed by Agarwal *et al.* in 2018 and provides an approximation ratio of 4, the best known one (see

also [27] and [4] for closely-related references that were published roughly concurrently with [3]). In addition, an attractive feature of **Sincronia** is that it relies on priority mechanisms which are supported by most transport layers. It is worth mentioning that **Sincronia** is not restricted to the offline setting but can also efficiently handle online coflow arrivals.

Although **Sincronia** provides a fairly efficient solution to the coflow scheduling problem, it requires *a priori* knowledge of the flow volumes, which is difficult to obtain in practice. Therefore, some authors have investigated the non-clairvoyant setting, where on a coflow arrival, only the number of flows and their input/output ports are revealed to the coflow scheduler, while their volumes remains unknown [11, 36, 15, 18, 34, 14, 6]. However, none of the proposed algorithms comes with performance guarantees, with the noticeable exception of the **BlindFlow** algorithm which was shown in [6] to be $8p$ -approximate, where p is the maximum number of flows that any coflow can have. Given the orders of magnitude of p in practice, this performance guarantee is of course much weaker than that provided by **Sincronia** in the clairvoyant setting.

It could be argued however that assuming absolutely no prior information on flow volumes, as was done in the above works on the non-clairvoyant setting, is excessively pessimistic. Indeed, historical data on previous executions of data-parallel applications could be exploited to get some insights on what the flow volumes could be. For instance, machine learning (ML) algorithms could be used to obtain predictions on flow volumes from historical data. The improvement of online scheduling algorithms via unreliable ML predictions is a recent line of research which has attracted a lot of studies in the last few years (see e.g. [23, 16]) and which was recently applied to non-clairvoyant coflow scheduling [7]. Another relevant possibility, which is explored in the present paper, is to assume that the unknown flow volumes are on-line realizations of known probability distributions, which can be estimated from flow sizes data collected during previous executions.

Stochastic coflow scheduling has received little attention up to now. To the extent of our knowledge, the only work that has addressed the coflow scheduling problem in the stochastic setting is [19], in which the authors study the non-preemptive coflow scheduling problem. The objective is to minimize the weighted expected completion time of coflows. The proposed approach uses a time-indexed linear programming relaxation first introduced in [29], and uses its solution to come up with a feasible schedule. The proposed algorithm is shown to achieve an approximation ratio of $(1 + 2 \log L)(1 + \sqrt{LCV_{max}})(1 + LCV_{max})(1 + CV_{max}^2)$, where L is the number of ports and CV_{max} is an upper bound on the coefficient of variation of processing times¹. We, on the other hand, consider the preemptive case.

1.1 Contributions

We consider the non-clairvoyant setting in which the size of each flow is unknown to the coflow scheduler and is revealed to it only upon flow completion. Although flow sizes are not known in advance, we assume that they are drawn from known probability distributions. It is shown that for any order-based rate-allocation policy, the vector of expected completion times satisfies certain inequalities, from which a polyhedral relaxation of the performance space of such coflow scheduling policies is obtained. We use this polyhedral relaxation to analyze the performance of a priority policy in which the priority order is computed with **Sincronia** from expected flow sizes. Specifically, it is shown that the approximation ratio of this policy with respect to the optimal priority policy is upper bounded by $4(1 + \sqrt{L}\sigma_{max}/\mu_{min})$ for any distribution of flow sizes with finite first and second moments, where μ_{min} is a lower bound on the expected processing times and σ_{max} is an upper bound on their standard deviations. Tighter bounds can be obtained if we have information on the distribution of processing times. For instance, we show that for normally distributed processing times, the approximation ratio of the **Sincronia**-based priority policy is upper bounded by $4(1 + \sqrt{2 \log L} CV_{max})$, where, as before, CV_{max} is an upper bound on the coefficient of variation of processing times. We also obtain asymptotic bounds for Pareto-distributed processing times.

¹The processing time of a coflow on a port is the total time required for transmitting all of its constituent flows on that port in isolation.

1.2 Organization

The paper is organized as follows. Section 2 introduces the main notations used throughout the paper and presents the stochastic coflow scheduling problem that we address. In Section 3, we establish a polyhedral relaxation of the performance space of order-based coflow scheduling policies and introduce our **Sincronia**-based priority policy. Upper bounds on the approximation ratio of this policy with respect to the optimal priority policy are proven in Section 4. Section 5 is devoted to numerical results. Finally, some conclusions are drawn and future research directions are discussed in Section 6.

2 Stochastic coflow scheduling

We consider a set $\mathcal{C} = \{1, 2, \dots, n\}$ of coflows, which are all present in the system at time 0. Each coflow k is a set F_k of flows, where flow j of coflow k is characterized by its source node, its destination node and its volume $V^{k,j}$. Each coflow k may also be assigned a weight w_k (default weight is 1). As for the network, we assume the *big switch model*² [12] and we let b_ℓ denote the capacity of port $\ell \in \mathcal{L}$. We let $F_{k,\ell}$ be the set of flows of coflow k that use port ℓ either as ingress or egress port, and denote by C_k the completion time of coflow k , that is, the time at which its last flow finishes.

In the clairvoyant setting where flow sizes $V^{k,j}$ are known to the coflow scheduler, the goal is to find a rate allocation $r(t)$ to flows, where $r^{k,j}(t) \in \mathbb{R}_+$ is the rate allocated to flow $j \in F_k$ at time t , which minimizes the weighted average coflow completion time $\sum_{k \in \mathcal{C}} w_k C_k$. Mathematically, the problem can be stated as follows:

$$\min_r \sum_{k \in \mathcal{C}} w_k C_k \tag{P1}$$

$$\text{s.t.} \quad \sum_{k \in \mathcal{C}} \sum_{j \in F_{k,\ell}} r^{k,j}(t) \leq b_\ell, \quad \forall \ell \in \mathcal{L}, \forall t \in \mathcal{T}, \tag{1}$$

$$\int_0^{C_k} r^{k,j}(t) dt \geq V^{k,j}, \quad \forall j \in F_k, \forall k \in \mathcal{C}, \tag{2}$$

where \mathcal{T} is the time horizon. Constraint (1) expresses that, at any instant t , the total rate that port ℓ assigns to flows cannot exceed its capacity b_ℓ . Constraint (2) ensures that the data of flows of each coflow k should be completely transmitted before its completion time C_k .

Here we consider the non-clairvoyant setting in which the size of each flow is unknown to the coflow scheduler and is revealed to it only upon flow completion. Although flow sizes are not known in advance, we assume that they are realized on-line according to given probability distributions. Throughout the paper, flow volumes $V^{k,j}$ are supposed to be statistically independent, with finite first and second moments. The aim is to find a rate-allocation policy that minimizes the weighted average completion time of coflows in expectation, which by linearity of expectation is $\sum_{k \in \mathcal{C}} w_k \mathbb{E}[C_k]$.

A rate allocation policy r specifies the rate allocated to each flow at any given time in such a way that constraint (1) is satisfied. We shall specifically consider work-conserving policies which allocate a constant rate to flows in between decision epochs, where decision epochs correspond to time 0 and to flow completion times. If t is a decision epoch, the decisions of such a policy may only depend on the "past up to time t ", which is given by the sets of data transfers already finished or being performed at t , as well as on the *a priori* knowledge of the input data of the problem. This includes of course the distributional information about remaining sizes of unfinished flows. In other words, admissible policies must be nonanticipative in the sense that scheduling decisions cannot depend on future information, such as the actual volumes of unfinished flows. A simple example of a work-conserving and nonanticipative scheduling policy

²This model abstracts out the datacenter network fabric as one big switch interconnecting servers. The underlying assumption is that the fabric core can sustain 100% throughput and only the ingress and egress ports are potential congestion points.

is the round-robin (RR) policy which allocates rate $r^{k,j}(t) = \min\{b_i/n_i(t), b_o/n_o(t)\}$ to flow $j \in F_k$ at time t , where i and o are the input and output ports of this flow, respectively, and $n_\ell(t)$ is the number of ongoing flows on port ℓ at time t .

Among work-conserving and nonanticipative rate-allocation policies, we restrict ourselves to *order-based policies*. Such policies define a partial order on the set of flows such that one flow precedes another one if they have either the same source or the same destination. Under order-based policies, no single port is transmitting multiple flows at the same time. The analysis of order-based policies is much simpler as it is possible to leverage a well-known reduction to the Concurrent Open Shop problem [17]. Moreover, a key advantage of these policies is that they are easily implementable by leveraging priority mechanisms supported by most transport layers. An important property of order-based policies is that the completion order of the flows (and hence of coflows) on all ports is fixed and independent of flow sizes. In the following, we let \mathcal{R} be the set of order-based policies. Note that the RR policy does not belong to \mathcal{R} as multiple flows may be served in parallel. However, \mathcal{R} contains all priority policies which schedule coflows in some priority order π . Such a policy allocates a positive rate to a flow of coflow $\pi(k')$ if and only if no flows of a coflow $\pi(k)$ such that $k < k'$ is available for transmission on the same input/output port.

The problem that we consider is therefore the following

$$\begin{aligned} \min_r \quad & \sum_{k \in \mathcal{C}} w_k \mathbb{E}[C_k(r)] & (\text{P2}) \\ \text{s.t.} \quad & r \in \mathcal{R}, & (3) \end{aligned}$$

where the random variable $C_k(r)$ represents the completion time of coflow k under order-based policy $r \in \mathcal{R}$. In the following, we denote by C^{OPT} the optimum value for this problem, that is, $C^{OPT} = \inf_{r \in \mathcal{R}} \sum_{k \in \mathcal{C}} w_k \mathbb{E}[C_k(r)]$.

3 Sincronia with expected processing times as inputs

In this section, we establish a generic bound on the approximation ratio of **Sincronia** for problem (P2) when ran with expected processing times as inputs. Following an approach proposed in [21], we first show in Section 3.1 that for any order-based policy r , the vector of expected completion times $\mathbb{E}[C_k(r)]$ satisfies certain inequalities. The latter inequalities yield a polyhedral relaxation of the performance space of order-based scheduling policies, as discussed in Section 3.2. It turns out that this polyhedral relaxation is exactly the same as the one used in **Sincronia** in the clairvoyant setting, except that the actual processing times are replaced by their expected values. In Section 3.3, we briefly describes how **Sincronia** works and propose a priority policy for stochastic coflow scheduling in which rates are allocated to coflows according to the priority order computed by **Sincronia** from expected processing times. We obtain performance guarantees for this policy in Section 3.4.

3.1 Stochastic Parallel Inequalities

Let $P_{\ell,k} = \left(\sum_{j \in F_{k,\ell}} V^{k,j}\right) / b_\ell$ be the total transmission time of coflow k at port ℓ in isolation. In the following, we denote by $\mu_{\ell,k}$ the expected value $\mathbb{E}[P_{\ell,k}]$ of this transmission time.

The so-called parallel inequalities are inequalities on the completion times of jobs which are valid for any feasible schedule on m machines [24, 26]. These inequalities were extended in [21] to stochastic parallel-machine scheduling. We adapt the proof of Theorem 3.1 in [21] to show that the same inequalities are also valid for the stochastic coflow scheduling problem.

Theorem 1. *Let $r \in \mathcal{R}$ be any order-based rate-allocation policy for coflow scheduling. Then*

$$\sum_{k \in A} \mu_{\ell,k} \mathbb{E}[C_k(r)] \geq f_\ell(A) = \frac{1}{2} \left\{ \left(\sum_{k \in A} \mu_{\ell,k} \right)^2 + \sum_{k \in A} \mu_{\ell,k}^2 \right\}, \quad (4)$$

for all $\ell \in \mathcal{L}$ and $A \subseteq \mathcal{C}$.

Proof. See Appendix A. □

3.2 LP Relaxation

It follows from Theorem 1 that for any order-based policy r the expected CCT values $\mathbb{E}[C_k(r)]$ lie in the polyhedron \mathcal{P} defined by inequalities (4). It implies that the following linear program provides a lower bound on C^{OPT} :

$$\text{Minimize } \sum_{k \in \mathcal{C}} w_k C_k \quad (\text{LP-Primal})$$

s.t

$$\sum_{k \in A} \mu_{\ell,k} C_k \geq f_{\ell}(A), \quad \text{for all } \ell \in \mathcal{L} \text{ and } A \subseteq \mathcal{C}, \quad (5)$$

$$C_k \geq 0, \quad \text{for all } k \in \mathcal{C}, \quad (6)$$

where $f_{\ell}(A)$ is defined in (4). We thus have $C^{LP} \leq C^{OPT}$, where C^{LP} is the optimum value of problem LP-Primal.

3.3 Sincronia-based priority policy

In order to schedule a set \mathcal{C} of coflows with known processing times $p_{\ell,k}$, **Sincronia** first computes a priority order over coflows, that is, a permutation π of \mathcal{C} . If $\pi(k) < \pi(k')$, then the flows of coflow k are transmitted in the data centre network with strict priority over the flows of coflow k' . **Sincronia** then uses a preemptive greedy rate allocation which assigns to a single flow the full rate of its incoming and outgoing ports at any given point of time and preserves this ordering, that is, it is guaranteed that coflows will be completed in the priority order (that is, $C_k \leq C_{k'}$ if $\pi(k) < \pi(k')$).

It turns out that problem (LP-Primal) is exactly the same as the one used in **Sincronia** to compute the priority order π [3], except that the actual processing times $p_{\ell,k}$ have been replaced by their expected values $\mu_{\ell,k}$. A natural approach to obtain an admissible scheduling policy in the stochastic setting is therefore to compute a priority order π using **Sincronia** with expected processing times $\mu_{\ell,k}$, instead of the unknown actual processing times $P_{\ell,k}$. The greedy rate allocation of **Sincronia** then gives an admissible rate-allocation policy for problem (P2).

Before establishing performance guarantees for the above **Sincronia**-based priority policy, we first briefly describe how **Sincronia** computes the priority order π from expected processing times $\mu_{\ell,k}$. **Sincronia** uses a primal-dual algorithm in which the primal problem is (LP-Primal), and the dual problem is as follows

$$\text{Maximize } \sum_{\ell \in \mathcal{L}} \sum_{A \subseteq \mathcal{C}} f_{\ell}(A) y_{\ell,A} \quad (\text{LP-Dual})$$

s.t

$$\sum_{A: k \in A} \sum_{\ell \in \mathcal{L}} \mu_{\ell,k} y_{\ell,A} \leq w_k, \quad \text{for all } k \in \mathcal{C}, \quad (7)$$

$$y_{\ell,A} \geq 0, \quad \text{for all } \ell \in \mathcal{L} \text{ and } A \subseteq \mathcal{C}. \quad (8)$$

The **Sincronia** primal-dual algorithm is shown in Algorithm 1. The dual variables $y_{\ell,A}$ are initialized to 0 and the set A of unscheduled coflows is initially set to \mathcal{C} . The algorithm assigns priorities to coflows in the order of increasing priority. At each iteration t , the algorithm determines the bottleneck port b , that is, the port ℓ for which the total expected completion time $\sum_{k \in A} \mu_{\ell,k}$ is maximum. It then determines the coflow k^* with the largest weighted expected processing time on the bottleneck. Coflow k^* is assigned priority t with $\pi(k^*) = t$ and it

is removed from the set of unscheduled coflows A . Before that, the dual variable $y_{b,A}$ and the primal variable C_{k^*} are set to their final values and the weights of the coflows using the bottleneck port are updated. As discussed in in [3], the complexity of Algorithm 1 is $O(n^2)$.

Algorithm 1 Sincronia primal-dual algorithm with expected processing times $\mu_{\ell,k}$ as inputs

```

1:  $y_{\ell,A} \leftarrow 0$  for all  $\ell \in \mathcal{L}$  and  $A \subseteq \mathcal{C}$ .
2:  $A \leftarrow \mathcal{C}$ 
3: for  $t = n \dots 1$  do
4:    $b \leftarrow \operatorname{argmax}_{\ell \in \mathcal{L}} \sum_{k \in A} \mu_{\ell,k}$  ▷ Bottleneck port
5:    $k^* \leftarrow \operatorname{argmin}_{k \in A} (w_k / \mu_{b,k})$  ▷ Largest weighted proc. time
6:    $C_{k^*} \leftarrow \sum_{k \in A} \mu_{b,k}$  ▷ Set primal variable
7:    $y_{b,A} \leftarrow w_{k^*} / \mu_{b,k^*}$  ▷ Set dual variable
8:    $w_k \leftarrow w_k - w_{k^*} \mu_{b,k} / \mu_{b,k^*}$  for all  $k \in A$  ▷ Update weights
9:    $\pi(t) \leftarrow k^*$  ▷ Set priority of coflow  $k^*$ 
10:   $A \leftarrow A \setminus \{k^*\}$  ▷ Remove  $k^*$  from unscheduled coflows
11: end for

```

The approximation ratio of Algorithm 1 is established in [3] and formally stated in Theorem 2 below.

Theorem 2. *Algorithm 1 provides a feasible solution to problem LP-Primal such that $C_{\pi(t)}^{SIN} = \max_{\ell} \sum_{\tau=1}^t \mu_{\ell,\pi(\tau)}$ for $t = 1, 2, \dots, n$. Moreover, the cost of this solution is at most two times C^{LP} , that is,*

$$\sum_{t=1}^n w_{\pi(t)} \max_{\ell} \sum_{\tau=1}^t \mu_{\ell,\pi(\tau)} \leq 2 C^{LP} \leq 2 C^{OPT}. \quad (9)$$

3.4 Performance guarantees for the Sincronia-based priority policy

Let π be the priority order computed by Algorithm 1 from expected processing times. We now consider the greedy rate allocation of Sincronia (see Algorithm 2 in [3]). We let C_k be the random completion time of coflow k under the greedy rate allocation when coflows are processed in the priority order π . It is proven in [3] that $c_{\pi(t)} \leq 2 \max_{\ell} \sum_{\tau=1}^t p_{\ell,\pi(\tau)}$ for any fixed realization of flow volumes, so that

$$\mathbb{E} [C_{\pi(t)}] \leq 2 \mathbb{E} \left[\max_{\ell} \sum_{\tau=1}^t P_{\ell,\pi(\tau)} \right] \text{ for all } t = 1, 2, \dots, n. \quad (10)$$

In the following, in order to combine inequalities (9) and (10) so as to obtain an upper bound on the approximation ratio of Sincronia, we define the quantity α as follows

$$\alpha = \max_{t=1,\dots,n} \frac{\mathbb{E} \left[\max_{\ell} \sum_{\tau=1}^t P_{\ell,\pi(\tau)} \right]}{\max_{\ell} \left(\sum_{\tau=1}^t \mu_{\ell,\pi(\tau)} \right)}. \quad (11)$$

Theorem 3. *The weighted average CCT of the Sincronia-based priority policy is upper bounded by $4 \times \alpha$ times the optimal one.*

Proof. We have

$$\sum_{t=1}^n w_{\pi(t)} \mathbb{E} [C_{\pi(t)}] \leq 2 \times \sum_{t=1}^n w_{\pi(t)} \mathbb{E} \left[\max_{\ell} \sum_{\tau=1}^t P_{\ell,\pi(\tau)} \right], \quad (12)$$

$$\leq 2\alpha \times \sum_{t=1}^n w_{\pi(t)} \max_{\ell} \left(\sum_{\tau=1}^t \mu_{\ell,\pi(\tau)} \right), \quad (13)$$

$$\leq 4\alpha \times C^{OPT}, \quad (14)$$

where the first inequality follows from (10), the second one follows from (11) and the last one follows from (9). \square

In the next section, we derive upper bounds on the quantity α in (11). In view of Theorem 3, performance guarantees for the Sincronia-based priority policy directly follow from these bounds.

4 Upper bounds on α

We first give in Section 4.1 an upper bound on α that holds for any distribution of the processing times. Distribution-dependent upper bounds on α are derived in Section 4.3 for some specific distributions.

4.1 Distribution-independent upper bound

For a given priority order θ , let the matrix $\mathbf{P}_k^\theta = [P_{\ell,j}]$ denote the $L_k \times k$ matrix with the processing times of flows ahead of coflow k on the L_k links used by coflow k in the order θ . For ease of notation, we use k instead of $\theta(k)$. It will be apparent from the context that \mathbf{P}_k^θ denotes the \mathbf{P} matrix for the k th coflow in the order θ which is in fact $\theta(k)$. Note that the matrix \mathbf{P}_k^θ can be written as $\mathbf{P}_k^\theta = \mathbf{M}_k^\theta + \mathbf{X}_k^\theta$, where $\mathbf{M}_k^\theta = \mathbb{E}\mathbf{P}_k^\theta$ and \mathbf{X}_k^θ is a zero-mean matrix. We observe from (11) that

$$\alpha = \max_k \frac{\mathbb{E}\|\mathbf{M}_k^\pi + \mathbf{X}_k^\pi\|}{\|\mathbf{M}_k^\pi\|}, \quad (15)$$

where $\|\cdot\|$ is the induced ∞ -norm and π is the order computed by Sincronia. An upper bound for α can be obtained by taking the worst-case order and using the triangular inequality for norms,

$$\alpha \leq \max_\theta \max_k \frac{\mathbb{E}\|\mathbf{M}_k^\theta + \mathbf{X}_k^\theta\|}{\|\mathbf{M}_k^\theta\|} \leq 1 + \max_\theta \max_k \frac{\mathbb{E}\|\mathbf{X}_k^\theta\|}{\|\mathbf{M}_k^\theta\|}. \quad (16)$$

Thus, computing an upper-bound on α can be reduced to the problem of computing the expectation of the ∞ -norm of random matrices whose elements are taken from the given distribution. These bounds will depend on the probability distribution of $X_{\ell,j}$ s. For a finite k but L going to ∞ , one can use extreme-value theory [22] which models how the maximum of a certain number of random variables behaves. When both k and L go to ∞ , we refer the reader to [1] for results on the asymptotic bounds on α . Below, we give an upper bound on α that is non-asymptotic but less sharp and which only depends upon the existence of certain moments.

Let $\mu_{min} = \min\{\mu_{\ell,j} : \mu_{\ell,j} > 0\}$ and let $p \in (1, \infty)$ be such that $m_p := \max_{\ell,j} E|X_{\ell,j}|^p < \infty$. That is, μ_{min} is the smallest expected flow-size among those flows that are not always 0, and m_p is the largest p th moment of the flow-sizes provided that it is finite.

Theorem 4 (Upper bound). *Let $\mu_{min} = \min_{\ell,j} \mu_{\ell,j}$ and let $m_p := \max_{\ell,j} E|X_{\ell,j}|^p < \infty$ for $p \in (1, \infty)$. Then,*

$$\alpha \leq 1 + L^{1/p} \frac{m_p^{1/p}}{\mu_{min}}.$$

This result says that if $P_{\ell,j}$ have a finite p th moment³, then the upper bound on approximation ratio is of the order of $L^{1/p}$. This bound is independent of the distribution and just requires a check on whether a certain moment exists or not. This bound is useful for distribution for which it is not straightforward to obtain a bound from (16) or by using extremal value theory.

For the proof of this theorem, we need the following simple intermediate result giving an upper bound on the second moment of the sum of absolute values of random variables in terms of an upper bound on the second moments of the individual random variables.

³Note that $P_{\ell,j}$ having a finite p th moment is equivalent to $|X_{\ell,j}|$ having a finite p th moment. So, to invoke Theorem 4 it is sufficient to perform the check on $P_{\ell,j}$.

Lemma 1. Let Y_i , $i = 1, 2, \dots, n$ be independent random variables with $E|Y_i|^p \leq m_p, \forall i$. Then

$$\mathbb{E} \left(\sum_i |Y_i| \right)^p \leq n^p m_p. \quad (17)$$

Proof. The result is a consequence of the Cauchy-Schwartz inequality which applied on $\sum_i |Y_i|$ gives

$$\left(\sum_i |Y_i| \right)^p \leq n^{p-1} \left(\sum_i |Y_i|^p \right)$$

Now taking expectation on both sides and using the linearity of the expectation operator and the fact that $\mathbb{E}|Y_i|^p \leq m_p$, we obtain the claimed result. \square

That the above bound is tight follows by taking $Y_i = \pm 1$ with probability 0.5.

Proof of Theorem 4. Let θ be an arbitrary order. We denote by $\mu_{\ell,j}^\theta$ the quantity $\mu_{\ell,\theta(j)}$. Using the standard inequality that the ∞ -norm is a lower bound on the p -norm and the fact that μ_{\min} is a lower bound on the means, we get

$$\frac{\mathbb{E} \|\mathbf{X}_k^\theta\|}{\|\mathbf{M}_k^\theta\|} = \frac{\mathbb{E} \max_\ell \sum_{j=1}^k |X_{\ell,j}^\theta|}{\max_\ell \sum_{j=1}^k \mu_{\ell,j}^\theta} \quad (18)$$

$$\leq \frac{\mathbb{E} \left(\sum_\ell \left(\sum_{j=1}^k |X_{\ell,j}^\theta| \right)^p \right)^{1/p}}{k \mu_{\min}} \quad (19)$$

$$\leq \frac{\left(\sum_\ell \mathbb{E} \left(\sum_{j=1}^k |X_{\ell,j}^\theta| \right)^p \right)^{1/p}}{k \mu_{\min}} \quad (20)$$

$$\leq \frac{(\sum_\ell k^p m_p)^p}{k \mu_{\min}} = L^{1/p} \frac{m_p^{1/p}}{\mu_{\min}}, \quad (21)$$

where (20) follows from Jensen's inequality and linearity of expectation, and (21) follows from Lemma 1 and the fact that $E|X_{\ell,j}|^p = m_p$. The claimed results then follows from 16 since the right-hand side of (21) is independent of θ and k . \square

Corollary 1 ($p = 2$). For $p = 2$, the upper bound becomes

$$\alpha \leq 1 + L^{1/2} \frac{\sigma_{\max}}{\mu_{\min}} \quad (22)$$

where $\sigma_{\max} = \max_{\ell,j} \sigma_{\ell,j}$ is the maximum standard deviation of the flow sizes $P_{\ell,j}$.

This bound in Theorem 4 is consistent with **Sincronia** in the sense that if the flow sizes are deterministic, then $|X|^p = 0$ for all $p > 1$, and we recover the approximation factor of **Sincronia** in the clairvoyant setting.

4.2 Asymptotically matching asymptotic lower bound

To show that the upper bound in Theorem 4 is asymptotically tight, we shall assume that $P_{\ell,j}$ are Pareto random variables and give asymptotic bounds as $L_k \rightarrow \infty$.

Let k be finite and $P_{\ell,j}$ be independent and identically distributed random variables with shape parameter $\zeta > 1$. This last condition is required for the mean of these random variables to be finite. We will need the fact that for the p th moment of $P_{\ell,j}$ to exist (and hence be able to use the result of Theorem 4), ζ has to be strictly larger than p .

A Pareto random variable of shape parameter ζ belongs to the class of regularly varying random variables of index ζ (see [20] and references therein for the definition and properties of the class of regularly varying random variables).

From the closure property of regularly varying random variables [20], $\sum_{j=1}^k P_{\ell,j}$ is again a regularly varying random variable of index ζ (which is the shape parameter of the constituent Pareto random variables). Further, both $P_{\ell,j}$ and $\sum_{j=1}^k P_{\ell,j}$ belong to the maximum domain of attraction of the Fréchet distribution of index ζ .

For this class of distributions, the maximum grows as $L^{1/\zeta}$ [20], that is,

$$\max_{\ell} \frac{\sum_{j=1}^k P_{\ell,j}}{\sum_{j=1}^k \mu_{\ell,j}} = \Theta(L^{1/\zeta}).$$

For a Pareto random variable with finite p th moment, ζ has to be greater than p . To get close to the upper bound in Theorem 4, one can take then take an $\zeta > p$ but as close to p as desired. This results in a small gap to the upper bound in Theorem 4 since we cannot set $\zeta = p$. We do not know how to bridge this gap either by improving the upper bound or by exhibiting a sequence of random variables with finite p th moment that have a growth rate of exactly $L^{1/p}$. The finite p th moment requirement is needed since the bound of Theorem 4 is obtained under that assumption. To keep the comparison fair, we need to take ζ strictly larger than p and incur a small gap in the two bounds.

4.3 Distribution-dependent tighter upper bounds

The bounds on α in Sections 4.1 and 4.2 are proportional to $L^{1/p}$ and only require the existence of the p th moment of the flow processing times. In practice, the processing times could be better behaved with existence of exponential moments. We now investigate how tighter bounds can be obtained when exponential moments (rather than just the p th moment) of $P_{\ell,j}$ are finite.

Gamma distributed processing times

We start with distributions with exponential tails, that is, $\mathbb{P}(X > t) \sim \exp(-ct)$. For these distributions, the moment generating function has a non-zero radius of convergence. We focus on gamma-distributed random variables that can also be seen as a sum of a finite number of exponentially distributed random variables.

Lemma 2. *Let $P_{\ell,j}$ be independent and follow a gamma distribution with scale parameter $\beta_{\ell,j}$ and shape parameter $s_{\ell,j}$. Then,*

$$\alpha \leq c + \log\left(c + \sqrt{c^2 - 1}\right). \quad (23)$$

where $c = 1 + \gamma \log(L)$, $\gamma = \left(\frac{\sigma_{max}}{\mu_{min}}\right)^2$ with $\sigma_{max}^2 = \max_{\ell,j} s_{\ell,j} \beta_{\ell,j}^2$ is the maximum variance of the flows and $\mu_{min} = \min_{\ell,j} s_{\ell,j} \beta_{\ell,j}$ is the smallest non-zero mean of flow-size.

Proof. See appendix B □

Combining Theorem 3 and Lemma 2, we have the following result regarding the approximation ratio of **Sincronia** for gamma-distributed processing times.

Corollary 2. *If all random variables $P_{\ell,j}$ are independent and gamma-distributed, the approximation of **Sincronia** with expected processing times as inputs is $4(c + \log(c + \sqrt{c^2 - 1}))$.*

Remark 1. *This bound on α does not directly depend upon CV_{max} . However, if $CV_{max} = 0$, then $\gamma = 1$ which leads to $\alpha = 1$. We thus recover the approximation ratio of 4 of **Sincronia** in the clairvoyant setting.*

Note in particular that for large number of ports, that is large L , this upper bound is roughly $\Theta(\log(L))$. The upper bound therefore increases only slowly with the number of ports in contrast to the $\Theta(L^{1/p})$ growth for the Pareto distribution in the previous subsection.

Normally distributed processing times

We now obtain an upper bound on α when the processing times $P_{\ell,j}$ are normally distributed.

Lemma 3. *If all random variables $P_{\ell,j}$ are independent and follow a normal distribution with mean $\mu_{\ell,j}$ and coefficient of variation $CV_{\ell,j}$, then $\alpha \leq 1 + \sqrt{2\log(L)}CV_{max}$, where $L = |\mathcal{L}|$ is the number of ports and $CV_{max} = \max_{\ell,j} CV_{\ell,j}$.*

Proof. See appendix C □

As a direct consequence of Theorem 3 and Lemma 3, we have the following result regarding the approximation ratio of Sincronia for normally distributed processing times.

Corollary 3. *If all random variables $P_{\ell,j}$ are independent and follow a normal distribution, the approximation ratio of the Sincronia-based priority policy is $4 \left\{ 1 + \sqrt{2\log(L)}CV_{max} \right\}$.*

Note in particular that when $CV_{max} = 0$, we retrieve the approximation ratio of Sincronia in the clairvoyant setting. Further, the upper bound only grows as $\Theta(\sqrt{\log(L)})$ which is even slower than the $\Theta(\log(L))$ growth for the gamma distribution. This is to be expected since the normal distribution has a lighter tail than the gamma distribution.

5 Numerical Results

In this section, we compare the average CCTs obtained with the Sincronia-based priority policy against that of the clairvoyant Sincronia algorithm and the lower bound C^{LP} on random problem instances. We also include in the comparison the round robin scheduling policy and a random order policy.

The random instances are generated using a coflow workload generator based on a real-world trace with 526 coflows. The trace was obtained from a one-hour long run of Hive/MapReduce jobs on a 3000-machine cluster with 150 racks at Facebook [12]. The workload generator allows upsampling the Facebook trace to desired number of ports L and coflows n , while still keeping workload characteristics similar to the original Facebook trace. The original workload generator, which is publicly available [2], was slightly modified in order to be able to generate flow sizes following a target probability distribution (e.g., a normal or Gamma distribution with specified mean and standard deviation).

In order to compare the performance of scheduling algorithms under various conditions, we generate multiple configurations by varying the following parameters:

- The number of ports is always assumed to be even (the first $L/2$ ports are input ports, while the other ports are output ports) and can take values between 8 and 80.
- The number of coflows n is either $n = L/2$, $n = L$ or $n = 3L/2$.
- When the flow sizes follow a specific target distribution, the mean flow size is 100 and the coefficient of variation CV is either $\frac{1}{2}$, 1 or 2. Otherwise, flow sizes are drawn from the empirical distribution of the Facebook trace.
- The weights w_k of all coflows are set to 1.

For each possible configurations of these parameters, we randomly generate problem instances. Note that in the non-clairvoyant setting, a problem instance specifies the set of ports and their capacities as well as the set of coflows, with the respective flows of each coflow, but does not specify the actual flow sizes which are unknown to a non-clairvoyant coflow scheduler. We therefore apply the following procedure for evaluating the performance of coflow scheduling algorithms:

1. We first generate 100 problem instances with the specified number of ports and number of coflows by randomly generating the set of flows F_k of each coflow k . These problem instances do not specify the actual flow sizes and thus represents the information available to a non-clairvoyant coflow scheduler.

L	n	UB	CL				NC				RR			
			mean	std	Q1	Q3	mean	std	Q1	Q3	mean	std	Q1	Q3
4	4	8.6	1.01	0.06	0.97	1.06	1.08	0.06	1.02	1.11	1.55	0.12	1.48	1.67
8	4	10	1.07	0.06	1.03	1.11	1.11	0.08	1.06	1.14	1.71	0.19	1.60	1.83
8	8	10	1.05	0.05	1.01	1.08	1.09	0.05	1.06	1.13	2.08	0.17	2.01	2.19
16	8	11.2	1.09	0.04	1.06	1.12	1.11	0.05	1.08	1.15	2.23	0.18	2.10	2.34
16	16	11.2	1.09	0.03	1.06	1.11	1.12	0.05	1.08	1.15	2.80	0.20	2.66	2.94
32	16	12.4	1.10	0.03	1.08	1.13	1.12	0.04	1.08	1.15	2.92	0.24	2.78	3.03

Table 1: C^{ALG}/C^{LP} for flow sizes following a Gamma distribution with $CV = 0.5$.

- For each problem instance, we randomly generate 1000 realizations of flow sizes according to the specified random distribution. Given a scheduling algorithm, the average CCTs obtained with these traffic realizations are averaged to estimate the expected average CCT of the algorithm over the corresponding problem instance.

5.1 Small instances

We first consider small instances for which the lower bound C^{LP} can be computed using Gurobi as linear-programming solver. This lower bound is directly computed for each of the 100 problem instances as this does not require the actual flow sizes, but only their mean values. For each problem instance, we also estimate the expected average CCT C^{ALG} obtained with the clairvoyant Sincronia (CL), with the non-clairvoyant Sincronia-based priority policy (NC), with round-robin (RR) and with a random order policy (RO) by averaging over the 1000 traffic realizations for this problem instance.

We present below the ratios C^{ALG}/C^{LP} obtained for three different random distributions of flow sizes (Gamma, Normal and Pareto) as well as for flow sizes following the empirical flow size distribution in the Facebook trace. These ratios provide upper bounds on the empirical approximation ratios of the considered algorithms.

Gamma distribution

Let us first consider the results obtained with a Gamma distribution for flow volumes. We consider three different coefficients of variations: $CV = \frac{1}{2}$, $CV = 1$ and $CV = 2$.

The mean, standard deviation as well as the first and third quartiles of the ratios C^{ALG}/C^{LP} are reported in Tables 1, 2 and 3 for flow sizes following a Gamma distribution with $CV = 0.5, 1$, and 2, respectively. Column UB in these tables shows the theoretical upper bound on C^{NC}/C^{LP} computed using (23). For $CV = 0.5$, we observe that for all considered configurations of the numbers of ports and coflows, non-clairvoyant Sincronia performs almost as well as the clairvoyant one. We also note that the variability of the ratio C^{ALG}/C^{LP} is also quite low for both algorithms, indicating that they perform very well for the vast majority of our random problem instances. We also observe that the ratio C^{RR}/C^{LP} has a much higher mean and standard deviation, which shows that the round-robin allocation performs poorly as compared to our NC priority policy when the variability of flow sizes is low. When the variability of the flow sizes increases (i.e., $CV = 1$ or $CV = 2$), the performance of the non-clairvoyant Sincronia degrades as compared to the clairvoyant one. The variability of the ratio C^{NC}/C^{LP} is also much higher than that observed in the case $CV = 0.5$. Nevertheless, the performance of NC remains fairly close to that of CL, and much better than that obtained with RR. For all three values of CV , the upper bound UB is pessimistic with values much bigger than 4 while the empirical ratios are between 1 and 2.

For a more detailed analysis, we also plot the histograms of the ratios C^{ALG}/C^{LP} obtained for the 100 random problem instances in the case $L = 32$ and $n = 16$ in Figure 1a and Figure 1b for $CV = 0.5$ and $CV = 2$, respectively. For $CV = 0.5$, clearly the histograms obtained for $ALG = CL$ and $ALG = NC$ are very similar. The values obtained with RR are much higher, and even higher than those obtained with a random priority order. For $CV = 2$, as compared

L	n	UB	CL				NC				RR			
			mean	std	Q1	Q3	mean	std	Q1	Q3	mean	std	Q1	Q3
4	4	15.6	0.96	0.09	0.90	1.02	1.14	0.09	1.09	1.20	1.38	0.15	1.28	1.49
8	4	19.5	1.10	0.10	1.03	1.18	1.20	0.11	1.13	1.28	1.63	0.21	1.50	1.78
8	8	19.5	1.04	0.08	1.00	1.10	1.18	0.08	1.14	1.23	1.90	0.19	1.77	2.03
16	8	23.1	1.13	0.06	1.10	1.18	1.21	0.07	1.17	1.26	2.17	0.19	2.05	2.29
16	16	23.1	1.12	0.04	1.09	1.15	1.20	0.06	1.16	1.24	2.65	0.18	2.51	2.76
32	16	26.6	1.14	0.04	1.11	1.18	1.18	0.05	1.15	1.22	2.84	0.22	2.70	2.95

Table 2: C^{ALG}/C^{LP} for flow sizes following a Gamma distribution with $CV = 1$.

L	n	UB	CL				NC				RR			
			mean	std	Q1	Q3	mean	std	Q1	Q3	mean	std	Q1	Q3
4	4	36.4	0.89	0.11	0.81	0.98	1.28	0.13	1.21	1.36	1.14	0.17	1.02	1.27
8	4	49.0	1.18	0.18	1.03	1.31	1.43	0.17	1.33	1.52	1.54	0.27	1.35	1.74
8	8	49.0	1.05	0.12	0.96	1.13	1.41	0.12	1.32	1.50	1.63	0.23	1.47	1.80
16	8	61.1	1.26	0.10	1.19	1.33	1.49	0.10	1.41	1.55	2.09	0.23	1.93	2.25
16	16	61.1	1.20	0.07	1.16	1.24	1.46	0.08	1.40	1.51	2.40	0.19	2.27	2.55
32	16	73.0	1.27	0.07	1.21	1.32	1.40	0.08	1.35	1.45	2.76	0.22	2.60	2.89

Table 3: C^{ALG}/C^{LP} for flow sizes following a Gamma distribution with $CV = 2$.

to the histogram for CL , the most probable values for NC are slightly shifted to the right by $+0.15$. As in the case $CV = 0.5$, the values obtained with RR are much higher, and even higher than those obtained with a random priority order.

Figure 2 summarizes the above results. It shows the average values (over 100 random instances) obtained for CL, NC, and RR. We first note that CL provides near-optimal results in all cases. For $CV = 0.5$, NC performs almost as well as CL and much better than RR. For $CV = 2$, there is a noticeable performance degradation as compared to CL, but the proposed algorithm still outperforms significantly RR. We also note that the gap between CL and NC decreases as the number of ports increases.

Normally distributed flow volumes

We now consider the results obtained for normally-distributed flow volumes. As before, we consider three different coefficients of variations: $CV = \frac{1}{2}$, $CV = 1$ and $CV = 2$. The mean, standard deviation as well as the first and third quartiles of the ratios C^{ALG}/C^{LP} are reported in Tables 4, 5 and 6 for $CV = \frac{1}{2}$, $CV = 1$ and $CV = 2$, respectively. Column UB in these tables is the theoretical upper bound on C^{NC}/C^{LP} computed using Cor. 3.

As for the gamma distribution, UB is quite pessimistic compared to the observed ratios. Regarding C^{ALG}/C^{LP} , we observe that when the variability of flow sizes is low the NC performs almost as well as CL and provides almost optimal solutions. We also observe that the variability of the ratio C^{NC}/C^{LP} is quite low in all configurations and that NC provides much better results than RR. These observations are confirmed by the histograms shown in Figure 3a which report the ratios C^{ALG}/C^{LP} obtained for the 100 random problem instances in the case $L = 32$ and $n = 16$. Although the repartition is slightly different, CL and NC provide similar results, which are much better than those obtained with RR or RO.

When the flow size variability increases, we observe, as in the case of the gamma distribution, a slight performance degradation of NC as compared to CL. Nevertheless, the results obtained with NC are much better than those obtained with RR. This can also be observed in the histograms shown in Figure 3b for 32 ports and 16 coflows. We observe that in this configuration, a random-order policy is much more efficient than RR.

Figure 4 shows the average values (over 100 random instances) obtained for CL, NC and RR. The main observations are similar to those made for a gamma distribution of flow sizes.




Figure 1: Histograms of the ratios C^{ALG}/C^{LP} for CL, NC, RR and a random-order policy (RO) for 100 random problem instances with $L = 32$ ports and $n = 16$ coflows. Flow sizes follow a Gamma distribution with either (a) $CV = 0.5$ or (b) $CV = 2$.

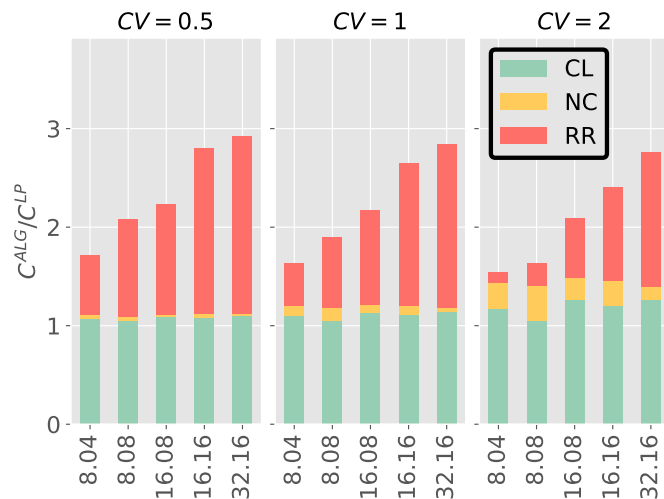


Figure 2: Average values of C^{ALG}/C^{LP} for flow sizes following a Gamma distribution with $CV = 0.5$ (left), $CV = 1$ (middle) or $CV = 2$ (right). The number of ports L and the number n of coflows are indicated on the x -axis as $L.n$.

L	n	UB	CL				NC				RR			
			mean	std	Q1	Q3	mean	std	Q1	Q3	mean	std	Q1	Q3
4	4	7.3	1.02	0.06	0.98	1.08	1.09	0.06	1.04	1.12	1.58	0.13	1.51	1.69
8	4	8.1	1.08	0.06	1.04	1.13	1.12	0.08	1.06	1.15	1.74	0.19	1.62	1.86
8	8	8.1	1.05	0.05	1.02	1.09	1.10	0.05	1.06	1.13	2.12	0.18	2.03	2.23
16	8	8.7	1.09	0.04	1.07	1.12	1.12	0.05	1.08	1.15	2.27	0.18	2.14	2.36
16	16	8.7	1.09	0.03	1.07	1.11	1.12	0.05	1.08	1.15	2.84	0.20	2.70	2.98
32	16	9.3	1.11	0.04	1.08	1.13	1.12	0.04	1.08	1.15	2.95	0.24	2.80	3.06

Table 4: C^{ALG}/C^{LP} for flow sizes following a Normal distribution with $CV = 0.5$.

L	n	UB	CL				NC				RR			
			mean	std	Q1	Q3	mean	std	Q1	Q3	mean	std	Q1	Q3
4	4	10.7	1.08	0.09	1.01	1.13	1.22	0.08	1.17	1.27	1.58	0.17	1.50	1.69
8	4	12.2	1.20	0.09	1.13	1.25	1.27	0.10	1.19	1.33	1.83	0.22	1.69	1.98
8	8	12.2	1.14	0.07	1.10	1.19	1.24	0.07	1.20	1.29	2.16	0.21	2.05	2.30
16	8	13.4	1.21	0.05	1.17	1.25	1.26	0.07	1.22	1.30	2.40	0.20	2.27	2.53
16	16	13.4	1.20	0.04	1.17	1.22	1.26	0.06	1.22	1.29	2.98	0.21	2.83	3.13
32	16	14.5	1.22	0.04	1.19	1.25	1.25	0.05	1.21	1.28	3.15	0.25	3.00	3.27

Table 5: C^{ALG}/C^{LP} for flow sizes following a Normal distribution with $CV = 1$.

L	n	UB	CL				NC				RR			
			mean	std	Q1	Q3	mean	std	Q1	Q3	mean	std	Q1	Q3
4	4	17.3	1.37	0.14	1.25	1.46	1.66	0.13	1.57	1.75	1.92	0.24	1.77	2.09
8	4	20.3	1.57	0.15	1.47	1.69	1.72	0.16	1.63	1.82	2.30	0.31	2.12	2.52
8	8	20.3	1.48	0.11	1.41	1.56	1.68	0.11	1.61	1.76	2.65	0.29	2.45	2.87
16	8	22.8	1.60	0.08	1.55	1.66	1.71	0.10	1.65	1.78	3.07	0.27	2.91	3.26
16	16	22.8	1.57	0.07	1.53	1.62	1.70	0.09	1.64	1.76	3.74	0.26	3.56	3.91
32	16	25.1	1.60	0.06	1.55	1.65	1.66	0.07	1.61	1.71	4.01	0.31	3.81	4.17

Table 6: C^{ALG}/C^{LP} for flow sizes following a Normal distribution with $CV = 2$.




Figure 3: Histograms of the ratios C^{ALG}/C^{LP} for CL, NC, RR and RO for 100 random problem instances with $L = 32$ ports and $n = 16$ coflows. Flow sizes follow a Normal distribution with either (a) $CV = 0.5$ or (b) $CV = 2$.




Figure 4: Average values of C^{ALG}/C^{LP} for flow sizes following a Normal distribution with $CV = 0.5$ (left), $CV = 1$ (middle) or $CV = 2$ (right). The number of ports L and the number n of coflows are indicated on the x -axis as $L.n$.

L	n	UB	CL				NC				RR			
			mean	std	Q1	Q3	mean	std	Q1	Q3	mean	std	Q1	Q3
4	4	8.0	1.02	0.05	0.97	1.06	1.06	0.05	1.03	1.10	1.59	0.11	1.54	1.69
8	4	9.7	1.07	0.05	1.03	1.11	1.10	0.08	1.05	1.13	1.72	0.19	1.60	1.85
8	8	9.7	1.04	0.04	1.01	1.07	1.09	0.05	1.05	1.12	2.11	0.17	2.02	2.21
16	8	12.0	1.09	0.04	1.06	1.11	1.11	0.05	1.08	1.15	2.23	0.18	2.11	2.34
16	16	12.0	1.08	0.03	1.06	1.10	1.12	0.05	1.08	1.15	2.79	0.21	2.65	2.94
32	16	15.3	1.10	0.03	1.08	1.13	1.12	0.04	1.08	1.15	2.90	0.24	2.75	3.02

Table 7: C^{ALG}/C^{LP} for flow sizes following a Pareto distribution with $CV = 0.5$.

Pareto distribution

The statistical results obtained for a Pareto distribution of flow sizes are reported in Tables 7, 8 and 9 for $CV = 0.5$, $CV = 1$ and $CV = 2$, respectively. Figures 5a and 5b show the statistical distributions of the ratios C^{ALG}/C^{LP} obtained for $L = 32$ ports, $n = 16$ coflows and $CV = 0.5$ or $CV = 2$, respectively. Column UB in these tables is the theoretical upper bound on C^{NC}/C^{LP} computed using (22).

As for the gamma and normal distributions, we observe that UB is pessimistic, and that NC is near-optimal and performs almost as well as CL when the variability of flow sizes is low, and remains quite close to it when the variability of flow sizes increases. In all considered configurations, NC outperforms RR, which in most cases is worst than a random-order policy.

Figure 6 shows the average values (over 100 random instances) obtained for CL, NC and RR. The main observations are similar to those made for gamma and normal distributions of

L	n	UB	CL				NC				RR			
			mean	std	Q1	Q3	mean	std	Q1	Q3	mean	std	Q1	Q3
4	4	12.0	0.99	0.05	0.94	1.03	1.08	0.06	1.03	1.11	1.52	0.11	1.47	1.62
8	4	15.3	1.08	0.06	1.03	1.13	1.14	0.09	1.08	1.19	1.69	0.20	1.56	1.82
8	8	15.3	1.04	0.05	1.01	1.08	1.13	0.06	1.10	1.18	2.03	0.17	1.96	2.13
16	8	20.0	1.10	0.04	1.08	1.13	1.17	0.06	1.13	1.21	2.19	0.18	2.05	2.29
16	16	20.0	1.09	0.04	1.07	1.11	1.17	0.06	1.13	1.21	2.71	0.20	2.57	2.85
32	16	26.6	1.12	0.04	1.09	1.15	1.18	0.05	1.14	1.21	2.85	0.24	2.69	2.96

Table 8: C^{ALG}/C^{LP} for flow sizes following a Pareto distribution with $CV = 1$.

L	n	UB		CL				NC				RR			
		mean	std	Q1	Q3	mean	std	Q1	Q3	mean	std	Q1	Q3		
4	4	20.0	0.98	0.06	0.94	1.02	1.10	0.07	1.06	1.13	1.48	0.11	1.41	1.58	
8	4	26.6	1.08	0.07	1.03	1.13	1.17	0.10	1.11	1.22	1.66	0.20	1.53	1.79	
8	8	26.6	1.03	0.06	1.01	1.07	1.17	0.08	1.12	1.22	1.98	0.17	1.90	2.08	
16	8	36.0	1.12	0.05	1.09	1.16	1.23	0.07	1.19	1.28	2.17	0.18	2.03	2.28	
16	16	36.0	1.10	0.04	1.08	1.13	1.24	0.06	1.19	1.27	2.66	0.20	2.53	2.80	
32	16	49.2	1.14	0.05	1.11	1.18	1.24	0.06	1.21	1.29	2.82	0.24	2.67	2.93	

Table 9: C^{ALG}/C^{LP} for flow sizes following a Pareto distribution with $CV = 2$.




Figure 5: Histograms of the ratios C^{ALG}/C^{LP} for CL, NC, RR and RO for 100 random problem instances with $L = 32$ ports and $n = 16$ coflows. Flow sizes follow a Pareto distribution with either (a) $CV = 0.5$ or (b) $CV = 2$.

flow sizes.

Facebook trace

We now present the results obtained with a flow size distribution derived from the empirical distribution of flow sizes in the Facebook trace. The statistical results obtained in this case are reported in Tables 10, and we show the specific distributions of the ratios C^{ALG}/C^{LP} obtained for $L = 32$ and $n = 16$ in Figure 7. Figure 8 shows the average values (over 100 random instances) obtained for CL, NC and RR.

Those results are completely different from those obtained for synthetic distribution of flow sizes. For all considered configurations, NC performs poorly as compared to CL, and is even significantly worse than RR. The ratio C^{NC}/C^{LP} is also much more variable than for synthetic distributions. For small values of L and n , both CL and RR provide average CCT which are lower than C^{LP} . We remark that RR can perform better than the lower bound of LP since it falls outside the class of order-based policies. The LP bound is valid only for policies that are

L	n	CL				NC				RR			
		mean	std	Q1	Q3	mean	std	Q1	Q3	mean	std	Q1	Q3
2	2	0.71	0.02	0.70	0.73	1.00	0.02	0.99	1.02	0.76	0.02	0.75	0.77
4	4	0.72	0.10	0.65	0.79	1.26	0.17	1.14	1.36	0.83	0.12	0.73	0.90
8	4	1.02	0.19	0.90	1.13	1.55	0.30	1.36	1.74	1.16	0.23	1.03	1.31
8	8	0.79	0.11	0.72	0.85	1.58	0.22	1.43	1.74	1.01	0.17	0.91	1.11
16	8	1.23	0.17	1.14	1.34	2.11	0.34	1.87	2.36	1.56	0.25	1.38	1.72
16	16	0.99	0.10	0.93	1.06	2.09	0.19	1.95	2.23	1.47	0.17	1.36	1.61
32	16	1.48	0.14	1.39	1.58	2.70	0.27	2.51	2.90	2.16	0.24	1.95	2.35

Table 10: C^{ALG}/C^{LP} for flow sizes distributed as in the Facebook trace.




Figure 6: Average values of the ratios C^{ALG}/C^{LP} for flow sizes following a Pareto distribution with $CV = 0.5$ (left), $CV = 1$ (middle) or $CV = 2$ (right). The number of ports L and the number n of coflows are indicated on the x -axis as $L.n$.




Figure 7: Histograms of the ratios C^{ALG}/C^{LP} for CL, NC, RR and RO for 100 random problem instances with $L = 32$ ports and $n = 16$ coflows. Flow sizes are distributed as in the Facebook trace.

within this class.

In all cases, RR performs much better than NC, which is not much more better than a random order scheduling. As is apparent from Figure 7, the performance of NC is similar to that of a random-order policy. This can be explained by the extreme variability of the flow size distribution in the Facebook trace which has a coefficient of variation of 3.27 for the aggregated flow-sizes of a coflow at a port. While this quantity is not exactly the coefficient of variation of the flow-sizes, we expect these two to be related. If the aggregate coefficient of variation is high, we expect that of the flow-sizes to be high as well.

In conclusion, for small instances and low variability flow-sizes, the proposed NC policy is close to optimal and is better than RR. However, for high variability flow-sizes, RR outperforms NC and can even be better than the lower bound obtained from LP.

5.2 Large instances

We now consider larger problem instances for which the lower bound C^{LP} cannot be computed. As the clairvoyant offline optimum is no more available, we choose as reference the clairvoyant Sincronia (CL), and present the ratios C^{ALG}/C^{CL} obtained with two different algorithms: the non-clairvoyant Sincronia-based priority policy (NC) and the round-robin rate allocation (RR).




Figure 8: Average values of C^{ALG}/C^{LP} for flow sizes distributed as in the Facebook trace. The number of ports L and the number n of coflows are indicated on the x -axis as $L.n$.

L	n	NC				RR			
		mean	std	Q1	Q3	mean	std	Q1	Q3
20	10	1.00	0.02	1.00	1.01	2.29	0.32	2.13	2.49
20	20	1.01	0.02	1.00	1.01	2.82	0.26	2.64	2.98
20	30	1.01	0.02	1.00	1.02	3.13	0.27	2.95	3.30
40	20	1.00	0.02	0.99	1.01	2.84	0.41	2.55	3.15
40	40	1.00	0.02	0.99	1.01	3.52	0.30	3.34	3.72
40	60	1.01	0.01	1.00	1.02	3.97	0.27	3.81	4.15
80	40	1.01	0.02	0.99	1.02	3.78	0.33	3.58	4.04
80	80	1.01	0.02	1.00	1.02	4.62	0.32	4.41	4.87
80	120	1.00	0.01	1.00	1.01	5.11	0.33	4.91	5.33

Table 11: C^{ALG}/C^{CL} for flow sizes following a Gamma distribution with $CV = 0.2$.

As before, we average over 1000 random traffic realizations for each problem instance. We present below the results obtained for a gamma distribution of flow sizes, as well as for flow sizes distributed as in the Facebook trace. The theoretical upper bound (UB) is not shown here since it was already quite pessimistic for small instances.

Gamma distribution

We consider three different coefficients of variations: $CV = 0.2$, $CV = 1$ and $CV = 5$.

The mean, standard deviation as well as the first and third quartiles of the ratios C^{ALG}/C^{CL} are reported in Tables 11, 12 and 13 for flow sizes following a Gamma distribution with $CV = 0.2$, $CV = 1$ and $CV = 5$, respectively. For $CV = 0.2$ and $CV = 1$, we observe that for all considered configurations of the numbers of ports and coflows, NC performs almost as well as CL on average, and that the variability of the ratio C^{NC}/C^{CL} is quite low. We also observe that the ratio C^{RR}/C^{CL} has a much higher mean and standard deviation, which shows that the round-robin allocation performs poorly as compared to our NC policy when the variability of flow sizes is low.

On the other hand, for $CV = 5$ (that is, flow sizes with high variability) and a switch with fewer and equal to 40 ports, the average performance of RR is better than but close to that of NC. However, for larger switches with 80 ports, RR becomes worse than NC. This observation is similar to that for small instances and $CV = 2$ (see Table 3) for which RR was better for smaller instances but worse for larger ones.

Figure 9 shows the histograms of the ratios C^{ALG}/C^{CL} obtained for the 100 random problem instances for a configuration with 80 ports and 120 coflows when $CV = 0.2$ (left) and when

L	n	NC				RR			
		mean	std	Q1	Q3	mean	std	Q1	Q3
20	10	1.06	0.04	1.03	1.08	1.94	0.20	1.82	2.10
20	20	1.08	0.03	1.06	1.10	2.37	0.18	2.29	2.51
20	30	1.09	0.03	1.07	1.11	2.64	0.16	2.54	2.74
40	20	1.05	0.02	1.04	1.06	2.36	0.30	2.16	2.58
40	40	1.06	0.02	1.04	1.07	2.94	0.22	2.82	3.08
40	60	1.07	0.01	1.06	1.08	3.30	0.21	3.17	3.44
80	40	1.04	0.02	1.03	1.05	3.13	0.21	3.02	3.27
80	80	1.05	0.02	1.04	1.06	3.83	0.24	3.67	4.02
80	120	1.05	0.01	1.04	1.06	4.24	0.25	4.08	4.40

Table 12: C^{ALG}/C^{CL} for flow sizes following a Gamma distribution with $CV = 1$.

L	n	NC				RR			
		mean	std	Q1	Q3	mean	std	Q1	Q3
20	10	1.53	0.13	1.44	1.62	1.24	0.06	1.19	1.28
20	20	1.80	0.13	1.71	1.86	1.40	0.06	1.37	1.45
20	30	1.94	0.13	1.85	2.01	1.53	0.06	1.50	1.57
40	20	1.49	0.10	1.43	1.57	1.33	0.07	1.28	1.38
40	40	1.66	0.09	1.60	1.72	1.56	0.06	1.53	1.61
40	60	1.75	0.08	1.70	1.80	1.70	0.06	1.67	1.74
80	40	1.37	0.06	1.32	1.40	1.52	0.07	1.47	1.57
80	80	1.52	0.07	1.47	1.56	1.80	0.08	1.75	1.86
80	120	1.60	0.07	1.54	1.64	1.97	0.08	1.93	2.02

Table 13: C^{ALG}/C^{CL} for flow sizes following a Gamma distribution with $CV = 5$.

$CV = 5$ (right). It is clear that NC performs quite well in both cases, whereas RR performs poorly for low variability flow sizes.

Figure 10 provides a synthetic view of the above results. It shows the average values (over 100 random instances) of C^{ALG}/C^{CL} obtained for NC and RR.

Facebook trace

The statistical results for flow size distribution derived from the empirical distribution of flow sizes in the Facebook trace is reported in Table 14, and we show the specific distributions of the ratios C^{ALG}/C^{CL} obtained for $L = 80$ and $n = 120$ in Figure 11.

For most configurations, NC performs poorly as compared to the CL, and is even significantly worse than RR. However, for instances with 80 ports, NC outperforms (and only slightly) RR. This observation is again similar to the one for small instances (see Table 10) for NC gets closer RR in performance as the instances become larger. Again, RR is better for smaller instances because of high variability of flow sizes. An intuitive explanation for this phenomenon is as follows: as the number of flows increases, the aggregate coefficient of variation at a port will decrease thus favoring NC over RR.

Figure 12 provides a synthetic view of the above results. It shows the average values (over 100 random instances) of C^{ALG}/C^{CL} obtained for the non-clairvoyant Sincronia (NC) and the round-robin allocation (RR).

6 Concluding remarks

We addressed the coflow scheduling problem in the stochastic setting, assuming that the objective is to minimize the weighted expected completion time of coflows. We proposed a simple priority policy in which the priority order is computed with **Sincronia** using the expected processing times instead of the unknown actual processing times. Performance guarantees are established for this policy using a polyhedral relaxation of the performance space of stochastic coflow scheduling policies, first for arbitrary probability distributions of processing times, and

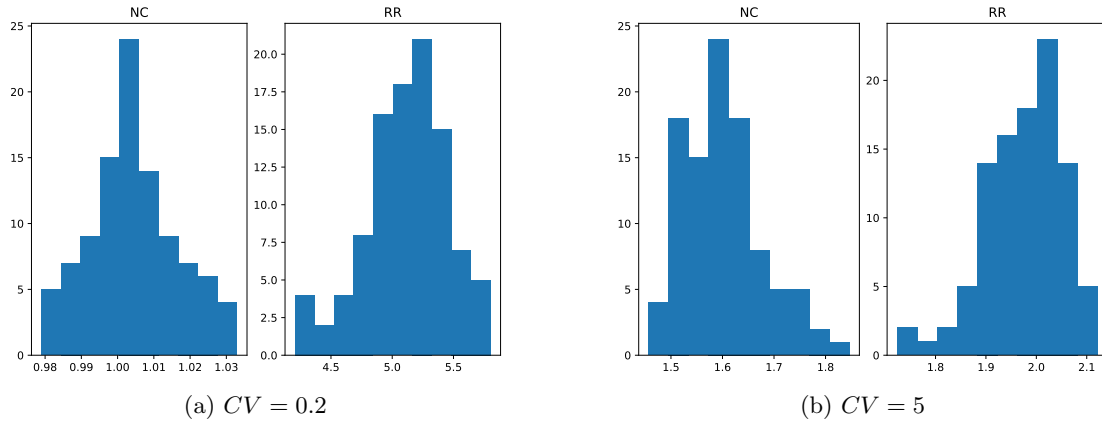


Figure 9: Histograms of the ratios C^{ALG}/C^{CL} for NC and RR for 100 random problem instances with $L = 80$ ports and $n = 120$ coflows. Flow sizes follow a Gamma distribution with either (a) $CV = 0.2$ or (b) $CV = 5$.




Figure 10: Average values of C^{ALG}/C^{CL} for flow sizes following a Gamma distribution with $CV = 0.2$ (left), $CV = 1$ (middle) or $CV = 5$ (right). The numbers of ports L and of coflows n are indicated on the x -axis as $L.n$.

L	n	NC				RR			
		mean	std	Q1	Q3	mean	std	Q1	Q3
20	10	1.70	0.12	1.62	1.76	1.25	0.06	1.21	1.30
20	20	2.15	0.14	2.05	2.24	1.45	0.06	1.42	1.50
20	30	2.43	0.16	2.29	2.52	1.61	0.06	1.57	1.65
40	20	1.65	0.10	1.57	1.71	1.36	0.08	1.31	1.41
40	40	1.98	0.11	1.92	2.05	1.62	0.07	1.58	1.66
40	60	2.17	0.11	2.09	2.24	1.80	0.07	1.75	1.84
80	40	1.57	0.07	1.52	1.62	1.56	0.09	1.49	1.61
80	80	1.81	0.08	1.75	1.86	1.86	0.09	1.80	1.92
80	120n	1.93	0.08	1.88	1.98	2.06	0.09	2.00	2.12

Table 14: Main statistical parameters of ratios C^{ALG}/C^{CL} for flow sizes distributed as in the Facebook trace.

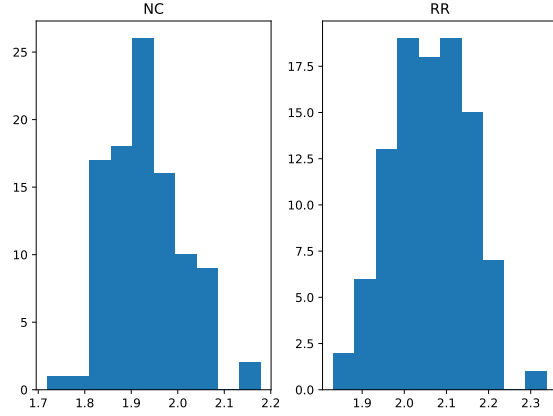


Figure 11: Histograms of the ratios C^{ALG}/C^{CL} for NC and RR for 100 random problem instances with $L = 80$ ports and $n = 120$ coflows. Flow sizes are distributed as in the Facebook trace.




Figure 12: Average values of C^{ALG}/C^{CL} for NC and RR for flow sizes distributed as in the Facebook trace. The numbers of ports L and of coflows n are indicated on the x -axis as $L.n$.

then for some specific distributions. Empirical results show that this priority policy performs very well in practice, and much better than the theoretical upper bounds suggest.

We see the following possible future directions emanating from this work. First, we plan to design preemptive scheduling policies leveraging the information on the attained service of coflows. Another interesting direction will be to propose policies that choose between RR and NC based on the variability of the flow-sizes. Further, it will also be interesting to compute the approximation ratio for a richer class of policies than the order-based ones in this paper. As we saw, certain policies like RR are not included in this class. Finally, we intend to establish tighter bounds for some other specific distributions of processing times, and hope to propose methods for computing the mean CCT in some asymptotic regimes such as large number of coflows and ports.

References

- [1] Radosław Adamczak, Joscha Prochno, Marta Strzelecka, and Michał Strzelecki. Norms of structured random matrices. *Mathematische Annalen*, 388(4):3463–3527, April 2023. URL: <http://dx.doi.org/10.1007/s00208-023-02599-6>, doi:10.1007/s00208-023-02599-6.
- [2] Saksham Agarwal and Rachit Agarwal. <https://github.com/sincronia-coflow/workload-generator>.
- [3] Saksham Agarwal, Rachit Agarwal, Shijin Rajakrishnan, David Shmoys, Akshay Narayan, and Amin Vahdat. Sincronia: Near-Optimal Network Design for Coflows. In *Proc. ACM SIGCOMM*, pages 16–29, 2018. doi:10.1145/3230543.3230569.
- [4] Saba Ahmadi, Samir Khuller, Manish Purohit, and Sheng Yang. On scheduling coflows. *Algorithmica*, 82:3604 – 3629, 2020.
- [5] Nikhil Bansal and Subhash Khot. Inapproximability of hypergraph vertex cover and applications to scheduling problems. In Samson Abramsky, Cyril Gavoille, Claude Kirchner, Friedhelm Meyer auf der Heide, and Paul G. Spirakis, editors, *Automata, Languages and Programming, 37th International Colloquium, ICALP 2010, Bordeaux, France, July 6-10, 2010, Proceedings, Part I*, volume 6198 of *Lecture Notes in Computer Science*, pages 250–261. Springer, 2010. doi:10.1007/978-3-642-14165-2_22.
- [6] Akhil Bhimaraju, Debanuj Nayak, and Rahul Vaze. Non-clairvoyant scheduling of coflows. In *18th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks, WiOPT 2020, Volos, Greece, June 15-19, 2020*, pages 81–88. IEEE, 2020. URL: <https://dl.ifip.org/db/conf/wiopt/wiopt2020/32.pdf>.
- [7] Olivier Brun, Balakrishna J. Prabhu, and Oumayma Haddaji. Prediction-based coflow scheduling. In *2024 22nd International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, pages 185–192, 2024. URL: <https://laas.hal.science/hal-04231370>.
- [8] Li Chen, Wei Cui, Baochun Li, and Bo Li. Optimizing coflow completion times with utility max-min fairness. In *Proc. IEEE INFOCOM*, pages 1–9, 2016.
- [9] Mosharaf Chowdhury, Samir Khuller, Manish Purohit, Sheng Yang, and Jie You. Near optimal coflow scheduling in networks. In *Proc. ACM SPAA*, pages 123–134, Phoenix, AZ, USA, June 22-24 2019.
- [10] Mosharaf Chowdhury and Ion Stoica. Coflow: A networking abstraction for cluster applications. In *Proc. ACM HotNets*, pages 31–36, Redmond, Washington, 2012.
- [11] Mosharaf Chowdhury and Ion Stoica. Efficient coflow scheduling without prior knowledge. *SIGCOMM Comput. Commun. Rev.*, 45(4):393–406, August 2015.

- [12] Mosharaf Chowdhury, Yuan Zhong, and Ion Stoica. Efficient Coflow Scheduling with Varys. In *Proc. ACM SIGCOMM*, pages 443–454, 2014. doi:10.1145/2619239.2626315.
- [13] Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, 2008.
- [14] Fahad R. Dogar, Thomas Karagiannis, Hitesh Ballani, and Antony Rowstron. Decentralized task-aware scheduling for data center networks. In *Proc. ACM SIGCOMM*, pages 431–442, New York, NY, USA, 2014. Association for Computing Machinery. doi:10.1145/2619239.2626322.
- [15] Yuanxiang Gao, Hongfang Yu, Shouxi Luo, and Shui Yu. Information-agnostic coflow scheduling with optimal demotion thresholds. In *2016 IEEE International Conference on Communications (ICC)*, pages 1–6, 2016. doi:10.1109/ICC.2016.7511241.
- [16] Sungjin Im, Ravi Kumar, Mahshid Montazer Qaem, and Manish Purohit. Non-clairvoyant scheduling with predictions. In *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '21*, pages 285–294, New York, NY, USA, 2021. Association for Computing Machinery. doi:10.1145/3409964.3461790.
- [17] Samir Khuller, Jingling Li, Pascal Sturmfels, Kevin Sun, and Prayaag Venkat. Select and permute: An improved online framework for scheduling to minimize weighted completion time. In Michael A. Bender, Martin Farach-Colton, and Miguel A. Mosteiro, editors, *LATIN*, volume 10807 of *Lecture Notes in Computer Science*, pages 669–682. Springer, 2018.
- [18] Libin Liu, Hong Xu, Chengxi Gao, and Peng Wang. Bottleneck-aware coflow scheduling without prior knowledge. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 50–55, 2020. doi:10.1109/INFOCOMWKSHPS50562.2020.9163039.
- [19] Ruijiu Mao, Vaneet Aggarwal, and Mung Chiang. Stochastic non-preemptive co-flow scheduling with time-indexed relaxation. In *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, pages 385–390, 2018. doi:10.1109/INFOCOMW.2018.8406913.
- [20] T. Mikosch. *Regular variation, subexponentiality and their applications in probability theory*. Report Eurandom. Eurandom, 1999.
- [21] Rolf H. Möhring, Andreas S. Schulz, and Marc Uetz. Approximation in stochastic scheduling: The power of LP-based priority policies. *J. ACM*, 46(6):924–942, nov 1999. doi:10.1145/331524.331530.
- [22] C. Klüppelberg P. Embrechts and T. Mikosch. *Modelling Extremal Events for Insurance and Finance*. Springer, 2011.
- [23] Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving online algorithms via ml predictions. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. URL: <https://proceedings.neurips.cc/paper/2018/file/73a427badebe0e32caa2e1fc7530b7f3-Paper.pdf>.
- [24] M. Queyranne. Structure of a simple scheduling polyhedron. *Mathematical Programming*, 58:263–285, 1993.
- [25] Sushant Sachdeva and Rishi Saket. Optimal inapproximability for scheduling problems via structural hardness for hypergraph vertex cover. In *2013 IEEE Conference on Computational Complexity*, pages 219–229, 2013. doi:10.1109/CCC.2013.30.

- [26] Andreas S. Schulz. Scheduling to minimize total weighted completion time: Performance guarantees of lp-based heuristics and lower bounds. In *Conference on Integer Programming and Combinatorial Optimization*, 1996.
- [27] Mehrnoosh Shafiee and Javad Ghaderi. An improved bound for minimizing the total weighted completion time of coflows in datacenters. *IEEE/ACM Trans. Netw.*, 26(4):1674–1687, 2018. doi:10.1109/TNET.2018.2845852.
- [28] Li Shi, Yang Liu, Junwei Zhang, and Thomas Robertazzi. Coflow scheduling in data centers: routing and bandwidth allocation. *IEEE Trans. Parallel Distrib. Syst.*, 32(11):2661–2675, 2021. doi:10.1109/TPDS.2021.3068424.
- [29] Martin Skutella, Maxim Sviridenko, and Marc Uetz. Unrelated machine scheduling with stochastic processing times. *Mathematics of Operations Research*, 41(3):851–864, 2016. arXiv:https://doi.org/10.1287/moor.2015.0757, doi:10.1287/moor.2015.0757.
- [30] Flemming Topsoe. Some bounds for the logarithmic function. *RGMA research report collection*, 7(2), 2004. URL: <https://rgmia.org/papers/v7n2/pade.pdf>.
- [31] Junchao Wang, Huan Zhou, Yang Hu, Cees De Laat, and Zhiming Zhao. Deadline-aware coflow scheduling in a dag. In *Proc. IEEE CloudCom*, pages 341–346, 2017.
- [32] Zhiliang Wang, Han Zhang, Xingang Shi, Xia Yin, Yahui Li, Haijun Geng, Qianhong Wu, and Jianwei Liu. Efficient scheduling of weighted coflows in data centers. *IEEE Transactions on Parallel and Distributed Systems*, 30(9):2003–2017, 2019. doi:10.1109/TPDS.2019.2905560.
- [33] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, Ion Stoica, et al. Spark: Cluster computing with working sets. *HotCloud*, 10(10-10):95, 2010.
- [34] 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. Association for Computing Machinery. doi:10.1145/2934872.2934880.
- [35] Tong Zhang, Fengyuan Ren, Jiakun Bao, Ran Shu, and Wenxue Cheng. Minimizing coflow completion time in optical circuit switched networks. *IEEE Trans. Parallel Distrib. Syst.*, 32(2):457–469, 2021.
- [36] Tong Zhang, Fengyuan Ren, Ran Shu, and Bo Wang. Scheduling coflows with incomplete information. In *2018 IEEE/ACM 26th International Symposium on Quality of Service (IWQoS)*, pages 1–10, 2018. doi:10.1109/IWQoS.2018.8624126.

A Proof of Theorem 1

Proof. Let r be a work-conserving and nonanticipative rate-allocation policy and define the random variable $Z_{\ell,k} = C_k(r) - P_{\ell,k}$. Consider any fixed realization of flow volumes. The latter yields a specific realization of the processing times $p_{\ell,k}$ and of the completion times $c_k(r)$ (we use lowercase letters to denote the particular values that our random variables take for this specific realization of flow volumes). Assume that, for this realization, coflows complete in the order π , that is, $c_{\pi(1)}(r) \leq c_{\pi(2)}(r) \leq \dots \leq c_{\pi(n)}(r)$. As coflow $\pi(k)$ completes after all coflows $\pi(j)$, $j = 1, \dots, k - 1$, we clearly have $c_{\pi(k)}(r) \geq \sum_{j \leq k} p_{\ell, \pi(j)}$ for all ports ℓ . It yields $z_{\ell, \pi(k)} \geq \sum_{j < k} p_{\ell, \pi(j)} \delta_{\pi(j)}^A$ for any subset $A \subseteq C$, where $\delta_{\pi(j)}^A$ is 1 if $\pi(j) \in A$ and 0 otherwise. We thus obtain

$$\begin{aligned}
\sum_{j \in A} p_{\ell,j} z_{\ell,j} &= \sum_k p_{\ell,\pi(k)} z_{\ell,\pi(k)} \delta_{\pi(k)}^A, \\
&\geq \sum_k \sum_{j < k} p_{\ell,\pi(k)} \delta_{\pi(k)}^A p_{\ell,\pi(j)} \delta_{\pi(j)}^A, \\
&= \frac{1}{2} \sum_{j,k:j \neq k} p_{\ell,\pi(k)} \delta_{\pi(k)}^A p_{\ell,\pi(j)} \delta_{\pi(j)}^A, \\
&= \frac{1}{2} \sum_{j,k \in A, j \neq k} p_{\ell,k} p_{\ell,j}, \tag{24}
\end{aligned}$$

and the lower bound is independent of the considered order. Conditioning on the order π and taking expectation on both sides of (24), it follows that

$$\begin{aligned}
\mathbb{E} \left[\sum_{j \in A} P_{\ell,j} Z_{\ell,j} \mid \text{order } \pi \right] &\geq \frac{1}{2} \sum_{j,k \in A, j \neq k} \mathbb{E} [P_{\ell,k} P_{\ell,j}], \\
&= \frac{1}{2} \sum_{j,k \in A, j \neq k} \mu_{\ell,k} \mu_{\ell,j}, \\
&= \frac{1}{2} \left(\sum_{k \in A} \mu_{\ell,k} \right)^2 - \frac{1}{2} \sum_{k \in A} \mu_{\ell,k}^2,
\end{aligned}$$

where the first equality follows from the independence of processing times. In general, for an arbitrary rate-allocation policy r , several completion orders are possible. However, if we assume that r is an order-based policy, the order π in which coflows complete is unique and depends only on the partial order implemented by this policy, and not on the actual realization of the processing times. Moreover, as processing times are independent, the random variables $P_{\ell,k}$ and $Z_{\ell,k}$ are statistically independent in this case, so that $\mathbb{E} [P_{\ell,k} Z_{\ell,k}] = \mu_{\ell,k} \mathbb{E} [Z_{\ell,k}]$. Hence

$$\mathbb{E} \left[\sum_{j \in A} P_{\ell,j} Z_{\ell,j} \mid \text{order } \pi \right] = \mathbb{E} \left[\sum_{j \in A} P_{\ell,j} Z_{\ell,j} \right] = \sum_{j \in A} \mu_{\ell,j} \mathbb{E} [Z_{\ell,j}],$$

so that

$$\sum_{j \in A} \mu_{\ell,j} \mathbb{E} [Z_{\ell,j}] \geq \frac{1}{2} \left(\sum_{k \in A} \mu_{\ell,k} \right)^2 - \frac{1}{2} \sum_{k \in A} \mu_{\ell,k}^2,$$

and we can conclude the proof with $\mathbb{E} [Z_{\ell,j}] = \mathbb{E} [C_j(r)] - \mu_{\ell,j}$. \square

B Proof of Lemma 2

The proof follows the standard approach of using the moment generating function. We give the details here since a similar proof is also used for Gaussian processing times.

Proof. First, observe that

$$\max_{\ell,j} \beta_{\ell,j} = \max_{\ell,j} \frac{\sigma_{\ell,j}^2}{\mu_{\ell,j}} \leq \gamma \mu_{\min}. \tag{25}$$

Let $X_{\ell,k} = \sum_{j=1}^k P_{\ell,j}$ and $Y_k = \max_{\ell} \frac{X_{\ell,k}}{\sum_i \mu_{\ell,i}}$. Using Jensen's inequality we obtain

$$\begin{aligned} e^{t\mathbb{E}[Y_k]} &\leq \mathbb{E} [e^{tY_k}] = \mathbb{E} \left[\max_{\ell} e^{t \frac{X_{\ell,k}}{\sum_i \mu_{\ell,i}}} \right] \leq \sum_{\ell} \mathbb{E} \left[e^{t \frac{X_{\ell,k}}{\sum_i \mu_{\ell,i}}} \right] \\ &= \sum_{\ell} \prod_{j=1}^k \mathbb{E} \left[e^{\frac{t}{\sum_i \mu_{\ell,i}} P_{\ell,j}} \right] \end{aligned} \quad (26)$$

with the last inequality following from the independence of $P_{\ell,j}$. Since $P_{\ell,j} \sim \Gamma(s_{\ell,j}, \beta_{\ell,j})$,

$$\begin{aligned} \mathbb{E}[e^{tP_{\ell,j}}] &= (1 - \beta_{\ell,j}t)^{-s_{\ell,j}} = (1 - \beta_{\ell,j}t)^{-\frac{\mu_{\ell,j}}{\beta_{\ell,j}}} \\ &\leq (1 - \gamma\mu_{\min}t)^{-\frac{\mu_{\ell,j}}{\gamma\mu_{\min}}}, \quad t < \gamma\mu_{\min}^{-1}, \end{aligned} \quad (27)$$

where the second equality follows from the fact that $s_{\ell,j}\beta_{\ell,j} = \mu_{\ell,j}$ and the last inequality is a consequence of the fact that the RHS is increasing in $\beta_{\ell,j}$ and $\beta_{\ell,j} < \gamma\mu_{\min}$ (see (25)). Substituting (27) in (26), for $t < \beta_{\max}^{-1}$,

$$\begin{aligned} \sum_{\ell} \prod_{j=1}^k \mathbb{E} \left[e^{\frac{t}{\sum_i \mu_{\ell,i}} P_{\ell,j}} \right] &\leq \sum_{\ell} \left(\frac{1}{1 - \frac{\gamma\mu_{\min}t}{\sum_j \mu_{\ell,j}}} \right)^{\frac{\sum_j \mu_{\ell,j}}{\gamma\mu_{\min}}} \\ &\leq L \left(\frac{1}{1 - \gamma t} \right)^{\frac{1}{\gamma}}, \end{aligned}$$

where the last inequality uses the fact that the RHS is decreasing in $\sum_j \mu_{\ell,j}$ so that we can substitute $\sum_j \mu_{\ell,j}$ by μ_{\min} . Thus,

$$\mathbb{E}[Y_k] \leq \frac{\gamma \log(L) - \log(1 - \gamma t)}{\gamma t}, \quad \text{for } t < \gamma^{-1}.$$

The minimum of the RHS is obtained for

$$\gamma \hat{t} : \frac{1}{1 - \gamma \hat{t}} = \frac{\gamma \log(L) - \log(1 - \gamma \hat{t})}{\gamma \hat{t}} \quad (28)$$

and

$$\mathbb{E}[Y_k] \leq \frac{1}{1 - \gamma \hat{t}}.$$

We now show that the RHS in (23) is an upper bound for $\frac{1}{1 - \gamma \hat{t}} =: z$. Rearranging (28), z is the solution of

$$z - \log(z) = 1 + \gamma \log L. \quad (29)$$

Note that the LHS is an increasing function of z and any z satisfying the equality is larger than 1. Applying the bound [30]

$$\log(z) \leq \frac{(z-1)(z+1)}{2z}, \quad \text{for } z \geq 1,$$

in the LHS of (29) we get

$$z - \frac{(z-1)(z+1)}{2z} \leq z - \log(z) = c,$$

which leads to

$$z \leq c + \sqrt{c^2 - 1} \Leftrightarrow \log(z) \leq \log(c + \sqrt{c^2 - 1})$$

Substituting this bound for $\log(z)$ in (29), we obtain the stated result. \square

C Proof of Lemma 3

Proof. To simplify notations, let $X_\ell = \sum_{j=1}^k P_{\ell,j}$ and $Y = \max_\ell X_\ell$. We follow a standard approach by first deriving a bound on the moment-generating function of Y :

$$\exp(t\mathbb{E}[Y]) \leq \mathbb{E}[\exp(tY)], \quad (30)$$

$$= \mathbb{E}\left[\max_\ell \exp(tX_\ell)\right], \quad (31)$$

$$\leq \sum_\ell \mathbb{E}[\exp(tX_\ell)], \quad (32)$$

$$= \sum_\ell \mathbb{E}\left[\prod_{j=1}^k \exp(tP_{\ell,j})\right], \quad (33)$$

$$= \sum_\ell \prod_{j=1}^k \mathbb{E}[\exp(tP_{\ell,j})], \quad (34)$$

where inequality (30) follows from Jensen's inequality and equality (34) follows from the independence assumption. As we assume a normal distribution for the $P_{\ell,j}$, we have $\mathbb{E}[\exp(tP_{\ell,j})] = \exp\left(t\mu_{\ell,j} + \frac{\sigma_{\ell,j}^2}{2}t^2\right)$, where $\sigma_{\ell,j}^2 = CV_{\ell,j}^2\mu_{\ell,j}^2$ is the variance. It yields

$$\begin{aligned} \exp(t\mathbb{E}[Y]) &\leq \sum_\ell \exp\left(t\sum_{j=1}^k \mu_{\ell,j} + \frac{1}{2}\left(\sum_{j=1}^k \sigma_{\ell,j}^2\right)t^2\right), \\ &\leq L \exp\left(t\mu + \frac{\sigma^2}{2}t^2\right), \end{aligned}$$

where $\mu = \max_\ell \sum_{j=1}^k \mu_{\ell,j}$ and $\sigma^2 = \max_\ell \sum_{j=1}^k \sigma_{\ell,j}^2$. By taking the logarithm on both sides, we thus have $\mathbb{E}[Y] \leq \frac{1}{t} \log(L) + \mu + \frac{\sigma^2}{2}t$ for all $t > 0$. As the minimum of the right-hand side is obtained for $t = \sqrt{2 \log(L)}/\sigma$, we obtain $\mathbb{E}[Y] \leq \mu + \sqrt{2 \log(L)}\sigma$, that is,

$$\mathbb{E}\left[\max_\ell \sum_{j=1}^k P_{\ell,j}\right] \leq \max_\ell \sum_{j=1}^k \mu_{\ell,j} + \sqrt{2 \log(L)} \max_\ell \sqrt{\sum_{j=1}^k CV_{\ell,j}^2 \mu_{\ell,j}^2}. \quad (35)$$

It yields

$$\mathbb{E}\left[\max_\ell \sum_{j=1}^k P_{\ell,j}\right] \leq \max_\ell \sum_{j=1}^k \mu_{\ell,j} + \sqrt{2 \log(L)} CV_{max} \max_\ell \sqrt{\sum_{j=1}^k \mu_{\ell,j}^2}, \quad (36)$$

$$\leq \max_\ell \sum_{j=1}^k \mu_{\ell,j} + \sqrt{2 \log(L)} CV_{max} \max_\ell \sum_{j=1}^k \mu_{\ell,j}, \quad (37)$$

so that $\alpha \leq 1 + \sqrt{2 \log(L)} CV_{max}$, as claimed. \square