

Asymptotically Optimal Approximation Algorithms for Coflow Scheduling

Hamidreza Jahanjou
Northeastern University
Boston, MA 02115
hamid@ccs.neu.edu

Erez Kantor
Northeastern University
Boston, MA 02115
erezk@ccs.neu.edu

Rajmohan Rajaraman
Northeastern University
Boston, MA 02115
rraj@ccs.neu.edu

ABSTRACT

Many modern datacenter applications involve large-scale computations composed of multiple data flows that need to be completed over a shared set of distributed resources. Such a computation completes when all of its flows complete. A useful abstraction for modeling such scenarios is a *coflow*, which is a collection of flows (e.g., tasks, packets, data transmissions) that all share the same performance goal.

In this paper, we present the first approximation algorithms for scheduling coflows over general network topologies with the objective of minimizing total weighted completion time. We consider two different models for coflows based on the nature of individual flows: circuits, and packets. We design constant-factor polynomial-time approximation algorithms for scheduling packet-based coflows with or without given flow paths, and circuit-based coflows with given flow paths. Furthermore, we give an $O(\log n / \log \log n)$ -approximation polynomial time algorithm for scheduling circuit-based coflows without given flow paths (here n is the number of network edges).

We obtain our results by developing a general framework for coflow schedules, based on interval-indexed linear programs, which may extend to other coflow models and objective functions and may also yield improved approximation bounds for specific network scenarios. We also present an experimental evaluation of our approach for circuit-based coflows that show a performance improvement of at least %22 on average over competing heuristics.

1 INTRODUCTION

In modern datacenter network applications, a large-scale computation over a big data set is often composed of multiple tasks or multiple data flows that need to be completed over a shared set of distributed resources. Such a computation completes when all the data flows in the computation complete. A useful abstraction for modeling such scenarios is that of a *coflow* [5], which is a collection of flows (e.g., tasks, packets, data transmissions) that all share the

same performance goal. Coflows arise frequently in distributed computing and datacenters [2, 5, 6, 13, 23, 28, 29]. A prominent example is parallel data processing; in certain MapReduce computations, the reduce phase at a particular reducer can begin only after all the relevant data from the map phase has arrived at the reducer [10]. Similar examples occur in Dryad [14], and Spark [27].

In this paper, we study the problem of designing coflow schedules that minimize total weighted completion time. We consider two different models for coflows based on the nature of individual flows: circuits and packets.

- **Circuit-based coflows** a flow is a connection request (data transmission) from source to sink on a network.
- **Packet-based coflows** a flow is a packet that needs to be routed from source to sink on a network.

Before going into formal definitions, Figure 1 illustrates a circuit-based coflow scheduling scenario. Three coflows A , B , and C require bandwidth assignment and scheduling on a triangle network with unit edge capacities. Coflow A consists of two flows A_1 of size 2 and A_2 of size 1 whereas the other two have one flow each of size 1. Three solutions are proposed. In (s1), each flow is given a bandwidth of $1/2$; all flows are scheduled to run in parallel. The total completion time is $4 + 2 + 4 = 10$. In (s2), coflow A has the highest priority followed by coflow B and C . The total completion time is $2 + 2 + 4 = 8$. But, it is possible to do even better by observing that flow C can run at the same time as either flow A_2 or flow B ; this gives rise to an optimal solution, shown in (s3), with a total completion time of $4 + 2 + 1 = 7$.

We now present the models and problem formulations, our main results and techniques, and a review of related work.

1.1 Models and Notation

Regardless of the model, the following notation is used. The set of all flows is denoted by $F = \{f_1, \dots, f_n\}$. We define a *coflow* F_i to be a subset $\{f_1^i, \dots, f_{n_i}^i\}$ of flows such that the set of all coflows, denoted by \mathcal{F} , provides a partition of F . Each flow f_i has a source s_i , a destination d_i , and, depending on the model, a size σ_i . Additionally, each flow has a release time r_i at which point it becomes available. Notice that in our formulation each flow has a release time whereas in previous works, release times are at the level of coflows.

We also have a network modeled as a directed graph $G = (V, E)$ which, in the case of circuit-based coflows, also features edge capacities $\{c(e)\}_{e \in E}$.

The completion time of a flow f_i is denoted by c_i . The completion time C_k of a coflow F_k is defined to be the completion time of its

Supported in part by NSF grants Grant CSR-1217981 and CCF-1422715, and a Google Research Award.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SPAA'17, July 24–26, 2017, Washington, DC, USA.

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-4593-4/17/07...\$15.00

<https://doi.org/10.1145/3087556.3087567>

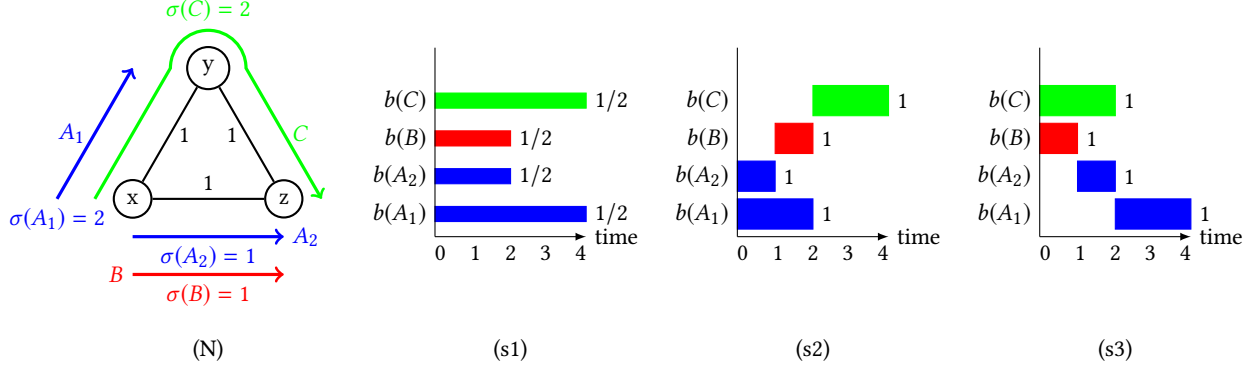


Figure 1: Three possible bandwidth assignments for the set of coflows $\{A, B, C\}$. Coflow A has two flows (A_1 of size 2 and A_2 of size 1); the other two have each one flow of size 1. The objective is to minimize sum of completion times. (N) The network is a triangle where each edge has unit capacity. Flows are drawn around the edges along with their size. (s1) A bandwidth of 1/2 is assigned to each flow; total completion time is 10. (s2) Priority is given to coflow A then to coflow B and then to coflow C; total completion time is 8. (s3) An optimal solution; total completion time is 7.

last flow; that is,

$$C_k = \max_{f \in F_k} c_f.$$

The goal is to minimize the weighted sum of coflow completion times

$$\bar{C} = \sum_k \omega_k C_k = \sum_k \omega_k \max_{f \in F_k} c_f. \quad (1)$$

Note that if there is only one coflow, the problem reduces to makespan minimization. Similarly, if every coflow has one flow, the problem reduces to minimizing total weighted (flow) completion time. Therefore, in a sense, coflow scheduling bridges two well-studied minimization objectives. In particular, this means that hardness results for either of these problems also hold for coflow scheduling. Below, we consider the particulars of each model.

Circuits. In this model, a flow f_i is a connection request (data transmission) from a source node s_i to a destination node d_i on a network $G = (V, E)$ with edge capacities $\{c(e)\}_{e \in E}$. A coflow is a set of connection requests. To each connection request in F , we need to assign a path p_i (if not given already) and a bandwidth b_i such that \bar{C} is minimized without violating edge capacities. In general, the bandwidth of a connection request is an integrable function of time, $b_i(t)$, and its completion time c_i is defined as the smallest value satisfying

$$\int_0^{c_i} b_i(t) dt = \sigma_i,$$

where σ_i is the size of the connection request.

As we will see, one can assume, without loss of generality, that the bandwidth functions are piece-wise constant. We can further subdivide this problem into two cases: (a) paths are given for each connection request and we only need to assign bandwidths, (b) paths are not given; hence, each connection request requires routing and bandwidth assignment.

Packets. In this case, each flow f_i is a packet residing (at its release time r_i) on its source node s_i waiting to be routed to its destination node $d_i \in V - \{s_i\}$. A coflow is a set of packets. Each link

Model	Paths	Approx.	Hardness
Packet-based	given	$O(1)$	APX-hard
	not given		NP-hard
Circuit-based	given	$O(\frac{\log E }{\log \log E })$	$\Omega(\frac{\log E }{\log \log E })$
	not given		$\Omega(\frac{\log E }{\log \log E })$

Table 1: Approximation ratio of our algorithms and corresponding lower bounds for different models.

$e \in E$ can serve at most one packet at a time (i.e. edge capacities are set to 1). A coflow F_k completes at time C_k when all of its packets have reached their destination. The goal is to minimize \bar{C} .

This model can be seen as a discrete version of circuit-based coflows. As in the previous model, we will consider two subproblems: (a) a path is given for each packet; only scheduling is needed (b) paths are not given; both routing and scheduling are required for each packet.

1.2 Our Results and Techniques

We give asymptotically optimal approximation algorithms for minimizing total weighted completion time under the different models defined above.

- We present $O(1)$ -approximation algorithms for circuit-based coflows with given paths, packet-based coflows with given paths, and packet-based coflows where paths are not given.
- We present an $O(\log n / \log \log n)$ -approximation algorithm for circuit-based coflows where paths are not given.
- We also present an experimental evaluation of our approach for circuit-based coflows which shows %22 or more performance improvement on average over competing heuristics.

Note that, all of the above problems are at least NP-hard. Table 1 summarizes our results.

A noteworthy contribution of our work is a common framework for developing coflow schedules, which applies to all the models

we study. The framework consists of three parts. First, we reformulate the coflow scheduling problem as an instance of minimizing total weighted flow completion time problem with precedence constraints in form of depth-1 in-trees. Second, we devise an interval-index linear program for each problem instance, and then reduce the problem of finding flow schedules that minimize total weighted completion time (with aforementioned precedences) to multiple instances of the problem of finding flow schedules that minimize makespan. Third, we perform rounding. The specifics of each step varies across models. In particular, the last step is achieved by a careful rounding of the associated linear program solution. Though our approximation bounds are optimal to within constant factors, we have not attempted to optimize these constants. We believe that our framework provides a general approach for solving coflow scheduling problems under diverse models and objectives.

The algorithms given for circuit-based coflows provide the first provable approximation bounds, and apply in the general setting of arbitrary capacitated networks, arbitrary source-sink pairs, arbitrary demands, and arbitrary release times.

To the best of our knowledge, packet-based coflows have not been studied earlier, though related models have been analyzed, as we discuss in previous work.

1.3 Previous Work

Coflows. The notion of coflow was first proposed by Chowdhury and Stoica as “a networking abstraction to express the communication requirements of prevalent data parallel programming paradigms” [5, 6]. In [7], Chowdhury et al. present effective heuristics for coflow scheduling without release times on a non-blocking switch network. They also show that coflow scheduling is NP-hard for $P \times P$ switches, $P \geq 2$, via a reduction from the concurrent open shop scheduling problem. Approximation algorithms for coflow scheduling on a non-blocking switch with (coflow) release times are given in [15, 23]. Zhao et al. consider coflow scheduling over general network topologies; thereby, routing the flows becomes an additional requirement [29]. They provide heuristics based on the observation that simultaneous routing and scheduling of coflows is required to obtain good performance.

Machine scheduling and task-based coflows. Although coflow scheduling is a relatively recent development, the related problem of scheduling jobs consisting of tasks on parallel machines where a job completes when all of its tasks complete, has been studied before. This model is known in the literature as *order scheduling*; the goal is to minimize the (weighted) sum of job completion times. In their survey, Leung et al., classify order scheduling problems into three categories based on capabilities of the machines [21]. Specifically, in the fully-dedicated case, each machine can process only one type of task. This model, sometimes denoted by $G|| \sum \omega_j C_j$, is also known as the concurrent open shop scheduling since, in contrast to the latter, tasks of a job can be processed at the same time [11]. In the fully-flexible case, all machines can process all types of task. Lastly, in the arbitrary case, unrestricted subsets of machines can process arbitrary subsets of tasks.

Machine scheduling problems are described [12] by three values $\alpha|\beta|\gamma$ where α represents the machine environment, β describes job constraints and γ specifies the objective function to be minimized.

The task-based coflow scheduling model, where a flow is simply a task, corresponds to the fully-flexible case on (unrelated) parallel machines. For this model, Blocher and Chhajed observe that the problem is strongly NP-complete even when all machines are identical [4]. Correra, Skutella and Verschae give a 13.5 approximation algorithm for this problem on unrelated machines [9].

Circuit-based coflows. Circuit-based coflow scheduling with fixed paths is a generalization of $1|r_j, pmtn| \sum \omega_j C_j$ which is strongly NP-hard [18]. On the other hand, the version with variable paths is related to congestion minimization. Chuzhoy et al. show that congestion minimization on directed graphs is hard to approximate within $\Omega(\frac{\log |E|}{\log \log |E|})$ unless NP has $N^{O(\log \log N)}$ -time randomized algorithms [8].

Packet-based coflows. The problem of scheduling packets on store-and-forward networks with given paths is equivalent to the job shop scheduling problem with unit processing times. Leighton, Maggs and Rao [19] show that there exists a schedule achieving a makespan of $O(C+D)$ where C and D denote maximum edge congestion and maximum path dilation respectively. In [20], Leighton et al. give a polynomial-time algorithm for finding such a schedule. Peis et al. [22] show that for all $\epsilon > 0$ there is no $(6/5 - \epsilon)$ -approximation algorithm for the packet scheduling problem with given paths (minimizing makespan) unless $P = NP$. An $O(1)$ -approximation algorithm for $|r_{ij}, p_{ij} = 1| \sum_S C_S$ is given in [24]; the authors consider a generalized minsum objective function.

On the other hand, when paths are not given, Srinivasan and Teo give the first constant-ratio approximation algorithm for packet routing and scheduling to minimize makespan [26]. Koch et al. [16] extend this result to messages which consist of packets having the same source and sink.

The rest of this paper is organized as follows. We study circuit-based coflows in §2 and packet-based coflows in §3. We present experimental results in §4 and conclude in §5.

2 CIRCUIT-BASED COFLOWS

We now consider our main model. In this setting, we have a network $G = (V, E)$ with edge capacities $\{c(e)\}_{e \in E}$. A coflow F_i has a weight $\omega_i \in \mathbb{R}^{\geq 0}$ and consists of a number of connection requests (data transmissions) f_j^i . Each connection request has a source $s_j^i \in V$, a sink $d_j^i \in V$, a release time $r_j^i \in \mathbb{R}^{\geq 0}$ and a size $\sigma_j^i \in \mathbb{R}^{\geq 0}$. A solution must specify a path p_j^i from source to sink (if not given), and a bandwidth function b_j^i for each connection request f_j^i . In general, the bandwidth $b_j^i(t)$ is an integrable function of time such that

$$\int_{r_j^i}^{c_j^i} b_j^i(t) dt = \sigma_j^i, \quad (2)$$

where c_j^i , the completion time of f_j^i , is the smallest value for which the equality holds. In addition, we require that edge capacities are respected. Hence, any solution must satisfy

$$\forall e \in E : \sum_{f_j^i \in P(e)} b_j^i(t) \leq c(e) \quad (3)$$

for all t , where $P(e)$ is the set of paths that use edge e .

There are two natural versions of this problem. In one version, a path p_j^i is given as a part of the input for each connection request f_j^i and we need to assign a bandwidth b_j^i to each request with the goal of minimizing total weighted (coflow) completion time \bar{C} without violating edge capacities. Note that any network topology in which there is a unique path between pairs of vertices, e.g. trees or non-blocking switches, falls into this category. In another version, paths are not given; hence, assigning a (single) path to each connection request becomes an additional requirement.

The rest of this section is organized as follows. First, we present two key lemmas used in deriving our results. Next, in §2.1, we give a 17.6-approximation algorithm for the circuit-based coflow problem with given paths. Finally, in §2.2, an $O(\frac{\log |E|}{\log \log |E|})$ -approximation algorithm is presented for the case where paths are not given.

Two lemmas. The following two lemmas are simple but crucial. The first one (Lemma 2.1) states that, without loss of generality, we can assume that bandwidth functions are piece-wise constant.

LEMMA 2.1. *Suppose that there exist a capacitated network (G, c) and a set of feasible flows $\{f_1, \dots, f_n\}$, where each flow f_i is specified by a triple $(p_i, b_i(t), r_i)$ consisting of its path, bandwidth, and release time. Given two time points $T_1 < T_2$ such that $r_i \leq T_1$ for all $i \in [n]$, there exists a set of feasible constant bandwidths $\{b'_i\}_{1 \leq i \leq n}$ with the same delivery amount over the time interval $[T_1, T_2]$.*

PROOF. Let σ_i denote the amount of flow delivered on path p_i ; that is, $\sigma_i = \int_{T_1}^{T_2} b_i(t) dt$. Define new bandwidths $b'_i = \sigma_i / (T_2 - T_1)$. Clearly, the volume delivered remains the same and release times are respected. Moreover, for any edge e , the sum of flow bandwidths passing through e satisfies $\sum_{i \in P(e)} b'_i \leq c(e)$, where $P(e)$ is the set of paths that use e . Indeed, assume that there exists an edge e' for which $\sum_{i \in P(e')} b'_i > c(e')$. Then, it follows that

$$\begin{aligned} \sum_{i \in P(e')} b'_i &= \sum_{i \in P(e')} \sigma_i / (T_2 - T_1) > c(e') \\ &\Rightarrow \sum_{i \in P(e')} \frac{\int_{T_1}^{T_2} b_i(t) dt}{T_2 - T_1} > c(e') \\ &\Rightarrow \sum_{i \in P(e')} \int_{T_1}^{T_2} b_i(t) dt > (T_2 - T_1) c(e') \end{aligned}$$

which is a contradiction since the total volume delivered can not be more than $(T_2 - T_1)c(e')$. \square

The second lemma (Lemma 2.2) shows that it is possible to discretize the bandwidth usage.

LEMMA 2.2. *Suppose that there exist a capacitated network (G, c) , a path p , and a set of feasible flows $\{f_1, \dots, f_n\}$ over the path, where each flow f_i has a bandwidth $b_i(t)$ and a release time r_i . Given two time points $T_1 < T_2$ such that $r_i \leq T_1$ for all $i \in [n]$, there exists a set of feasible bandwidths $\{b'_i\}_{1 \leq i \leq n}$ such that at any point in time $t \in [T_1, T_2]$, there is at most one active flow (i.e. with non-zero bandwidth). Furthermore, for every flow, the delivery amount over the time interval $[T_1, T_2]$ is as before.*

PROOF. Let $c_m = \min_{e \in p} c(e)$ denote the minimum edge capacity of the path. Also, let σ_i denote the volume corresponding to b_i delivered over the time interval; that is, $\sigma_i = \int_{T_1}^{T_2} b_i(t) dt$.

The idea is to assign full bandwidth to flows one at a time. Set $t_0 = 0$ and $t_i = \sum_{k=0}^{i-1} \frac{\sigma_k}{c_m}$, for $i \in [n]$. Define the new bandwidths as follows. For $i \in [n]$,

$$b''_i(t) = \begin{cases} c_m, & \text{if } t \in (T_1 + t_{i-1}, T_1 + t_i] \\ 0, & \text{otherwise} \end{cases}$$

It is obvious that $b''_i(t)$ is feasible and can be non-zero only in the interval $(T_1 + t_{i-1}, T_1 + t_i]$. Furthermore, this schedule does not require any additional time. Indeed,

$$\begin{aligned} t_n &= \sum_{k=0}^n \frac{\sigma_k}{c_m} = \sum_{k=0}^n \frac{\int_{T_1}^{T_2} b_k(t) dt}{c_m} \\ &= \int_{T_1}^{T_2} \sum_{k=0}^n \frac{b_k(t)}{c_m} dt \leq \int_{T_1}^{T_2} dt = T_2 - T_1. \end{aligned}$$

\square

2.1 Paths Are Given

In this setting, each connection request comes with a path. We give a 17.6-approximation algorithm for this problem. We begin with an observation.

OBSERVATION 2.3. *The circuit-based coflow scheduling with given paths is strongly NP-complete.*

This follows directly from the NP-completeness of the scheduling problem $1|pmtn, r_i| \sum \omega_i c_i$ which seeks to minimize total weighted completion time on a single machine with preemption and release times [18]. Indeed, given an instance of the scheduling problem, consider a single edge $e : s \bullet \rightarrow \bullet t$ such that every job j with a processing time of p_j corresponds to a coflow consisting of just one flow of size $\sigma_j = p_j$ from s to t and with the corresponding release time. Now, Lemma 2.2 states that we can turn a solution returned by solving the circuit-based coflow problem into a preemptive schedule: simply determine the order of release and completion times on the time line and apply the lemma to each interval.

We follow our general framework for coflow scheduling. In order to get the best approximation ratio, we optimize over some parameters such as interval length.

Reformulation. We transform the problem of minimizing total weighted coflow completion time to an instance of minimizing total weighted flow completion time. To capture completion times at the level of coflows, we (a) introduce a dummy flow f_0^i in each coflow F_i , and (b) assign new weights ω'_{ij} to each flow. In more detail, the dummy flow comes with precedence constraints from $f_j^i, j \neq 0$ to f_0^i , requiring that in each coflow, f_0^i finishes last. In effect, the precedence graph is a forest of n depth-1 in-trees. We set the weight to ω_i for the dummy flow f_0^i (i.e. $\omega'_{i0} = \omega_i$) and to zero for all other flows (i.e. $\omega'_{ij} = 0, j \neq 0$). Finally, note that dummy flows do not have source, destination or size but are required to finish no sooner than other flows in the same coflow. Now, the new objective function to be minimized is $\sum_i \sum_j \omega'_{ij} c_j^i$.

The linear program. In this step, we devise an interval-indexed linear program. The time line is divided into segments $[0, 1], (1, 1+\epsilon],$

$(1+\epsilon, (1+\epsilon)^2], \dots, (\tau_\ell, \tau_{\ell+1}]$ for $\ell \in \{0, 1, \dots, L\}$, where $\tau_\ell = (1+\epsilon)^{\ell-1}$, $\tau_0 = 0$, $\epsilon > 0$ will be determined later, and L is a sufficiently large integer.

In the linear program below, $x_{j\ell}^i$ is the portion of connection request f_j^i completed in the ℓ^{th} interval $(\tau_\ell, \tau_{\ell+1}]$ and $b_{j\ell}^i$ is the bandwidth of p_j^i (i.e. f_j^i 's path) during the ℓ^{th} interval. Notice that release times are assumed to occur at the start of intervals; we show that this restriction costs us no more than a factor of $1 + \epsilon$. Constraints (6) ensure that the dummy connection request finishes last in each coflow. Constraints (7) calculate the bandwidth of each path p_j^i for every time interval (which is the fraction of flow size delivered, divided by interval length). Constraints (8) ensure that no capacity violation occurs; note that $P(e)$ denotes the set of connection requests whose paths use edge e . Finally, constraints (9) ensure that release times are respected.

$$\text{Minimize } \sum_i \sum_j \omega_{ij}' c_j^i \quad \text{subject to}$$

$$\sum_\ell x_{j\ell}^i = 1 \quad \forall i, j \quad (4)$$

$$\sum_{\ell \leq L} \tau_\ell x_{j\ell}^i \leq c_j^i \quad \forall i, j \quad (5)$$

$$c_j^i \leq c_0^i \quad \forall i, j \quad (6)$$

$$b_{j\ell}^i = \sigma_j^i x_{j\ell}^i / \tau_\ell \quad \forall i, j, \ell \quad (7)$$

$$\sum_{f_j^i \in P(e)} b_{j\ell}^i \leq c(e) \quad \forall \ell, e \quad (8)$$

$$r_j^i > \tau_{\ell+1} \Rightarrow x_{j\ell}^i = 0 \quad \forall i, \ell \quad (9)$$

$$x_{j\ell}^i \geq 0 \quad \forall i, j, \ell, e \quad (10)$$

LEMMA 2.4. $\frac{1}{1+\epsilon}$ times the optimal value of the above linear program is a lower bound on the weighted sum of completion times for circuit-based coflows with paths and release times.

PROOF. Consider a schedule $\langle p_j^i, b_{j\ell}^i \rangle$ achieving average coflow completion time \hat{C} . We show that it corresponds to a feasible solution to the linear program with an objective value $\leq (1 + \epsilon)\hat{C}$. Start by setting the variables $x_{j\ell}^i$ for each connection request f_j^i according to the interval in which they finish. Next, note that release times are arbitrary; hence, a connection request may arrive in the middle of an interval. On the other hand, constraints (7) require constant bandwidth in each interval. To remedy this, move all release times to the end of the interval in which they occur and update $x_{j\ell}^i$ if necessary; this clearly increases \hat{C} by no more than a factor of $1 + \epsilon$.

Finally, for each interval $(\tau_\ell, \tau_{\ell+1}]$, apply Lemma 2.1 to get constant bandwidth for each connection request. More precisely, we set

$$b_{j\ell}^i = \frac{\int_{\tau_\ell}^{\tau_{\ell+1}} b_j^i(t) dt}{\tau_\ell}$$

for each f_j^i . Clearly, constraints (7) are satisfied; furthermore, the lemma ensures that constraints (8) stay satisfied as well. This setting of variables satisfies all the remaining constraints as well. \square

Rounding. The next step after solving the linear program is to construct a schedule. Let $\{\hat{x}_{j\ell}^i\}_{ij\ell}$ denote an optimal solution to the LP. Set $\hat{b}_{j\ell}^i$ according to equation (7). For some $0 < \alpha \leq 1$, define the α -interval h_{ij}^α of a connection request f_j^i to be the interval $(\tau_{h_{ij}^\alpha}, \tau_{h_{ij}^\alpha+1}]$ which contains the earliest point in time that a cumulative α -fraction of the connection request has been completed; that is

$$h_{ij}^\alpha = \min \left\{ \ell : \sum_{t=0}^{\ell} x_{jt}^i \geq \alpha \right\}.$$

The exact value of α will be determined later.

In the schedule, each connection request f_j^i will run entirely in $(\tau_{h_{ij}^\alpha+D}, \tau_{h_{ij}^\alpha+D+1}]$ —that is, the D^{th} interval after its α -interval. The parameter $D \in \mathbb{N}$, the *displacement factor*, will be determined later. Note that the length of this interval is $(1 + \epsilon)^{D-1}$ times the cumulative length of intervals from the beginning to the end of the α -interval.

Fix a k . Let $S[k]$ denote the set of connection requests scheduled to run in $(\tau_k, \tau_{k+1}]$. Clearly, the α -interval of all these connection requests is $k - D$. Consider a connection request $f_j^i \in S[k]$. According to the optimal solution obtained by solving the LP, its bandwidth function is

$$\hat{b}_j^i(t) = \begin{cases} \hat{b}_{j0}^i, & \text{if } t \in [0, 1], \\ \hat{b}_{j1}^i, & \text{if } t \in (\tau_1, \tau_2], \\ \vdots & \vdots \\ \hat{b}_{jk}^i, & \text{if } t \in (\tau_{k-D}, \tau_{k-D+1}]. \end{cases}$$

We now define a new constant bandwidth

$$\tilde{b}_j^i = \sum_{\ell=0}^{\ell=k-D} \frac{\alpha^{-1} \hat{b}_{j\ell}^i \tau_\ell}{\tau_k} = \sum_{\ell=0}^{\ell=k-D} \frac{\alpha^{-1} \hat{b}_{j\ell}^i}{(1 + \epsilon)^{k-\ell}}. \quad (11)$$

Observe that we have scaled and added up each bandwidth value depending on how far its interval is from the k^{th} interval. For each connection request in $S[k]$, at least an α fraction of its size σ_j^i was delivered by the end of the α -interval. But now, with the new bandwidth, the entire flow is delivered over the stretched time line. More precisely,

$$\forall f_j^i \in S[k] : \int_0^{\tau_{h_{ij}^\alpha+1}} \hat{b}_j^i(t) dt = \int_0^{\tau_{k-D+1}} \tilde{b}_j^i(t) dt = \alpha \tilde{b}_j^i \tau_k.$$

This also means that the new bandwidth and the displacement factor have to satisfy

$$D \geq \left\lceil \log_{1+\epsilon} \frac{1}{\alpha} \right\rceil + 1. \quad (12)$$

Furthermore, we must ensure that no capacity violation occurs. In the schedule, for every k and for every edge e ,

$$\begin{aligned} \sum_{f_j^i \in P(e)} \tilde{b}_j^i &= \sum_{f_j^i \in P(e)} \sum_{\ell=0}^{\ell=k-D} \frac{\alpha^{-1} \hat{b}_{j\ell}^i}{(1+\epsilon)^{k-\ell}} \\ &= \sum_{\ell=0}^{\ell=k-D} \frac{\alpha^{-1}}{(1+\epsilon)^{k-\ell}} \sum_{f_j^i \in P(e)} \hat{b}_{j\ell}^i \\ &\leq c(e) \sum_{\ell=0}^{\ell=k-D} \frac{\alpha^{-1}}{(1+\epsilon)^{k-\ell}}, \end{aligned}$$

where the inequality follows from constraints (8). Thus, we require that $\sum_{\ell=0}^{\ell=k-D} \frac{\alpha^{-1}}{(1+\epsilon)^{k-\ell}} \leq 1$, for all k . Hence,

$$\frac{1}{\epsilon(1+\epsilon)^{D-1}} \leq \alpha. \quad (13)$$

Finally, note that the blow up factor in completion time for each connection request in $S[k]$, taking into account the $(1+\epsilon)$ -factor loss in the LP, is at most

$$\frac{(1+\epsilon)\tau_{k+1}}{(1-\alpha)\tau_{k-D}} = \frac{(1+\epsilon)^{D+2}}{1-\alpha}. \quad (14)$$

Numerically minimizing (14) subject to (13) and displacement factor constraint (12), we get an approximation factor of 17.5319 for $\alpha = 0.5$, $D = 3$, and $\epsilon \approx 0.5436$.

2.2 Paths Are Not Given

In this section, we consider circuit-based coflow scheduling when paths are not given. Our $O(\frac{\log |E|}{\log \log |E|})$ approximation ratio for this case is asymptotically tight since an $\Omega(\frac{\log |E|}{\log \log |E|})$ hardness follows from hardness of congestion minimization on directed graphs [8]. Recall that, in the congestion minimization problem, given a directed graph and source-sink pairs, the goal is to connect each pair by a path such that congestion, the maximum number of paths crossing an edge, is minimized. But, regarding the source-sink pairs as connection requests of unit size in a single coflow, the congestion is equivalent to makespan, where all edge capacities are set to 1. So, essentially, congestion minimization can be seen as a special case of our problem. We now go over our algorithm.

Algorithm 1: CIRCUIT-BASED COFLOW SCHEDULING

input : Coflows \mathcal{F} , Network (G, c)

- 1 Construct the LP \mathcal{L} ;
- 2 Solve \mathcal{L} ; perform rounding and scaling to get flow values x_e^f and completion times c_f ;
- 3 **foreach** flow $f \in \mathcal{F}$ **do**
- 4 PathSet[f] = FlowDecomposition(f);
- 5 Path[f] = Rounding(PathSet[f]);
- 6 **end**

return : Flow paths and ordering based on c_f

Reformulation. Identical to §2.1.

The linear program. The linear program is similar to §2.1, where we have set $\epsilon = 1$. Notice that constraints (18)-(21) are the usual flow constraints per time interval.

$$\begin{aligned} &\text{Minimize } \sum_i \sum_j \omega_{ij}^i c_j^i && \text{subject to} \\ &\sum_{\ell} x_{j\ell}^i = 1 && \forall i, j \end{aligned} \quad (15)$$

$$\sum_{\ell \leq L} \tau_{\ell} x_{j\ell}^i \leq c_j^i \quad \forall i, j \quad (16)$$

$$c_j^i \leq c_0^i \quad \forall i, j \quad (17)$$

$$\sum_{\substack{e \in N(v) \\ v \notin \{s_j^i, d_j^i\}}} x_{ij\ell}^e = 0 \quad \forall i, j, \ell \quad (18)$$

$$\sum_{e \in N(d_j^i)} x_{ij\ell}^e = \sigma_j^i x_{j\ell}^i / \tau_{\ell} \quad \forall i, j, \ell \quad (19)$$

$$\sum_{e \in N(s_j^i)} x_{ij\ell}^e = - \sum_{e \in N(d_j^i)} x_{ij\ell}^e \quad \forall i, j, \ell \quad (20)$$

$$\sum_i \sum_j x_{ij\ell}^e \leq c(e) \quad \forall \ell, e \quad (21)$$

$$r_j^i > \tau_{\ell+1} \Rightarrow x_{j\ell}^i = 0 \quad \forall i, j, \ell \quad (22)$$

$$x_{ij\ell}^e, x_{j\ell}^i \geq 0 \quad \forall i, j, \ell, e \quad (23)$$

Next, we show that the optimal value of the linear program gives us a relaxed lower bound.

LEMMA 2.5. *Half the optimal value of the above linear program provides a lower bound on the weighted sum of completion times for circuit-based coflows with release times.*

PROOF. Following the proof of Lemma 2.4, move release times if necessary, set the variables $x_{j\ell}^i$ and calculate bandwidths per interval. Next, set the value of flow variables $x_{ij\ell}^e$ based on the paths p_j^i and the calculated bandwidths. This setting of variables satisfies all constraints. \square

Rounding. This step involves scaling followed by flow decomposition and randomized rounding. Let $\langle \hat{x}_{ij\ell}^e, \hat{x}_{\ell}^i \rangle$ denote an optimal solution to the LP. Analogous to the previous case in §2.1, define the half interval h_j^i of each connection request f_j^i . All connection requests whose half interval is H will be scheduled to run entirely in interval $(\tau_{H+3}, \tau_{H+4}]$. In analogy with §2.1, set $\alpha = \frac{1}{2}$ and $D = 3$. Now, fix a k and let $S[k]$ denote the set of connection requests scheduled to run in $(\tau_k, \tau_{k+1}]$. Perform scaling and adding up of flows for every edge to get updated flow variables

$$\tilde{x}_{ij}^e = \sum_{\ell=0}^{\ell=k-3} \frac{\hat{x}_{ij\ell}^e}{2^{k-\ell-1}}. \quad (24)$$

The calculations in §2.1 apply here as well (where we use inequalities (21) instead of (8)). Hence, the entire size of every connection request is delivered in the new settings and all edge capacities are respected. The next step is to apply the well-known flow decomposition theorem (see e.g. [1]) to the flow variables \tilde{x}_{ij}^e to get a set

of flow paths $P_{ij} = \{p_j^i\}$ from s_j^i to d_j^i for each connection request $f_j^i \in S[k]$.

At this point, after computing the flow paths, we have for each k a set of connection requests $S[k]$ expressed in terms of flow paths. Unfortunately, this is not enough since each connection request must have exactly one path from source to sink. What we have, instead, is a set of flow paths $P_{ij} = \{p_{j1}^i, \dots, p_{jn_{ij}}^i\}$ for each connection request f_j^i . This brings us to the final step.

Let b_j^i denote the total amount of flow in P_{ij} :

$$b_j^i = \sum_{k \in [n_{ij}]} |p_{jk}^i|.$$

We apply the randomized rounding technique of Raghavan and Thompson [25] to choose a single path from P_{ij} : we randomly choose a flow path from P_{ij} where the probability of p_{jk}^i being selected is $|p_{jk}^i|/b_j^i$. All of the flow in P_{ij} will go over the chosen path. Let \bar{b}_e denote the total flow passing through e after the rounding. Using a Chernoff-Hoeffding bound, we can show that for every $e \in E$ and every $\delta > 0$:

$$\Pr [\bar{b}_e \geq (1 + \delta)c(e)] \leq \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{c(e)/|E|}.$$

Hence,

$$\begin{aligned} \Pr \left[\max_e \frac{\bar{b}_e}{c(e)} \geq (1 + \delta) \right] &\leq \sum_e \Pr \left[\frac{\bar{b}_e}{c(e)} \geq (1 + \delta) \right] \\ &\leq \sum_e \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{c(e)/|E|} \\ &\leq |E| \left(\frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^{(\min_e c(e))/|E|}. \end{aligned}$$

If we set $\delta = O\left(\frac{\log |E|}{\log \log |E|}\right)$, we obtain that the bandwidth used exceeds the capacity of any edge by at most $O\left(\frac{\log |E|}{\log \log |E|}\right)$ with probability at most $1/n^c$ for some constant $c > 0$. Consequently, by scaling down the bandwidth allocated by $1 + \delta$, and scaling up in time by the same factor, we obtain a feasible set of connection requests with a blowup in completion times at most $O\left(\frac{\log |E|}{\log \log |E|}\right)$.

3 PACKET-BASED COFLOWS

In this model, a flow corresponds to a *single* packet (flow size is 1) and a coflow is a collection of packets. The schedule is based on discrete time steps. During a step, each packet can either stay at a node (in a queue) or move along an edge with the restriction that at most one packet can use the edge at a time. In §3.1, we observe that when paths are given, we can apply known results for job shop scheduling. Next, in §3.2, we consider the case where paths are not given. We present an $O(1)$ -approximation algorithm for this case.

3.1 Paths Are Given

When a path p_j^i is given for each packet f_j^i , the only remaining problem is scheduling. It is well-known that the problem of scheduling packets on a store-and-forward network with fixed paths is an instance of job-shop scheduling. Specifically, each packet f_j^i can be regarded as a job which needs to be executed on a number of

machines (edges) in a prescribed order (given by p_j^i) such that a machine can process one job at a time. Consequently, the packet-based coflow problem with given paths, is an instance of shop scheduling $J|r_j, p_{ij} = 1 | \sum_S \omega_S C_S$ with unit processing times and a *generalized* min-sum objective; hence, the algorithm of [24] immediately gives an $O(1)$ -approximation.

THEOREM 3.1 ([24]). *There exists a polynomial-time $O(1)$ approximation algorithm for the (generalized) min-sum job-shop problem $J(P)|r_j, p_{ij} = 1 | \sum_j \omega_j C_j$.*

3.2 Paths Are Not Given

This is the more general case where one needs to both route and schedule packets. We give a polynomial-time $O(1)$ -approximation algorithm for this problem.

An ingredient in our solution is the notion of *time-expanded graphs* first introduced by Ford and Fulkerson [17]. In more detail, given a (directed) graph $G = (V, E)$ and n coflows $\{F_i = \{f_1^i, \dots, f_{n_i}^i\}\}_{i \in [n]}$, we construct a time-expanded graph G^T as follows. Each node in the time-expanded graph is labeled by a tuple (v, t) where v is a node in G and $0 \leq t \leq T$ is the timestamp. Therefore, the edge set of G^T is

$$\begin{aligned} E(G^T) = &\{(u, t), (v, t + 1)\} : (u, v) \in E(G), 0 \leq t \leq T - 1\} \\ &\cup \{(v, t), (v, t + 1)\} : v \in V(G), 0 \leq t \leq T - 1\} \end{aligned}$$

where the second group of edges are called queue edges and their role is to simulate packets waiting for one or more rounds at a node. Let us denote by E' the set $E(G^T)$ excluding the queue edges.

The overall idea of the algorithm is as follows. By a net-flow we mean a network flow in the ordinary sense. We express each packet as multiple net-flows to the packet's destination in different times in G^T ; this will become clear shortly. The remaining steps involve writing an interval-indexed linear program, assigning packets to intervals, and performing routing and scheduling packets for each interval.

Reformulation. Consider a packet $f_j^i \in F_i$ to be routed from s_j^i to d_j^i in G . Corresponding to this packet, we send a set of net-flows f_{ij}^t in G^T : net-flow f_{ij}^t has node (s_j^i, r_j^i) as source (demand is -1) and node (d_j^i, t) as destination where $t \in \{r_i + 1, \dots, T\}$ and sum of demands of these nodes is 1 (the exact values will be determined by solving the linear program below). Also, we set T to a sufficiently large integer (say $|E| \sum_{i \in [n]} n_i$). As before, we introduce dummy flows f_0^i and update weights ω'_i accordingly.

The linear program. Let $T' = \log T$, $\tau_0 = 1$ and $\tau_\ell = 2^{\ell-1}$ for $\ell \in [T']$. Let c_j^t denote the earliest time when a net-flow f_{ij}^t in G^T reaches its destination; that is $c_j^t = \min_t c_{ij}^t$ where c_{ij}^t is the completion time of f_{ij}^t . Variables $x_e^{f_{ij}^t}$ denote the amount of flow corresponding to f_{ij}^t over edge $e \in E(G^T)$. In constraints (25), $f_{ij\ell}$ is the total amount of flow, corresponding to net-flows f_{ij}^t , that enter their destination nodes (d_j^i, t) for $t \in (\tau_\ell, \tau_{\ell+1}]$. By $E((d_{ij}, t))$ we mean the set of incoming edges to (d_{ij}, t) in constraints (25). Moreover, b_{ij}^t is the demand of node (d_j^i, t) and $\mathcal{N}_{ij}^{f_{ij}^t} x_e^{f_{ij}^t} = b_{ij}^t$ specify the obvious requirements of sending a net-flow of size b_{ij}^t

from (s_j^i, r_i) to (d_j^i, t) for $r_i \leq t \leq T$, as well as flow conservation at nodes other than source and sink.

$$\begin{aligned} & \text{Minimize } \sum_{i \in [n]} \sum_{j=0}^{n_i} \omega'_{ij} c_j^i & \text{subject to} \\ & \sum_{\tau_\ell < t \leq \tau_{\ell+1}} \sum_{e \in E((d_{ij}, t))} x_e^{f_{ij}^t} = f_{ij\ell} & \forall i, j, \ell \end{aligned} \quad (25)$$

$$\sum_{\ell \leq T'} \tau_\ell f_{ij\ell} \leq c_{ij} \quad \forall i, j \quad (26)$$

$$c_{ij} \leq c_{i0} \quad \forall i, j \quad (27)$$

$$\sum_{i \in [n]} \sum_{j=0}^{n_i} \sum_{t \leq \tau_{\ell+1}} x_e^{f_{ij}^t} \leq \tau_{\ell+1} \quad \forall \ell, e \quad (28)$$

$$\sum_{e \in E'} \sum_{t \leq \tau_{\ell+1}} x_e^{f_{ij}^t} \leq \tau_{\ell+1} \quad \forall i, j, \ell \quad (29)$$

$$\sum_{t \in T} b_{ij}^t = 1 \quad \forall i, j \quad (30)$$

$$\mathcal{N}^{f_{ij}^t} x^{f_{ij}^t} = b^{f_{ij}^t} \quad \forall i, j \quad (31)$$

$$0 \leq x_e^{f_{ij}^t} \leq 1 \quad \forall i, j, t, e \quad (32)$$

Constraints (28) and (29) restrict respectively the congestion and dilation of packets which reach their destination in interval $[0, \tau_{\ell+1}]$. As can be seen, queue edges don't contribute to path dilation.

LEMMA 3.2. *The optimal value of the linear program (25)-(32) with weights $\omega'_{i0} = \omega_i$ and $\omega'_{ij} = 0, j \neq 0$, is a lower bound on the optimal weighted sum of packet-based coflow completion times.*

PROOF. Given a set $S = \langle \{p_j^i\}, \{q_j^i\} \rangle_{i \in [n], j \in [n_i]}$ of paths $\{p_j^i\}$ and schedules $\{q_j^i\}$, we show that S constitutes a feasible solution to the LP. Without loss of generality, let g denote the maximum amount of time any packet has to wait in the schedule; the existence of $g \in O(1)$ follows from [19] (note that g is the maximum queue size.) Each packet f_j^i is routed, on a path of length $|p_j^i|$, from $s_j^i \in V$ to $d_j^i \in V$ according to the schedule $\{q_j^i\}$ which specifies the amount of time $\in [0, g]$ the packet has to wait before crossing the next edge on its path. Let T_j^i denote the time step when f_j^i reaches its sink. Furthermore, we define $T = \max_{ij} T_j^i$.

For each packet f_j^i , we construct an integral flow $f_{ij}^{T_{ij}}$ in the time-expanded graph G^T . Consider a packet f_j^i with release time r_j^i in G going from s_j^i to d_j^i on a path specified by p_j^i according to the schedule q_j^i and arriving at time T_j^i . The corresponding flow, $f_{ij}^{T_{ij}}$, in G^T from (s_j^i, r_j^i) to (d_j^i, T_j^i) is constructed in steps. We keep a current time step t which is initially set to r_j^i . For each edge $e = (u, v)$ on the path p_j^i , starting from s_j^i , the flow f_j^i follows the edge(s)

$$((u, t), (u, t+1)) \rightarrow \dots \rightarrow ((u, t+h), (v, t+h+1)),$$

in G^T where h is the delay specified by q_j^i before crossing the edge. Next, we update $t \leftarrow t+1$ and repeat until we reach d_j^i . Clearly,

T_j^i is the sum of the path length $|p_j^i|$ and total amount of delay prescribed by the schedule q_j^i .

It follows from the construction that constraints (31) are satisfied. Furthermore, since the schedule does not allow more than one packet to cross an edge at the same time, the path followed by a flow is exclusive to that flow except, possibly, the edges $((u, t), (u, t+1))$ which effectively simulate queues. At this point, we can set the value of variables $x_e^{f_{ij}^t}$ according to the flow $f_{ij}^{T_{ij}}$. We also set $c_j^i = T_j^i$ and $c_0^i = \max_{j \in [n_i]} c_j^i$. It is easy to verify that constraints (25)-(27) are satisfied. Regarding congestion constraints (28), observe that, for all ℓ , packets that have reached their destination by time $\tau_{\ell+1}$, can not cause an edge to have a congestion value more than $\tau_{\ell+1}$. The same argument applies to dilation constraints (29). \square

Rounding. At a high-level, this step consists of two parts (a) Assign the packets to time intervals; (b) Route and schedule packets in each interval by applying the algorithm of Srinivasan and Teo [26].

Let $\hat{S} = \langle \hat{x}_e^{f_{ij}^t}, \hat{f}_{ij\ell}, \hat{c}_j^i \rangle$ denote an optimal solution to the LP (25)-(32). We start by filtering the solution and assigning packets to intervals. Specifically, define the half-interval of a packet f_j^i to be the interval $(\tau_{h_j^i}, \tau_{h_j^i+1}]$ such that $h_j^i = \min \{ \ell : \sum_{t < \ell} \hat{f}_{ij\ell} \geq \frac{1}{2} \}$; then, set

$$\bar{f}_{ij\ell} = \begin{cases} \hat{f}_{ij\ell} / (1 - \sum_{t \geq \ell} \hat{f}_{ij\ell}), & \text{if } \tau_{\ell+1} \leq h_j^i, \\ 0, & \text{if } \tau_{\ell+1} \geq h_j^i. \end{cases}$$

Also, set the other variables $\bar{x}_e^{f_{ij}^t}$ and \bar{c}_j^i according to (25) and (26) respectively. Clearly, the blowup in each variable and the right hand side of (28)-(29) is at most 2.

Next, we assign each packet f_j^i to its half-interval. Let $P[\ell]$ denote the set of packets that have been assigned to the interval $(\tau_\ell, \tau_{\ell+1}]$. At this stage, we go over all time intervals sequentially and, for each one, route and schedule all packets in it. Fix an ℓ and consider all packets $f_j^i \in P[\ell]$. Collapse the portion of the time-expanded graph G^T , corresponding to the times $t \leq \tau_{\ell+1}$, back to G by combining nodes and edges which differ only in time stamp and by removing queue edges altogether.

Observe that the removal of queue edges will not cause any problem since any amount of flow passing through them will have left the queue by the end of the last interval. Now, the filtered solution $\bar{S} = \langle \bar{x}_e^{f_{ij}^t}, \bar{f}_{ij\ell}, \bar{c}_j^i \rangle$ satisfies

$$\sum_{i, j: f_j^i \in P[\ell]} \bar{x}_e^{f_{ij}^t} \leq \tau_{\ell+2} \quad \forall e \in E \quad (33)$$

$$\sum_{e \in E} \bar{x}_e^{f_{ij}^t} \leq \tau_{\ell+2} \quad \forall f_j^i \in P[\ell] \quad (34)$$

$$\bar{\mathcal{N}}^{f_{ij}^t} \bar{x}^{f_{ij}^t} = \bar{b}^{f_{ij}^t} \quad \forall f_j^i \in P[\ell] \quad (35)$$

$$0 \leq \bar{x}_e^{f_{ij}^t} \leq 1 \quad \forall f_j^i \in P[\ell] \quad \forall e \in E \quad (36)$$

where the constraints $\bar{\mathcal{N}}^{f_{ij}^t} \bar{x}^{f_{ij}^t} = \bar{b}^{f_{ij}^t}$ express the flow requirements in the collapsed graph. Formally,

LEMMA 3.3. For every $\ell \in [T']$, the set of flows $\bar{x}_e^{f_{ij}}$ corresponding to packets f_j^i in $P[\ell]$ satisfy the LP in (33)-(36) with respect to the (collapsed) graph G .

This LP is exactly the one considered in [26] and by applying their theorem, we can route and schedule all packets in $P[\ell]$ such that the last packet arrives in time $O(\tau_{\ell+2})$.

THEOREM 3.4 ([26]). There are constants $c', c'' > 0$ such that the following holds. For any packet routing problem on any network, there is a set of paths and a corresponding schedule that can be constructed in polynomial time such that the routing time is at most c' times the optimal, and the maximum queue size at each edge is bounded by c'' .

We do this for $P[1]$, $P[2]$ and so on. Note that the coflow release times are respected. The optimal completion time of a packet $f_j^i \in P[\ell]$ in the LP (25)-(32) is $\hat{c}_j^i = O(\tau_{\ell+1})$. On the other hand, the completion time of the same packet in our algorithm is

$$\tilde{c}_j^i = \sum_{t \leq \ell} O(\tau_{t+1}) = \sum_{t \leq \ell} O(2^t) = O(2^\ell) = O(\tau_{\ell+1}). \quad (37)$$

Consequently, the weighted completion time of our algorithm is within a $O(1)$ factor of the optimal solution.

4 EXPERIMENTS

In this section, we describe some experiments conducted to evaluate the practical performance of our circuit-based algorithm in §2.2.

4.1 Methodology

We use a simulator to assess the performance of our algorithm in practice. Packet-level simulators are not suitable for this setting due to their high level of overhead [3, 29]. Therefore, like previous works [3, 7, 29], we developed a flow-based simulator. At a high level, the simulator is an event queue. Each flow corresponds to an event which happens at its release time. The simulator chooses the next flow based on the ordering prescribed by a scheduling algorithm or scheme. A second event occurs when a flow completes; at which time, its reserved bandwidth is released.

We use a 128-server Fat-Tree with 1Gbs links as the network topology. Due to the complexity of the linear program to be solved, simulations of large instances were prohibitively slow. Each coflow instance is randomly generated with flow release times, flow sizes, and coflow weights based on Poisson distributions. Each result is the average of 10 tries.

4.2 Implementation

We implement Algorithm 1 in §2.2 with some minor tweaks. In particular, to avoid wasting time and bandwidth, each flow starts as soon as it can (in the order prescribed by the linear program), as opposed to starting at the beginning of its half-interval.

The path decomposition algorithm tries to minimize the number of paths per flow by finding the “thickest” paths; this is done using a well-known version of Dijkstra’s shortest-path algorithm. We implemented the above algorithm in C++ using Lemon libraries and IBM CPLEX 12.6.3. to solve the linear program.

4.3 Performance

We measure the weighted sum of completion times. As in [29], we compare our algorithm to the following schemes.

- Baseline: flows are routed and ordered randomly.
- Schedule-only: flows are routed randomly; ordering is by minimum completion time which is computed as the ratio of flow size to path bandwidth.
- Route-only: flows are routed for achieving good load balance and edge utilization; ordering is arbitrary.

We investigate the impact of two parameters. First is the number of flows in each coflow, which we refer to as *coflow width*. Second is the number of coflows.

Coflow width. In this case, we fix the number of coflows to 10 and run experiments for coflow widths in $\{4, 8, 16, 32\}$. As illustrated in Figure 2, LP-BASED completes *every instance* faster than the baseline, Schedule-only, and Route-only schedules. On average, the improvement over the baseline, Schedule-only and Route-only is by %126, %96, and %22, respectively.




Figure 2: Changing the coflow width; the number of coflows is set to 10.

Number of coflows. Using a fixed coflow width of 16, we vary the number of coflows from 10 to 25, in increments of 5. The results are illustrated in Figure 3. We find that again LP-BASED completes *every*

instance faster than the baseline as well as the Schedule-only and Route-only schedules. On average, the improvement over the baseline, Schedule-only and Route-only is by factors of %110, %72, and %26, respectively. We note that, as can be seen in the experiments,




Figure 3: Changing the number of coflows; coflow width is set to 16.

the worst-case approximation ratio of $O(\log |E|/\log \log |E|)$, does not happen in practice. The main reason is that in the fat-tree topology there are a few paths between pairs of servers; furthermore, uniform link bandwidth causes the path decomposition routine to return fewer flow paths. Indeed, in all of our experiments, the path decomposition routine returns one path per flow.

5 CONCLUSION AND OPEN PROBLEMS

We have presented algorithms for scheduling task-based coflows, packet-based coflows, and circuit-based coflows for general networks. All of our algorithms achieve asymptotically-optimal total weighted completion time. An immediate line of further research is to obtain improved constant-factor bounds for general networks and improved algorithms for special network classes that are of particular interest to datacenter applications. More broadly, we would like to pursue algorithms for other important objectives and coflow models. Specifically, we are interested in designing approximation algorithms for the total weighted response time objective, where

the response time is the difference between completion time and release time. Another line of research is the design of schedules for coflows in optical networks.

6 ACKNOWLEDGMENTS

We would like to thank the reviewers for their useful comments and pointing out the relevant work done in [9].

REFERENCES

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. 1993. *Network Flows: Theory, Algorithms, and Applications*.
- [2] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat. 2010. Hedera: Dynamic Flow Scheduling for Data Center Networks. In *NSDI*.
- [3] Mohammad Al-Fares, Sivasankar Radhakrishnan, Barath Raghavan, Nelson Huang, and Amin Vahdat. 2010. Hedera: Dynamic Flow Scheduling for Data Center Networks. In *USENIX*.
- [4] J. D. Blocher and D. Chhajed. 1996. The customer order lead-time problem on parallel machines. *Naval Research Logistics* 43, 5 (1996).
- [5] M. Chowdhury and I. Stoica. 2012. Coflow: A Networking Abstraction for Cluster Applications. In *HotNets*.
- [6] M. Chowdhury, M. Zaharia, J. Ma, M. I. Jordan, and I. Stoica. 2011. Managing data transfers in computer clusters with orchestra. In *SIGCOMM*.
- [7] M. Chowdhury, Y. Zhong, and I. Stoica. 2014. Efficient Coflow Scheduling with Varys. *SIGCOMM, Comput. Commun. Rev.* 44, 4 (2014).
- [8] J. Chuzhoy, V. Guruswami, S. Khanna, and K. Talwar. 2007. Hardness of routing with congestion in directed graphs. In *STOC*.
- [9] José R. Correa, Martin Skutella, and José Verschae. 2009. *The Power of Preemption on Unrelated Machines and Applications to Scheduling Orders*. Springer Berlin Heidelberg, Berlin, Heidelberg, 84–97. https://doi.org/10.1007/978-3-642-03685-9_7
- [10] J. Dean and S. Ghemawat. 2008. MapReduce: simplified data processing on large clusters. *Commun. ACM* 51, 1 (2008).
- [11] N. Garg, A. Kumar, and V. Pandit. 2007. Order Scheduling Models: Hardness and Algorithms. In *FSTTCS*.
- [12] R.L. Graham, E.L. Lawler, J.K. Lenstra, and A.H.G. Rinnooy Kan. 1979. Optimization and Approximation in Deterministic Sequencing and Scheduling: a Survey. *Discrete Mathematics*, Vol. 5.
- [13] C. Hong, M. Caesar, and B. Godfrey. Finishing flows quickly with preemptive scheduling. In *ACM Conference, SIGCOMM*.
- [14] M. Isard, M. Budiu, Y. Yu, A. Birrell, and D. Fetterly. 2007. Dryad: distributed data-parallel programs from sequential building blocks. In *Proc. of the EuroSys Conference*.
- [15] S. Khuller and M. Purohit. 2016. Improved Approximation Algorithms for Scheduling Co-Flows. In *SPAA*. Brief Announcement.
- [16] R. Koch, B. Peis, M. Skutella, and A.s Wiese. 2009. Real-Time Message Routing and Scheduling. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*. Lecture Notes in Computer Science, Vol. 5687.
- [17] Jr. L. R. Ford and D. R. Fulkerson. 1958. Constructing Maximal Dynamic Flows from Static Flows. *Operations Research* 6, 3 (1958).
- [18] J. Labetoulle, E.L. Lawler, J.K. Lenstra, and A.H.G. Rinnooy Kan. 1984. Preemptive Scheduling of Uniform Machines Subject to Release Dates. In *Progress in Combinatorial Optimization*, William R. Pulleyblank (Ed.). Academic Press.
- [19] F.T. Leighton, B.M. Maggs, and S.B. Rao. 1994. Packet routing and job-shop scheduling in $O(\text{congestion} + \text{dilation})$ steps. *Combinatorica* (1994).
- [20] F. T. Leighton, B.M. Maggs, and A.W. Richa. 1999. Fast Algorithms for Finding $O(\text{Congestion} + \text{Dilation})$ Packet Routing Schedules. *Combinatorica* 19, 3 (1999).
- [21] J. Y-T. Leung, H. Li, and M. Pinedo. 2005. *Order Scheduling Models: An Overview*.
- [22] B. Peis, M. Skutella, and A. Wiese. 2009. Packet Routing: Complexity and Algorithms. In *WAOA*.
- [23] Z. Qiu, C. Stein, and Y. Zhong. 2015. Minimizing the Total Weighted Completion Time of Coflows in Datacenter Networks. In *SPAA*.
- [24] M. Queyranne and M. Sviridenko. 2002. Approximation Algorithms for Shop Scheduling Problems with Minsum Objective. *Journal of Scheduling* (2002).
- [25] P. Raghavan and C. Thompson. 1987. Randomized Rounding: A Technique for Provably Good Algorithms and Algorithmic Proofs. *Combinatorica* 7 (1987).
- [26] A. Srinivasan and C.-P. Teo. 2001. A Constant-Factor Approximation Algorithm for Packet Routing and Balancing Local vs. Global Criteria. *SIAM J. Comput.* 30, 6 (2001).
- [27] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. 2010. Spark: Cluster Computing with Working Sets. In *HotCloud*.
- [28] D. Zats, T. Das, P. Mohan, D. Borthakur, and R. H. Katz. 2012. DeTail: reducing the flow completion time tail in datacenter networks. In *SIGCOMM*.
- [29] Y. Zhao, K. Chen, W. Bai, M. Yu, C. Tian, Y. Geng, Y. Zhang, D. Li, and S. Wang. 2015. Rapier: Integrating routing and scheduling for coflow-aware data center networks. In *INFOCOM*.