

Flow-time Optimization for Concurrent Open-Shop and Precedence Constrained Scheduling Models

Janardhan Kulkarni

Microsoft Research, Redmond, WA, USA
jakul@microsoft.com

Shi Li¹

Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY, USA
shil@buffalo.edu

Abstract

Scheduling a set of jobs over a collection of machines is a fundamental problem that needs to be solved millions of times a day in various computing platforms: in operating systems, in large data clusters, and in data centers. Along with makespan, *flow-time*, which measures the length of time a job spends in a system before it completes, is arguably the most important metric to measure the performance of a scheduling algorithm. In recent years, there has been a remarkable progress in understanding flow-time based objective functions in diverse settings such as unrelated machines scheduling, broadcast scheduling, multi-dimensional scheduling, to name a few.

Yet, our understanding of the flow-time objective is limited mostly to the scenarios where jobs have no dependencies. On the other hand, in almost all real world applications, think of MapReduce settings for example, jobs have dependencies that need to be respected while making scheduling decisions. In this paper, we take first steps towards understanding this complex problem. In particular, we consider two classical scheduling problems that capture dependencies across jobs: 1) *concurrent open-shop scheduling* (COSSP) and 2) *precedence constrained scheduling*. Our main motivation to study these problems specifically comes from their relevance to two scheduling problems that have gained importance in the context of data centers: *co-flow scheduling* and *DAG scheduling*. We design almost optimal approximation algorithms for COSSP and PCSP, and show hardness results.

2012 ACM Subject Classification Theory of computation → Scheduling algorithms

Keywords and phrases Approximation, Weighted Flow Time, Concurrent Open Shop, Precedence Constraints

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2018.16

Related Version A full version of the paper is available at <https://arxiv.org/abs/1807.02553>.

1 Introduction

Scheduling a set of jobs over a collection of machines is a fundamental problem that needs to be solved millions of times a day in various computing platforms: in operating systems, in large data clusters, and in data centers. Along with makespan, *flow-time*, which measures the length of time a job spends in a system before completing, is arguably the most important

¹ The work of the author is in part supported by NSF grants CCF-1566356 and CCF-1717134.



© Janardhan Kulkarni and Shi Li;

licensed under Creative Commons License CC-BY

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2018).

Editors: Eric Blais, Klaus Jansen, José D. P. Rolim, and David Steurer; Article No. 16; pp. 16:1–16:21



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

metric to measure the performance of a scheduling algorithm. In recent years, there has been a remarkable progress in understanding flow-time related objective functions in diverse settings such as unrelated machines scheduling [22, 31, 5, 29, 10], broadcast scheduling [7, 32, 9], multi-dimensional scheduling [28, 27], to name a few.

Yet, our understanding of the flow-time based objective functions is mostly limited to the scenarios where jobs have simple structures; in particular, each job is a single self contained entity. On the other hand, in almost all real world applications, jobs have more complex structures. Consider the MapReduce model for example. Here, each job consists of a set of *Map* tasks and a set of *Reduce* tasks. Reduce tasks cannot be processed unless Map tasks are completely processed². A MapReduce job is complete only when all Map and Reduce tasks are completed. Motivated by these considerations, in this paper, we consider two classical scheduling models that capture more complex job structures: 1) *concurrent open-shop scheduling* (COSSP), and 2) *precedence constrained scheduling* (PCSP). Our main reason to study these problems specifically comes from their relevance to two scheduling problems that have gained importance in the context of data centers: *co-flow scheduling and DAG scheduling*. We discuss more about how these problems relate to COSSP and PCSP in Section 1.3.

The objective function we consider in this paper is minimizing the *sum of general delay costs of jobs*, first introduced in an influential paper by Bansal and Pruhs [11] in the context of single machine scheduling. In this objective, for each job j we are given a *non-decreasing* function $g_j : \mathbb{Z}_+ \rightarrow \mathbb{Z}_+$, which gives the cost of completing the job at time t . The goal is to minimize $\sum_j g_j(C_j)$, where C_j is the completion time of j . A desirable aspect of the general delay cost functions is that they capture several widely studied flow-time and completion time based objective functions.

- Minimizing the sum of weighted flow-times of jobs. This is captured by the function $g_j(t) = w_j \cdot (t - r_j)$, where r_j is the release time of j .
- Minimizing the sum of weighted p th power of flow-times of jobs. This is captured by the function $g_j(t) = w_j \cdot (t - r_j)^p$.
- Minimizing the sum of weighted tardiness. This is captured by the function $g_j(t) = w_j \cdot \max\{0, (t - d_j)\}$, where d_j is the deadline of j .

In this paper, we design approximation algorithms for minimizing the sum of general delay costs of jobs for the concurrent open-shop scheduling and the precedence constrained scheduling problems.

1.1 Concurrent Open-shop Scheduling Problem (COSSP)

In COSSP, we are given a set of m machines and a set of n jobs. Each job j has a release time r_j . The main feature of COSSP is that a job consists of m operations $O_{1j}, O_{2j}, \dots, O_{mj}$, one for each machine $i \in [m]$. Each operation O_{ij} needs p_{ij} units of processing on machine i . We allow operations to have *zero processing lengths*. Throughout the paper, we assume without loss of generality that all our input parameters are positive integers. A job is complete only when *all its operations are complete*. That is, if C_j denotes the completion time of j , then p_{ij} units of operation O_{ij} must be processed in the interval $[r_j, C_j]$ on machine i . In the concurrent open-shop scheduling model, multiple operations of the same job *can be processed simultaneously* across different machines.

² In some MapReduce applications, Reduce tasks can begin after the completion of a subset of Map tasks.

The problem has a long history due to its applications in manufacturing, automobile and airplane maintenance and repair [53], etc., and has been extensively studied both in operations research and approximation algorithm communities [3, 16, 23, 43, 52, 51, 39, 8]. As minimizing makespan in COSSP model is trivial, much of the research has focused on the objective of minimizing the total weighted completion times of jobs. The problem was first considered by Ahmadi and Bagchi [3], who showed the NP-hardness of the problem. Later, several groups of authors, Chen and Hall [16], Garg et al. [23], Leung et al. [39], and Mastrolilli et al. [43], designed 2-approximation algorithms for the problem. Under Unique Games Conjecture, Bansal and Khot showed that this approximation factor cannot be improved for the problem [8].

Garg, Kumar, and Pandit [23] studied the more difficult objective of minimizing the total flow-time of jobs, and showed that the problem cannot be approximated better than $\Omega(\log m)$, by giving a reduction from the set cover problem. However, they did not give any approximation algorithm to the problem, and left it as an open problem. To the best of our knowledge, the problem has remained open ever since. In this paper, we make progress on this problem.

Let P denote the ratio of the maximum processing length among all the operations to the minimum non-zero processing length of among all the operations; that is, $P := \frac{\max_{i,j} \{p_{ij}\}}{\min_{i,j} \{p_{ij} : p_{ij} \neq 0\}}$.

► **Theorem 1.** *For the objective of minimizing the sum of general delay cost functions of jobs in the concurrent open-shop scheduling model, there exists a polynomial time $O(\log(m \log P))$ approximation algorithm.*

We obtain the above result by generalizing the algorithm in Bansal and Pruhs [11]. Note that when $m = 1$, our result gives a $O(\log \log P)$ approximation algorithm to the problem, matching the best known polynomial time result in [11]. Recently, for the special case of total weighted flow-time, a constant factor approximation algorithm was obtained by Batra, Garg, and Kumar [13] when $m = 1$. However, running time of their algorithm is pseudo-polynomial. Since the approach in [13] is very different from the one in [11], our result does not generalize [13].

As we discussed earlier, the general delay cost functions capture several widely studied performance metrics. Thus, we get:

► **Corollary 2.** *There is a polynomial time $O(\log(m \log P))$ approximation algorithm in the concurrent open-shop scheduling model for the following objective functions: 1) Minimizing the sum of weighted flow-times of jobs; 2) Minimizing the weighted ℓ_k -norms of flow-times of jobs; the approximation factor becomes $O((\log(m \log P))^{\frac{1}{k}})$; 3) Minimizing the sum of weighted tardiness of jobs.*

We give the proof Theorem 1 in Section 2.

1.2 Precedence Constrained Scheduling Problem (PCSP)

More complex forms of job structures are captured by the precedence constrained scheduling problem (PCSP), another problem that has a long history dating back to the seminal work of Graham [24]. Here, we have a set of m identical machines and a set of n jobs; each job j has a processing length $p_j > 0$, a release time $r_j > 0$. Each job j must be scheduled on exactly one of the m machines. The important feature of the problem is that they are precedence constraints between jobs that capture the computational dependencies across jobs. The precedence constraints are given by a partial order “ \prec ”, where a constraint $j \prec j'$ requires that job j' can only start after job j is completed. Our goal is to schedule (preemptively) each job on exactly one machine to minimize $\sum_j g_j(C_j)$.

Precedence constrained scheduling on identical machines to minimize the makespan objective is perhaps the most well-known problem in scheduling theory. Already in 1966, Graham showed that list scheduling gives a 2-approximation algorithm to the problem. Since then several attempts have been made to improve the approximation factor [37, 21]. However, Svensson [49] showed that problem does not admit a $2 - \epsilon$ approximation under a strong version of the Unique Games Conjecture introduced by Bansal and Khot [8]. An unconditional hardness of $(4/3 - \epsilon)$ is also known due to Lenstra and Rinnooy Kan [38]. Recently, Levey and Rothvoss [40] showed that it is possible to overcome these lowerbounds for the case when m is fixed. An LP-hierarchy lift of the time-index LP with a slightly super poly-logarithmic number of rounds provides a $(1 + \epsilon)$ approximation to the problem.

Another problem that is extensively studied in the precedence constrained scheduling model is the problem of minimizing the total weighted completion times of jobs. Note that this problem strictly generalizes the makespan problem, hence all the lowerbounds also extend to this problem. The current best approximation factor of 3.387 is achieved by a very recent result of Li [41]. The work builds on a 4-approximation algorithm due to Munier, Queyranne and Schulz ([45], [47]).

In a recent work, Agrawal *et al* [2] initiated the study of minimizing the total flow-time objective in DAG (Directed Acyclic Graphs) *parallelizability* model. In this model, each job is a DAG, and a job completes only when all the nodes in the DAG are completed. For this problem, they showed greedy online algorithms that are constant competitive when given $(1 + \epsilon)$ -speed augmentation. The DAG parallelizability model is a special case of PCSP. However, as there are no dependencies between jobs, and individual nodes of the DAG do not contribute to the total flow-time unlike in PCSP, complexity of the problem is significantly different from PCSP. For example, when there is only one machine, the DAG structure of individual jobs does not change the cost of the optimal solution, and hence the problem reduces to the standard single machine scheduling problem. Therefore, scheduling jobs using Shortest Remaining Processing Time (SRPT), where processing length of a DAG is its total work across all its nodes, is an optimal algorithm. (Within a DAG, the nodes can be processed in any order respecting the precedence constraints.)

On the other hand, we show a somewhat surprising result for the PCSP problem. We show that the problem of minimizing the total flow-times of jobs does not admit any reasonable approximation factor even on a *single machine*. This is in sharp contrast to makespan and the sum of weighted completion times objective functions that admit $O(1)$ -approximation algorithms even on multiple machines. Our hardness proof is based on a recent breakthrough work of Manurangsi [42] on approximating the Densest-k-Subgraph (DKS) problem.

► **Theorem 3.** *In the precedence constrained scheduling model, for the objective of minimizing the total flow-times of jobs on a single machine, no polynomial time algorithm can achieve an approximation factor better than $n^{\frac{1}{(\log \log n)^c}}$, for some universal constant $c > 0$, assuming the exponential time hypothesis (ETH).*

To circumvent this hardness result, we study the problem in the speed augmentation model, which can also be thought of as a bi-criteria analysis. In the speed augmentation model, each machine is given some small extra speed compared to the optimal solution. The speed augmentation model was introduced in the seminal work of Kalyanasundaram and Pruhs [34] to analyze the effectiveness of various scheduling heuristics in the context of online algorithms. However, the model has also been used in the offline setting to overcome strong lowerbounds on the approximability of various scheduling problems; see for example results on non-preemptive flow-time scheduling that use $O(1)$ -speed augmentation to obtain $O(1)$ -approximation ratio [6, 30].

Our second main result is an $O(1)$ -approximation algorithm for the problem in the speed augmentation model. Previously, no results were known for the flow-time related objective functions for PCSP.

► **Theorem 4.** *For the objective of minimizing the sum of general delay cost of jobs in the precedence constrained scheduling model on identical machines, there exists a polynomial time $O(1)$ -speed $O(1)$ approximation algorithm. Furthermore, the speed augmentation required to achieve an approximation factor better than n^{1-c} , for any $c > 0$, has to be at least the best approximation factor of the makespan minimization problem. The lowerbound on speed augmentation extends to any machine environment, such as related and unrelated machine environments.*

We give the proofs of above theorems in the full version of the paper.

1.3 Applications of COSSP and PCSP in Data Center Scheduling

Besides being fundamental optimization problems, COSSP and PCSP models are very closely related to the scheduling problems that arise in the context of data centers. In particular, COSSP is a special case of the *Coflow Scheduling* problem introduced in a very influential work of Choudary and Stoica [17, 19]. On the other hand, PCSP generalizes the DAG scheduling problem, again a widely studied problem in systems literature. In fact, the DAG scheduling model has been adopted by Yarn, the resource manager of Hadoop [1]. See [26, 25] and references there-in for more details. Due to space constraints, we will defer the detailed discussion of connections among these problems to Appendix A.

1.4 Overview of the Algorithms and Techniques


Both our algorithms are based on rounding linear programming relaxations of the problems. However, individual techniques are quite different, and hence we discuss them separately.

1.4.1 Open-shop Scheduling Problem

Our algorithm for COSSP is based on the geometric view of scheduling developed by Bansal and Pruhs [11] for the problem of minimizing the sum of general delay costs of jobs on a single machine, which is a special case of our problem when $m = 1$. The key observation that leads to this geometric view is that minimizing general delay costs is equivalent to coming-up with a set of *feasible deadlines* for all jobs. Moreover, testing the feasibility of deadlines further boils down to ensuring that for *every* interval of time, the total volume of jobs that have (arrival time, deadline) windows within that interval is not large compared to the length. By a string of nice arguments, the authors show that this deadline scheduling problem can be viewed as a capacitated geometric set cover problem called R2C, which stands for capacitated rectangle covering problem in 2-dimensional space.³ Further, they argue that an α -approximation algorithm for R2C problem can be used to obtain an α -approximation algorithm for the scheduling problem.

In R2C, we are given a set \mathcal{P} of points in 2 dimensions, where each $p \in \mathcal{P}$ is specified by coordinates (x_p, y_p) . Associated with each point $p \in \mathcal{P}$ is a demand $d_p > 0$. We are also

³ “C” stands for the capacitated version in which rectangles have capacities and points have demands. Later, we shall use “M” for the multi-cover version, where rectangles are uncapacitated (or have capacity 1) and points have different demands. We use “U” for the uncapacitated version, where rectangles are uncapacitated and all points have demand 1.



■ **Figure 1** The figure shows an instance of PR2C problem where there are two sets of points. For this instance the only feasible solution is to pick both the rectangles.

given a set \mathcal{R} of rectangles, where $r \in \mathcal{R}$ has the form $(0, x_r) \times (y_r^1, y_r^2)$. Each rectangle r has a capacity $c(r) > 0$ and a cost $w(r) > 0$. The goal is to choose a minimum-cost set of rectangles, such that for every point $p \in \mathcal{P}$, the total capacity of selected rectangles covering p is at least d_p .

Our problem, which we call PR2C, can be seen as a parallel version of R2C. In PR2C, we have m instances of R2C problem with a common set of rectangles. Namely, the i th instance is defined by $(\mathcal{P}_i, \mathcal{R})$, where each $p \in \mathcal{P}_i$ is associated with a demand d_p , each $r \in \mathcal{R}$ is associated with a capacity $c(i, r)$ and cost $w(r)$. Notice that a rectangle $r \in \mathcal{R}$ has the same cost across the m instances, but has different capacities in different instances. The goal is to find a minimum cost set $X \subseteq \mathcal{R}$ of rectangles that is a valid solution for every R2C instance $(\mathcal{P}_i, \mathcal{R})$. Using the arguments similar to [11], we also show that if there is an α -approximation to PR2C problem, it gives an α -approximation for the COSSP problem.

Thus, much of our work is about designing a good approximation algorithm for the PR2C problem. Our algorithm for PR2C is a natural generalization of the algorithm for R2C of Bansal and Pruhs in [11]. In [11], Bansal and Pruhs formulated an LP relaxation for R2C based on knapsack cover (KC) inequalities. In the LP, $x_r \in [0, 1]$ indicates the fraction of the rectangle r that is selected. If the LP solution picks a rectangle to some constant fraction, then we can also select the rectangle in our solution without increasing the cost by too much. After selecting these rectangles, some points in \mathcal{P} are covered, and some other points p will still have residual demands d'_p . These not-yet-covered points are divided into two categories: *light and heavy points*. Roughly speaking, a point p is heavy if it is mostly covered by rectangles r with $c(r) \geq d'_p$ in the LP solution; a point p is light if it is mostly covered by rectangles r with $c(r) < d'_p$. Heavy points and light points are handled separately. The problem of covering heavy points will be reduced to the R3U problem, a geometric weighted set-cover problem where elements are points in 3D space and sets are axis-parallel cuboids. On the other hand, the problem of covering light points can be reduced to $O(\log_2 P)$ R2M instances. Each R2M instance is a geometric weighted set multi-cover instance in 2D-plane. By appealing to the geometry of the objects produced by the scheduling instance, [11] prove that *union complexity of objects* in R3U and R2M instances is small. In particular, R3U instance has *union complexity* $O(\log P)$ and each R2M instance has union complexity $O(1)$. Using the technique of quasi-uniform sampling introduced in [50, 15] for solving geometric weighted set cover instances with small union complexity, Bansal and Pruhs obtain $O(\log \log P)$ and $O(1)$ approximation ratios for the problems of covering heavy and light points, respectively.

In our problem, we have m parallel R2C instances with a common set of rectangles. As in [11], for each instance, we categorize the points into heavy and light points based on the LP solution. The problem of handling heavy points then can be reduced to m R3U

instances, with the m sets of cuboids identified. However, we cannot solve these m R3U instances separately, because such a solution cannot be mapped back to a valid schedule for COSSP. Therefore, we combine the m instances of R3C into a single instance of a 4 dimensional problem. So, in the combined instance, our geometric objects, which we call hyper-4cuboids, contain (at most) one 4-cuboid (a cuboid in 4 dimensions) from each of the m instances. The goal is to choose a minimum cost set of objects to cover all the points. On the other hand, for the light points, [11] reduced the problem to $O(\log P)$ R2M instances. Again, this approach is not viable for our case as we need to solve all the instances in parallel. By a simple trick, we first merge the $O(\log P)$ instances into one R2M instance. We then have m R2M instances with a new sets of rectangles identified, which we map into a single 3-dimensional geometric multi-set cover problem.

In both cases we show that the union complexity of the objects in our geometric problems increase at most by a factor of m compared to the objects in [11]. Thus, we have $O(m \log P)$ union complexity for the problem for heavy points and $O(m)$ union complexity for the problem for light points. Using the technique of [11], we obtain $O(\log(m \log P))$ and $O(\log m)$ approximation ratios for heavy and light points respectively, resulting in an $O(\log(m \log P))$ overall approximation ratio.

1.4.2 Precedence Constrained Scheduling

Our algorithm for the precedence constrained scheduling problem works in two steps. In the first step, we construct a *migratory* schedule, in which a job may be processed on multiple machines. For migratory schedules, we can assume that all jobs have unit size by replacing each job j of size p_j with a precedence chain of p_j unit length jobs. Solving a natural LP relaxation for the problem gives us a completion time vector $(C_j)_j$. Then we run the list-scheduling algorithm of Munier, Queyranne and Schulz [45], and Queyranne and Schulz [47], that works for the problem with weighted completion time objective. Specifically, for each job j in non-decreasing order of C_j values, we insert j to the earliest available slot after r_j without violating the precedence constraints.

To analyze the completion time of job j , we focus on the schedule constructed by our algorithm after the insertion of j . A simple lemma is that at any time slot after r_j , we are making progress towards scheduling j in the schedule: either all machines are busy in the time slot, or we are processing a set of jobs whose removal will decrease the “depth” of j in the precedence graph. This lemma was used in [45, 47] to give their 3-approximation for the problem of minimizing weighted completion time (for unit-size jobs), and recently by Li [41] to give an improved 2.415-approximation for the same problem. With 3-speed augmentation, this leads to a schedule that completes every job j by the time C_j . With additional speed augmentation, this leads to an $O(1)$ -approximation for the problem with general delay cost functions. In the second step, we convert the migratory schedule into a non-migratory one, using some known techniques (e.g. [35], [20], [33]). The conversion does not increase the completion times of jobs, but requires some extra $O(1)$ -speed augmentation.

2 Concurrent Open-shop Scheduling

In this section we consider the concurrent open-shop scheduling problem. Recall that in COSSP, we are given a set of m machines and a set of n jobs. Each job has a release time r_j . A job consists of m operations $\{O_{ij}\}_i$, one for each machine $i \in [m]$. Each operation O_{ij} needs p_{ij} units of processing on machine i . A job finishes only when all its operations are completed. The goal is to construct a preemptive schedule that minimizes the sum of costs incurred by jobs: $\sum_j g_j(C_j)$.

As we mentioned earlier, our algorithm for COSSP is based on the geometric view of scheduling developed in the work of Bansal and Pruhs [11]. Similar to [11], we first reduce our problem to a geometric covering problem that we call Parallel Capacitated Rectangle Covering Problem (PR2C). We argue that an α -approximation to PR2C will give an α approximation to our problem. Then, we design an $O(\log(m \log P))$ -factor approximation algorithm to PR2C.

2.1 Reduction to PR2C Problem

In PR2C, we have m instances of R2C problem with a common set of rectangles. Input to the problem consists of sets $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m$, where each \mathcal{P}_i is a set of points in 2-dimensional space. Each point $p \in \mathcal{P}_i$ is specified by its coordinates (x_p, y_p) and has a demand $d_p > 0$. The input also consists of a set \mathcal{R} of rectangles. Each rectangle $r \in \mathcal{R}$ has the form $(0, x_r) \times (y_r^1, y_r^2)$, and has a cost $w(r)$. Each rectangle r also has a capacity $c(i, r)$ which depends on the point set \mathcal{P}_i . Notice that a rectangle $r \in \mathcal{R}$ has the same cost across the m instances, but has different capacities in different instances. The goal is to find a minimum cost set $X \subseteq \mathcal{R}$ of rectangles that is a valid solution for every R2C instance $(\mathcal{P}_i, \mathcal{R})$. Recall that a set of rectangles is a valid solution to an instance of R2C, if for every point the total capacity of rectangles that cover the point is at least the demand of the point.

We can capture the PR2C problem using an integer program. Let x_r be a binary variable that indicates if the rectangle $r \in \mathcal{R}$ is picked in the solution or not. Then, the following integer program IP (1 - 3) captures the PR2C problem.

$$\text{Minimize } \sum_{r \in \mathcal{R}} w(r)x_r \quad \text{s.t.} \quad (1)$$

$$\forall i, \forall p \in \mathcal{P}_i : \sum_{r \ni p} c(i, r)x_r \geq d_p, \quad (2) \quad \forall r \in \mathcal{R} : x_r \in \{0, 1\}. \quad (3)$$

To see the connection between COSSP and PR2C, we need to understand the structure of feasible solutions to COSSP. Consider a feasible schedule S to an instance of COSSP. Suppose the completion time of a job j is C_j in S . This implies that on every machine i , the job j completes its operation O_{ij} in the interval $[r_j, C_j]$. If S is a feasible schedule and $\{C_j\}_j$ is the completion times of jobs, then processing jobs using the Earliest Deadline First (EDF) algorithm on each machine ensures that all jobs finish by their deadline C_j . Thus, one of the main observations behind our reduction is that minimizing the sum of costs incurred by jobs is equivalent to coming up with a deadline C_j for each job j . Thus a natural approach is to formulate COSSP as a deadline scheduling problem. However, this raises the question: Is there a way to test if a set of deadlines $\{C_j\}_j$ is a feasible solution to COSSP? An answer to the question is given by the characterization of when EDF will find a feasible schedule (one where all jobs meet their deadlines) on a single machine. To proceed, we need to set up some notation.

For an interval $I = [t_1, t_2]$, which consists of all the time slots between t_1 and t_2 , let $J(I)$ denote the set of jobs that have release times in I : $J(I) = \{j | r_j \in [t_1, t_2]\}$. For a set of jobs J' and a machine i , let $P(i, J')$ denote the total processing length of operations O_{ij} of the jobs in the set J' . Now, we introduce the notion of *excess* for an interval.

► **Definition 5.** For a given interval I and a machine i , the excess of I on machine i , denoted by $\xi(i, I)$, is defined as $\max\{0, P(i, J(I)) - |I|\}$.

Following lemma states the condition under which EDF can schedule all the jobs within their deadlines.

► **Lemma 6.** *Given a set of jobs with release times and deadlines, scheduling operations on each machine according to EDF is feasible if and only if, for every machine i and every interval $I = [t_1, t_2]$, the total processing lengths of operations corresponding to jobs in $J(I)$ that have deadlines greater than t_2 is at least $\xi(i, I)$.*

Clearly, if the total processing lengths of operations corresponding to jobs in $J(I)$ that have deadlines greater than t_2 is less than $\xi(i, I)$, then no algorithm can meet all the deadlines, as the length of operations scheduled in the interval I is greater than $|I|$. Sufficiency of the above lemma follows from a bipartite graph matching argument, and we refer the reader to [11] for more details.

Lemma 6 leads to an integer programming formulation for COSSP, and we shall use this IP to define the PR2C instance. Let $L = \max_i \{\sum_j p_{ij}\}$; note that every reasonable schedule finishes by L . For every job j , and every integer $q \in [-1, \lceil \log_2 g_j(L) \rceil]$, let $t_{j,q}$ be the largest integer $t \in (r_j, L]$ such that $g_j(t) \leq 2^q$ (if such t does not exist, then $t_{j,q} = r_j$). Let $t_{j,-2} = r_j$. For every j and $q \in [-1, \lceil \log_2 g_j(L) \rceil]$, we have a variable $x_{j,q}$ indicating whether $C_j \in (t_{j,q-1}, t_{j,q}]$. Therefore, for a job j , the total number of $x_{j,q}$ variables in our IP will be at most $O(\log g_j(L))$. For every j and every $t > r_j$, let $q_{j,t}$ be the integer q such that $t \in (t_{j,q-1}, t_{j,q}]$.

Our IP for COSSP is as follows:

$$\begin{aligned} & \text{minimize} && \sum_j \sum_q \lfloor 2^q \rfloor x_{j,q}, && \text{s.t.} \\ \forall i, \forall I = [t_1, t_2] : & \sum_{j \in J(I)} p_{ij} x_{j, q_{j,t_2}} \geq \xi(i, I), && \text{and} && \forall j, \forall q : x_{j,q} \in \{0, 1\}. \end{aligned}$$

In Appendix B we shall argue that the value of above IP is within $O(1)$ factor of the optimum cost of the scheduling problem.

The above integer program hints towards how we can interpret COSSP geometrically as PR2C. Indeed, it is equivalent to PR2C problem. For every machine i in COSSP, we create a set \mathcal{P}_i in PR2C. For every interval $I = [t_1, t_2]$ with $\xi(i, I) > 0$, we associate a point $p_I = (t_1, t_2)$ in 2-dimensional space. The demand d_{p_I} of a point $p_I \in \mathcal{P}_i$ is equal to $\xi(i, I)$, the excess of interval I on machine i . This completes the description of the point sets $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m$ in PR2C. Now we define the rectangle set \mathcal{R} . We shall create a rectangle for each job j and each q for which $x_{j,q}$ is defined. Notice that $x_{j,q}$ appears on the left side of Constraint (11) if $r_j \in [t_1, t_2]$ and $t_2 \in (t_{j,q-1}, t_{j,q}]$. Thus, we shall let the rectangle for the variable $x_{j,q}$ to be $(0, r_j] \times (t_{j,q-1}, t_{j,q}]$. This rectangle has cost 2^q ; for the instance \mathcal{P}_i , it has a capacity p_{ij} . Notice that the rectangle for $x_{j,q}$ covers a point (t_1, t_2) if and only if $t_1 \leq r_j \leq t_{j,q-1} < t_2 \leq t_{j,q}$; in other words, the job j is released in the interval $[t_1, t_2]$, and its completion time is greater than t_2 , which is exactly what we want. Thus, the IP (10 - 12) is equivalent to our PR2C problem. We now forget about the COSSP problem and focus exclusively on designing a good algorithm for the PR2C problem.

2.2 Algorithm For PR2C Problem

Our algorithm for PR2C is based on rounding a linear programming relaxation of the IP (1 - 3). As pointed out in [11], the linear programming relaxation of the IP (1 - 3) obtained by relaxing the variables $x_r \in [0, 1]$ has a large integrality gap even when there is only one set of points. We strengthen our LP by adding the so called KC inequalities, introduced first in the work of Carr *et al* ([14]). Towards that we need to define $c(i, S)$, which indicates the

total capacity of rectangles in $S \subseteq \mathcal{R}$ with respect to a point set \mathcal{P}_i : $c(i, S) = \sum_{r \in S} c(i, r)$. We are ready to write the LP.

$$\text{Minimize } \sum_{r \in \mathcal{R}} w(r)x_r \quad \text{s.t.} \quad (4)$$

$$\forall i, \forall p \in \mathcal{P}_i, \forall S \subseteq \mathcal{R}: \sum_{r \in \mathcal{R} \setminus S: p \in r} \min\{c(i, r), \max\{0, d_p - c(i, S)\}\} \cdot x_r \geq d_p - c(i, S) \quad (5)$$

$$\forall r \in \mathcal{R}: x_r \geq 0 \quad (6)$$

Let us focus on the KC inequalities Eq.(5) for a point set \mathcal{P}_i . Fix a point $p \in \mathcal{P}_i$ and a set of rectangles $S \subseteq \mathcal{R}$. Recall that p has a demand of d_p . Suppose, in an integral solution all the rectangles in S are chosen. Then, they contribute at most $c(i, S)$ towards satisfying the demand of the point p . The remaining rectangles still need to account for $d_p - c(i, S)$. Notice also that in the KC constraints we truncate the capacity of a rectangle r to $\min\{c(i, r), \max\{0, d_p - c(i, S)\}\}$. This ensures that the LP solution does not cheat by picking a rectangle with a large capacity to a tiny fraction. Clearly, the truncation has no effect on the integral solution.

There are exponentially many KC constraints in our LP. However, using standard arguments, we can solve the LP in polynomial time to get a $(1 + \epsilon)$ -approximate solution, for any $\epsilon > 0$, which suffices for our purposes to obtain a logarithmic approximation to COSSP; see [11, 14] for more details. The rest of this section is devoted to rounding the LP solution.

Weighted Geometric Set Multi-Cover Problem. The main tool used in our rounding algorithm is the result by [12] for the *weighted geometric set multi-cover problem*. Similar to the standard set cover problem, in this problem, we are given a set U of points and a set M consisting of subsets of U ; typically, the sets in M are defined by geometric objects. Further, each point $p \in U$ has a demand d_p , and a set $r \in M$ has a weight $w(r)$. The goal is to find the minimum weight set $S \subseteq M$ such that every point p is covered by at least d_p number of subsets in S . Note crucially that the sets in M do not have capacities. If the sets have capacities, then the problem becomes similar to PR2C. The interesting aspect of geometric set cover problems is that sets in M are defined by geometric objects, hence subsets in M have more structure than in the standard set cover problem. In particular, if the sets in M have low union complexity, then the problem admits a better than $O(\log |M|)$ approximation algorithm. Now we introduce the concept of union complexity of sets.

► **Definition 7.** Given a set S of n geometric objects, the union complexity of S is the number of edges in the arrangement of the boundary of S .

We will not get into the details of the definition, and we refer the readers interested in knowing more to [44] for an excellent introduction or to [50, 15, 12]. For our purposes, it suffices to know that the union complexity of 3-dimensional objects is the *total number of vertices, edges, and faces on the boundary*. It turns out that the geometric objects in our rounding algorithm are 4-dimensional objects. However, by appropriate projection, we reduce our problem of bounding the union complexity of 4-dimensional objects to that of 3-dimensional ones.

For geometric objects with low union complexity, building on the breakthrough works of Chan *et al* [15] and Varadarajan [50], Bansal and Pruhs [12] showed the following theorem.

► **Theorem 8.** *Suppose the union complexity of every $S \subseteq M$ of size k is at most $kh(k)$. Then, there is a polynomial time $O(\log h(n))$ approximation algorithm for the weighted geometric set multi-cover problem. Further, the approximation factor holds even against a feasible fractional solution.*

Rounding: Our rounding algorithm is based on reducing our problem to several instances of the geometric set cover problem, and then appealing to Theorem 8. Consider an optimal solution $\vec{x} = \{x_r\}_r$ to the LP (4 - 6). Let $\beta > 1/20$ be some constant. We first scale the solution by a factor of $1/\beta$, and let $\vec{x}' = \{x'_r\}_r$ be this scaled solution. Clearly, scaling increases the cost of LP solution at most by a factor of $1/\beta$. Our rounding consists of three steps.

In the first step, we pick all rectangles r for which $x'_r \geq 1$. Let \mathcal{S} denote this set. Let $\mathcal{P}_i(\mathcal{S})$ denote the set of points that are covered by rectangles in \mathcal{S} . We modify the point sets \mathcal{P}_i for $i = 1, 2, \dots, m$, by removing the points that are already covered by \mathcal{S} . For all $i \in [m]$, let $\mathcal{P}'_i = \mathcal{P}_i \setminus \mathcal{S}$. In the second step, for each \mathcal{P}'_i we classify the points in the set into two types, *heavy* or *light*, based on the LP solution. For heavy points, we create an instance of the geometric set cover problem, and then appeal to Theorem 8. The main technical difficulty here is to bound the union complexity of the geometric objects in our instance. For the light points, we create another separate instance of the geometric set multi-cover problem, and then apply the Theorem 8. Finally, we obtain our solution by taking the union of the rectangles picked in all three steps.

Fix a set \mathcal{P}'_i of uncovered points. For a point $p \in \mathcal{P}'_i$, let $\mathcal{S}(p)$ denote the set of rectangles in \mathcal{S} that contain p . That is, $\mathcal{S}(p) = \mathcal{S} \cap \{r : r \in \mathcal{R}, p \in r\}$. For a point p , define the residual demand of p as $d_p - c(i, \mathcal{S}_p)$. From the definition of set \mathcal{P}'_i , for every point $p \in \mathcal{P}'_i$, we note that $d_p - c(i, \mathcal{S}_p) > 0$. We now apply the KC inequalities on the set $\mathcal{S}(p)$ for each point p . From Eq.(5) we have, for all $p \in \mathcal{P}'_i$:

$$\sum_{r \in \mathcal{R} \setminus \mathcal{S}(p): p \in r} (\min\{c(i, r), d_p - c(i, \mathcal{S}(p))\}) \cdot x_r \geq d_p - c(i, \mathcal{S}(p))$$

which implies that our scaled solution \vec{x}' satisfies for all $p \in \mathcal{P}'_i$:

$$\sum_{r \in \mathcal{R} \setminus \mathcal{S}(p): p \in r} (\min\{c(i, r), d_p - c(i, \mathcal{S}(p))\}) \cdot x'_r \geq \frac{d_p - c(i, \mathcal{S}(p))}{\beta}.$$

Note that for all r , $x'_r \in [0, 1]$; otherwise, we would have picked those rectangles in \mathcal{S} . Next we round the residual demands of points and the capacities of rectangles as follows: For a point $p \in \mathcal{P}'_i$, let \tilde{d}_p denote the residual demand of p rounded up to the nearest power of 2. On the other hand, we round down the capacities $c(i, r)$ of rectangles $r \in \mathcal{R}$ to the nearest power of 2; let $\tilde{c}(i, r)$ denote this new rounded down capacities. Since \vec{x}' is scaled by a factor of $1/\beta$ we still have that for all $p \in \mathcal{P}'_i$,

$$\sum_{r \in \mathcal{R} \setminus \mathcal{S}(p): p \in r} (\min\{\tilde{c}(i, r), \tilde{d}_p\}) \cdot x'_r \geq \frac{\tilde{d}_p}{4\beta} \geq 3\tilde{d}_p. \tag{7}$$

To classify the points into heavy and light, we need the notion of class.

► **Definition 9.** Let $\tilde{c}_{\min}(i) = \min_{r \in \mathcal{R}} \{\tilde{c}(i, r)\}$. A rectangle $r \in \mathcal{R}$ is a class k rectangle with respect to a point set \mathcal{P}_i if $\tilde{c}(i, r) = 2^k \cdot \tilde{c}_{\min}(i)$. We say that a point $p \in \mathcal{P}'_i$ is a class k point if $\tilde{d}_p = 2^k \cdot \tilde{c}_{\min}(i)$.

Recall that the capacities of rectangles for different point sets can be different. Therefore, the class of a rectangle depends on the point set \mathcal{P}'_i . Now we categorize points in \mathcal{P}'_i into heavy and light as follows.

► **Definition 10.** A point $p \in \mathcal{P}'_i$ belonging to class k is heavy if its demand is satisfied by the rectangles of class at least k in the LP solution. Otherwise, we say that the point is light.

From the definition, for all heavy points $p \in \mathcal{P}'_i$ we have

$$\sum_{r \in \mathcal{R} \setminus \mathcal{S}: \tilde{c}(i,r) \geq \tilde{d}_p, p \in r} x'_r \geq 1, \quad (8)$$

and from Eq.(7) all light points $p \in \mathcal{P}'_i$ satisfy

$$\sum_{r \in \mathcal{R} \setminus \mathcal{S}: \tilde{c}(i,r) < \tilde{d}_p, p \in r} \tilde{c}(i,r) x'_r \geq (1/4\beta - 1) \tilde{d}_p = 2\tilde{d}_p. \quad (9)$$

Let $\mathcal{P}^1_i \subseteq \mathcal{P}'_i$ denote the set of heavy points and let $\mathcal{P}^2_i \subseteq \mathcal{P}'_i$ denote the set of light points. Let $\mathcal{R}' = \mathcal{R} \setminus \mathcal{S}$. We create two separate instances, a heavy instance and a light instance of the PR2C problem, corresponding to the set of heavy points and the set of light points. The capacities of rectangles in these instances will be their rounded down values: that is, the capacity of a rectangle $r \in \mathcal{R}'$ for the point sets \mathcal{P}^*_i will be $\tilde{c}(i,r)$. Similarly, the demand of a point $p \in \mathcal{P}^*_i$ will be its residual demand $d_p - c(*, \mathcal{S}(p))$. Notice that the LP solution \vec{x}' restricted to \mathcal{R}' is a feasible solution for both $\{\mathcal{P}^1_i\}_i$ and $\{\mathcal{P}^2_i\}_i$. We round the solution \vec{x}' for these two instances of PR2C problem separately, and take their union. The algorithm for the light instance will be deferred in the Appendix C; for the instance we get an $O(\log m)$ approximation ratio. We now show how to solve the heavy instance.

2.3 Heavy Instance and Hyper-4cuboid Covering Problem

We round the heavy instance $(\{\mathcal{P}^1_i\}_i, \mathcal{R}')$ by reducing it to a weighted geometric set cover problem called hyper-4cuboid covering problem (HCCP). HCCP is a generalization of the R3U problem considered in [11]. In HCCP, we are given a set U of points $p = (p_1, p_2, p_3, p_4)$ in 4-dimensional space. We are also given a set M of geometric objects. Each object $v \in M$ is a set of m disjoint 4-dimensional cuboids. Formally, $v = \{r_1, r_2, \dots, r_m\}$, where each r_i is a 4-dimensional cuboid (4-cuboid) of the form $[2i, 2i + 1] \times [0, x] \times [y_1, y_2] \times [0, z]$. We call v a hyper-4cuboid. Each hyper-4cuboid has a cost $w(v)$. The goal is to find the minimum cost set $S \subseteq M$ such that S covers all the points in U . We say that a hyper-4cuboid v covers a point p if $\exists r_i \in v, p \in r_i$.

Now we show that there is a reduction from the PR2C problem on the heavy instance to the HCCP problem. For a point $p \in \mathcal{P}^1_i$ with coordinates (x, y) , we create a point $p' \in U$ with coordinates $(2i + 1/2, x, y, \tilde{d}_p)$. Note that the last coordinate of p' is determined by \tilde{d}_p , which is the residual demand of point p . And the first coordinate of p' is determined by the index of the set \mathcal{P}^1_i . For every rectangle $r \in \mathcal{R}'$, we create a hyper-4cuboid $v_r \in M$, which contains exactly one 4cuboid for each point set \mathcal{P}^1_i . For a given index $i \in [m]$ and a rectangle $r = [0, x] \times [y_1, y_2]$, there is a 4cuboid r_i of the form $[2i, 2i + 1] \times [0, x] \times [y_1, y_2] \times [0, \tilde{c}(i, r)]$. Note that the last coordinate is determined by the capacity of rectangle r towards the point set \mathcal{P}^1_i . The cost of hyper-4cuboid v_r is same as the cost of rectangle r .

► **Lemma 11.** *Suppose there is a feasible solution of cost α to the PR2C problem on the heavy instance, where for each $i \in [m]$ and for each $p \in \mathcal{P}^1_i$, the demand of the point p is completely satisfied by the rectangles of class at least the class of p . Then there is a feasible solution of cost at most α for the corresponding instance of HCCP. Similarly, a solution of cost α to the HCCP problem gives a solution of the same cost for the heavy instance of PR2C.*

We defer the proof of the lemma to Appendix D. Now, observe crucially that there is a fractional solution of cost at most the cost of LP solution to the heavy instance satisfying the requirements of Lemma 11. This is true because the LP solution \vec{x}' satisfies the inequality Eq.(8). Therefore to apply Theorem 8, it remains to quantify the union complexity of hyper-4cuboids in our reduction.

► **Lemma 12.** *The union complexity of any q hyper-4 cuboids in M is at most $O(qm \log P)$, where $P = \frac{\max\{p_{ij}\}}{\min\{p_{ij}\}}$.*

Proof. Consider a set $S \subseteq M$ of q hyper-4 cuboids. For $i = 1, 2, \dots, m$, define \mathcal{L}_i as the set of all 4-cuboids at level i ; Formally, $\mathcal{L}_i \subseteq \{\bigcup_{v \in S} r_i(v)\}$, such that every 4-cuboid $r \in \mathcal{L}_i$ has the first dimension equal to $[2i, 2i + 1]$. In other words, \mathcal{L}_i is the set of all 4-cuboids which have the same first dimension $[2i, 2i + 1]$. Now, observe from our construction that for any pair i, i' , the objects in \mathcal{L}_i and $\mathcal{L}_{i'}$ do not intersect. This is because 4-cuboids at different levels are separated by a distance of 1 in the first dimension. Therefore, the union complexity of the subset S is equal to m times the maximum of union complexity of 4-cuboids in $\{\mathcal{L}_i\}_i$. However, 4-cuboids in the sets \mathcal{L}_i all have the same form; thus, it suffices to bound the union complexity for any \mathcal{L}_i .

Fix some i , and consider the 4-cuboids in the set \mathcal{L}_i . Notice carefully that all the 4-cuboids in \mathcal{L}_i share the same first dimension $[2i, 2i + 1]$. This implies that we can ignore the first dimension as it does not add any edges to the arrangement of objects in \mathcal{L}_i . Now consider the projection of 4-cuboids to the remaining 3 dimensions; These are cuboids of the form $[0, x] \times [y_1, y_2] \times [0, z]$. For these type of cuboids, [11] proves that the union complexity is at most $O(q\theta)$, where θ is the number of distinct values taken by z . In our reduction, the values of z correspond to the capacities of rectangles. Recall that the capacities of rectangles in PR2C are defined based on the processing lengths of operations. As we round down the capacities to the nearest power of 2, the number of distinct values taken by z is at most $\log P$, where $P = \frac{\max\{p_{ij}\}}{\min\{p_{ij}\}}$. Putting everything together, we complete the proof. ◀

From Lemma 11, 12 and Theorem 8, we get the following.

► **Lemma 13.** *There is a solution of cost $O(\log(m \log P))$ times the cost of LP solution for the HCCP problem, and hence for the PR2C on the heavy instance.*

Proof of Theorem 1. Our final solution for PR2C problem is obtained by taking the union of all the rectangles picked in our solutions for the heavy and the light instances, and the set \mathcal{S} . Recall that in \mathcal{S} , we pick all the rectangles for which $x_r > 1/\beta$, for $\beta = 12$. The cost of our solution is at most $(O(\log(m \log P)) + O(\log m) + 1/\beta)$ times the cost of LP (4 - 6), which implies an $O(\log(m \log P))$ approximation algorithm. This completes the proof. ◀

References

- 1 Apache hadoop.
- 2 Kunal Agrawal, Jing Li, Kefu Lu, and Benjamin Moseley. Scheduling parallel DAG jobs online to minimize average flow time. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 176–189, 2016. doi:10.1137/1.9781611974331.ch14.
- 3 Reza Ahmadi, Uttarayan Bagchi, and Thomas A Roemer. Coordinated scheduling of customer orders for quick response. *Naval Research Logistics (NRL)*, 52(6):493–512, 2005.
- 4 Saba Ahmadi, Samir Khuller, Manish Purohit, and Sheng Yang. On scheduling coflows. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 13–24. Springer, 2017.
- 5 S. Anand, Naveen Garg, and Amit Kumar. Resource augmentation for weighted flow-time explained by dual fitting. In *SODA*, pages 1228–1241, 2012. URL: <http://portal.acm.org/citation.cfm?id=2095213&CFID=63838676&CFTOKEN=79617016>.
- 6 Nikhil Bansal, Ho-Leung Chan, Rohit Khandekar, Kirk Pruhs, Cliff Stein, and Baruch Schieber. Non-preemptive min-sum scheduling with resource augmentation. In *Foundations*

- of *Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 614–624. IEEE, 2007.
- 7 Nikhil Bansal, Moses Charikar, Ravishankar Krishnaswamy, and Shi Li. Better algorithms and hardness for broadcast scheduling via a discrepancy approach. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 55–71. Society for Industrial and Applied Mathematics, 2014.
 - 8 Nikhil Bansal and Subhash Khot. Inapproximability of hypergraph vertex cover and applications to scheduling problems. *Automata, Languages and Programming*, pages 250–261, 2010.
 - 9 Nikhil Bansal, Ravishankar Krishnaswamy, and Viswanath Nagarajan. Better scalable algorithms for broadcast scheduling. In *ICALP (1)*, pages 324–335, 2010. doi:10.1007/978-3-642-14165-2_28.
 - 10 Nikhil Bansal and Janardhan Kulkarni. Minimizing flow-time on unrelated machines. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 851–860, 2015. doi:10.1145/2746539.2746601.
 - 11 Nikhil Bansal and Kirk Pruhs. The geometry of scheduling. In *IEEE Symposium on the Foundations of Computer Science*, pages 407–414, 2010.
 - 12 Nikhil Bansal and Kirk Pruhs. Weighted geometric set multi-cover via quasi-uniform sampling. In *European Symposium on Algorithms*, pages 145–156. Springer, 2012.
 - 13 Jatin Batra, Naveen Garg, and Amit Kumar. Constant factor approximation algorithm for weighted flow time on a single machine in pseudo-polynomial time. *CoRR*, abs/1802.07439, 2018.
 - 14 Robert D Carr, Lisa Fleischer, Vitus J Leung, and Cynthia A Phillips. Strengthening integrality gaps for capacitated network design and covering problems. In *SODA*, pages 106–115, 2000.
 - 15 Timothy M Chan, Elyot Grant, Jochen Könemann, and Malcolm Sharpe. Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 1576–1585. Society for Industrial and Applied Mathematics, 2012.
 - 16 Zhi-Long Chen and Nicholas G Hall. Supply chain scheduling: Conflict and cooperation in assembly systems. *Operations Research*, 55(6):1072–1089, 2007.
 - 17 Mosharaf Chowdhury and Ion Stoica. Coflow: A networking abstraction for cluster applications. In *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*, pages 31–36. ACM, 2012.
 - 18 Mosharaf Chowdhury and Ion Stoica. Efficient coflow scheduling without prior knowledge. In *ACM SIGCOMM Computer Communication Review*, volume 45, pages 393–406. ACM, 2015.
 - 19 Mosharaf Chowdhury, Yuan Zhong, and Ion Stoica. Efficient coflow scheduling with varys. In *ACM SIGCOMM Computer Communication Review*, volume 44, pages 443–454. ACM, 2014.
 - 20 Nikhil Devanur and Janardhan Kulkarni. A unified rounding algorithm for unrelated machines scheduling. In *Preprint*, <https://users.cs.duke.edu/kulkarni/papers/tardines.pdf>, 2017.
 - 21 Devdatta Gangal and Abhiram Ranade. Precedence constrained scheduling in $(2 - \frac{7}{3}p + 1)$ -optimal. *Journal of Computer and System Sciences*, 74(7):1139–1146, 2008.
 - 22 Naveen Garg and Amit Kumar. Minimizing average flow-time : Upper and lower bounds. In *FOCS*, pages 603–613, 2007. doi:10.1109/FOCS.2007.42.

- 23 Naveen Garg, Amit Kumar, and Vinayaka Pandit. Order scheduling models: hardness and algorithms. *FSTTCS 2007: Foundations of Software Technology and Theoretical Computer Science*, pages 96–107, 2007.
- 24 Ronald L Graham. Bounds for certain multiprocessing anomalies. *Bell Labs Technical Journal*, 45(9):1563–1581, 1966.
- 25 Robert Grandl, Mosharaf Chowdhury, Aditya Akella, and Ganesh Ananthanarayanan. Altruistic scheduling in multi-resource clusters. In *OSDI*, pages 65–80, 2016.
- 26 Robert Grandl, Srikanth Kandula, Sriram Rao, Aditya Akella, and Janardhan Kulkarni. GRAPHENE: packing and dependency-aware scheduling for data-parallel clusters. In *12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016, Savannah, GA, USA, November 2-4, 2016.*, pages 81–97, 2016. URL: https://www.usenix.org/conference/osdi16/technical-sessions/presentation/grandl_graphene.
- 27 Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. Competitive algorithms from competitive equilibria: non-clairvoyant scheduling under polyhedral constraints. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, pages 313–322, 2014. doi:10.1145/2591796.2591814.
- 28 Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. Competitive flow time algorithms for polyhedral scheduling. In *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 506–524, 2015. doi:10.1109/FOCS.2015.38.
- 29 Sungjin Im, Janardhan Kulkarni, Kamesh Munagala, and Kirk Pruhs. Selfishmigrate: A scalable algorithm for non-clairvoyantly scheduling heterogeneous processors. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 531–540, 2014. doi:10.1109/FOCS.2014.63.
- 30 Sungjin Im, Shi Li, Benjamin Moseley, and Eric Torng. A dynamic programming framework for non-preemptive scheduling problems on multiple machines. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1070–1086. Society for Industrial and Applied Mathematics, 2015.
- 31 Sungjin Im and Benjamin Moseley. Online scalable algorithm for minimizing ℓ_k -norms of weighted flow time on unrelated machines. In *SODA*, pages 95–108, 2011. URL: http://www.siam.org/proceedings/soda/2011/SODA11_008_ims.pdf.
- 32 Sungjin Im and Benjamin Moseley. An online scalable algorithm for average flow time in broadcast scheduling. *ACM Transactions on Algorithms*, 8(4):39, 2012. doi:10.1145/2344422.2344429.
- 33 Sungjin Im and Benjamin Moseley. General profit scheduling and the power of migration on heterogeneous machines. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2016, Asilomar State Beach/Pacific Grove, CA, USA, July 11-13, 2016*, pages 165–173, 2016.
- 34 Bala Kalyanasundaram and Kirk Pruhs. Speed is as powerful as clairvoyance. *Journal of the ACM*, 47(4):617–643, 2000.
- 35 Bala Kalyanasundaram and Kirk R Pruhs. Eliminating migration in multi-processor scheduling. *Journal of Algorithms*, 38(1):2–24, 2001.
- 36 Samir Khuller and Manish Purohit. Brief announcement: Improved approximation algorithms for scheduling co-flows. In *Proceedings of the 28th ACM Symposium on Parallelism in Algorithms and Architectures*, pages 239–240. ACM, 2016.
- 37 Shui Lam and Ravi Sethi. Worst case analysis of two scheduling algorithms. *SIAM Journal on Computing*, 6(3):518–536, 1977.
- 38 J. K. Lenstra and A. H. G. Rinnooy Kan. Complexity of scheduling under precedence constraints. *Oper. Res.*, 26(1):22–35, 1978. doi:10.1287/opre.26.1.22.

- 39 Joseph Y-T Leung, Haibing Li, and Michael Pinedo. Scheduling orders for multiple product types to minimize total weighted completion time. *Discrete Applied Mathematics*, 155(8):945–970, 2007.
- 40 Elaine Levey and Thomas Rothvoss. A $(1 + \epsilon)$ -approximation for makespan scheduling with precedence constraints using lp hierarchies. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, pages 168–177. ACM, 2016.
- 41 Shi Li. Scheduling to minimize total weighted completion time via time-indexed linear programming relaxations. In *Proceedings of the 2017 IEEE 58rd Annual Symposium on Foundations of Computer Science*, FOCS '17, 2017.
- 42 Pasin Manurangsi. Almost-polynomial ratio hardness of approximating densest k-subgraph. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 954–961, 2017. doi:10.1145/3055399.3055412.
- 43 Monaldo Mastrolilli, Maurice Queyranne, Andreas S Schulz, Ola Svensson, and Nelson A Uhan. Minimizing the sum of weighted completion times in a concurrent open shop. *Operations Research Letters*, 38(5):390–395, 2010.
- 44 Jiří Matoušek. *Lectures on discrete geometry*, volume 212. Springer Science & Business Media, 2002.
- 45 Alix Munier, Maurice Queyranne, and Andreas S. Schulz. *Approximation Bounds for a General Class of Precedence Constrained Parallel Machine Scheduling Problems*, pages 367–382. Springer Berlin Heidelberg, Berlin, Heidelberg, 1998. doi:10.1007/3-540-69346-7_28.
- 46 Zhen Qiu, Cliff Stein, and Yuan Zhong. Minimizing the total weighted completion time of coflows in datacenter networks. In *Proceedings of the 27th ACM symposium on Parallelism in Algorithms and Architectures*, pages 294–303. ACM, 2015.
- 47 Maurice Queyranne and Andreas S. Schulz. Approximation bounds for a general class of precedence constrained parallel machine scheduling problems. *SIAM J. Comput.*, 35(5):1241–1253, 2006. doi:10.1137/S0097539799358094.
- 48 Mehrnoosh Shafiee and Javad Ghaderi. An improved bound for minimizing the total weighted completion time of coflows in datacenters. *arXiv preprint arXiv:1704.08357*, 2017.
- 49 Ola Svensson. Conditional hardness of precedence constrained scheduling on identical machines. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 745–754. ACM, 2010.
- 50 Kasturi Varadarajan. Epsilon nets and union complexity. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*, pages 11–16. ACM, 2009.
- 51 Edouard Wagner and Chelliah Sriskandarajah. Openshops with jobs overlap. *European Journal of Operational Research*, 71(3):366–378, 1993.
- 52 Guoqing Wang and TC Edwin Cheng. Customer order scheduling to minimize total weighted completion time. *Omega*, 35(5):623–626, 2007.
- 53 Jaehwan Yang. *Scheduling with batch objectives*. PhD thesis, Ohio State University, 1998.

A Applications of COSSP and PCSP in Data Center Scheduling

Besides being fundamental optimization problems, COSSP and PCSP models are very closely related to the scheduling problems that arise in the context of data centers. We briefly describe the relevance of COSSP and PCSP to scheduling in data centers.

Coflow Scheduling. The COSSP problem is a special case of the *coflow scheduling abstraction* introduced in a very influential work of Choudary and Stoica [17, 19], in the context of scheduling data flows. They defined coflow as a collection of parallel flows with a

common performance goal. Their main motivation to introduce coflow was that in big-data systems, MapReduce systems for example, communication is structured and takes place across machines in successive computational stages. In most cases, communication stage of jobs cannot finish until all its flows have completed. For example, in MapReduce jobs, a reduce task cannot begin until *all* the map tasks finish. Although coflow abstraction was introduced to model scheduling flows in big data networks, it can also be applied to job scheduling in clusters; see [25] for example.

Therefore, we describe a slightly more general version of the coflow abstraction. Here, we are given a set of m machines (or m resources). Each job j consists of a set of operations $\{O_{ij}\}$ for $i = 1, 2, \dots$. Associated with each operation O_{ij} is a demand vector $D_{ij} = (d_{ij}(1), d_{ij}(2), \dots, d_{ij}(m))$, where each $d_{ij}(k) \in \{0, 1\}$. The demand vector indicates the subset of machines or resources the operation requires. An operation can be executed only when all the machines in the demand vector are allocated to it. Moreover, for each operation we are also given a processing length p_{ij} . The goal is to schedule all operations such that at any time instant the capacity constraints on machines are not violated: that is, each machine is allocated to exactly one operation. A job completes only when all its operations have finished.

The coflow problem studied by Choudary and Stoica [17, 19] corresponds to the case where the demand vectors of all jobs have exactly two 1s. That is, every operation needs two machines to execute. The machines typically correspond to input and output ports in a communication link. On the other hand, if the demand vector D_{ij} of each operation O_{ij} consists of exactly one non-zero entry, then the coflow scheduling is equivalent to COSSP.

In the past few years, coflow scheduling has attracted a lot of research both in theory and systems communities. In practice, several heuristics are known to perform well [17, 19, 18, 25] for the problem. The theoretical study of coflow scheduling was initiated by Qiu, Stein, and Zhong [46]. By exploiting its connections to COSSP, they designed a constant factor approximation algorithm to the objective of minimizing the total weighted completion times of jobs. Building on this work, better approximation algorithms were designed in [36, 4, 48]. Unfortunately, the techniques developed in these works do not seem to extend to the flow-time related objectives.

DAG Scheduling. Another problem that has attracted a lot of research in practice is the DAG scheduling problem. In this problem, we are given a set of machines (or clusters), and a set of jobs. Each job has a weight that captures the priority of a job. Each job j is represented by a directed acyclic graph (DAG). Each node of a DAG represents a task – a single unit of computation – that needs to be executed on a *single machine*. Each task has a release time and a processing length. An edge $j \rightarrow j'$ in the DAG indicates that the task j' depends on the task j , and j' can not begin its execution unless j finishes. The goal is to schedule jobs/DAGs on the machines so as to minimize the total weighted flow-time of jobs. Interestingly, this is one of the models of job scheduling that has been adopted by Yarn, the resource manager of Hadoop [1]. (Hadoop is a popular implementation of MapReduce framework.) Because of this, DAG scheduling has been a very active area of research in practice; see [26, 25] and references there-in for more details.

It is not hard to see that DAG scheduling problem is a special case of the precedence constrained scheduling problem. The union of the individual DAGs of jobs can be considered as one DAG, with appropriately defined release times and weights for each node. Furthermore, if jobs have no weight, and the release times of all tasks in the same DAG are equal, then the DAG scheduling model described above is same as the DAG parallelizability model

studied by Agrawal *et al* [2]. In fact, they design $(1 + \epsilon)$ -speed $O(1/\epsilon)$ -competitive *online algorithm* for the problem. On the other hand, our approximation algorithm for PCSP gives an approximation algorithm for the DAG scheduling problem even with weights and arbitrary release times for individual tasks.

B Integer Program for COSSP

Our IP for COSSP is as follows:

$$\text{Minimize } \sum_j \sum_q [2^q] x_{j,q} \quad \text{s.t.} \quad (10)$$

$$\forall i, \forall I = [t_1, t_2] : \sum_{j \in J(I)} p_{ij} x_{j,q_i,t_2} \geq \xi(i, I), \quad (11) \quad \forall j, \forall q : x_{j,q} \in \{0, 1\}. \quad (12)$$

We shall argue that the value of above IP is within a factor of $O(1)$ times the optimum cost of the scheduling problem. First, given a feasible schedule S for COSSP with completion time vector $(C_j)_j$, we construct a solution to the above integer program with a small cost. For each j and integer q such that $2^{q-1} \leq g_j(C_j)$, we let $x_{j,q} = 1$; set all other x variables to 0. Clearly, the cost of the solution to the IP is at most $O(1)$ times the cost of S . Moreover, Constraint (11) is satisfied: we have that $\sum_{j \in J(I): C_j > t_2} p_{i,j} \geq \xi(i, I)$. If $C_j > t_2$ for some $j \in J(I)$, then $x_{j,q_i,t_2} = 1$, as $2^{q_i,t_2-1} \leq g_j(t_2) \leq g_j(C_j)$.

On the other hand, if we are given an optimum solution to the above IP, we can convert it into a schedule for COSSP of cost at most the cost of the IP. For every j , let the completion time C_j of the job to be $t_{j,q}$, where q is the largest number such that $x_{j,q} = 1$; such a q must exist in order to satisfy the Constraint (11) for the interval $t_1 = t_2 = r_j$. Then the cost of our schedule is at most the cost of IP. On the other hand, Constraint (11) says that $\sum_{j \in J(I)} p_{ij} x_{j,q_i,t_2} \geq \xi(i, I)$, implying $\sum_{j \in J(I), C_j > t_2} p_{ij} \geq \xi(i, I)$, as $x_{j,q_i,t_2} = 1$ implies $C_j \geq t_2$.

► **Remark.** We will not discuss the representation of arbitrary delay functions $g_j(t)$ and how to compute $t_{j,q}$ here. We refer the readers to [11] for more details. Running time of all our algorithms will be polynomial in $n, m, \log L$ and $\max_j \{\log g_j(L)\}$.

C Covering Light Points and Geometric Multi Cover by Hyper-cuboids

Here, we design an $O(\log m)$ approximation algorithm to PR2C problem restricted to the light instance: $(\{\mathcal{P}_i^2\}_i, \mathcal{R}')$. There is one main technical difference between our algorithm for the light case from the algorithm of Bansal and Pruhs [11]. Bansal and Pruhs divide the light instance into $\log P$ different levels, where in each level they solve 2-dimensional geometric multi-cover problem, and show a $O(1)$ approximation to each level. This in turn leads to a constant factor approximation algorithm to the light instance. However, in our problem we cannot solve the instances separately. Therefore, using a simple trick, we map our problem into a single instance of 3-dimensional problem, and show a $(\log m)$ approximation to it. This is precisely the reason we loose the factor $O(\log m)$ for the light case, as the union complexity of our objects becomes m times the union complexity of the objects in [11].

Recall that for every point $p \in \mathcal{P}_i^2$, the LP solution satisfies the inequality Eq.(9). Again, our idea is to reduce the instance to a geometric uncapacitated set multi-cover problem, and then appeal to Theorem 8. The geometric objects produced by our reduction are *sets of cuboids*. Hence we abbreviate this problem as GMCC.

In GMCC, we are given a set U of points in 3-dimensional space, and a set M of geometric objects. Each geometric object $v \in M$ is a collection of m disjoint cuboids, which we call as hyper-cuboids. Each point $p \in U$ has a demand e_p , and each hyper-cuboid

$v \in M$ has a cost b_v . The individual cuboids constituting a hyper-cuboid v have the form $[x, x + 1] \times [0, y] \times [z_1, z_2]$. The objective is to find the minimum cost subset $S \subseteq M$ of hyper-cuboids such that for every point $p \in U$, there are at least e_p number of hyper-cuboids in S that contain the point p .

Now we give the reduction \mathcal{H} from an instance of PR2C to an instance of GMCC. Fix some $i \in [m]$. Let $T > 0$ be some large constant, and recall $P = \frac{\max\{p_{ij}\}}{\min\{p_{ij}\}}$. For every light point $p \in \mathcal{P}_i^2$ with coordinates (x_p, y_p) , we create $\log P$ different points, each of them shifted in the second dimension by a distance of T . The $\log P$ different points corresponding to a point $p \in \mathcal{P}_i^2$ have coordinates $(2i + 1/2, kT + x_p, y_p)$, for $k = 1, 2, \dots, \log P$. To complete the description of point set U , we need to define the demands of each point, which we will do after describing the cuboids in M .

For every rectangle $r \in \mathcal{R}'$, with dimensions $[0, r_j] \times [t_1, t_2]$, we create one hyper-cuboid $v_r \in M$. The cost of hyper-cuboid v_r is same as the cost of rectangle r . The hyper-cuboid v_r contains m different cuboids, one for each point set $\mathcal{P}_1^2, \mathcal{P}_2^2, \dots, \mathcal{P}_m^2$. Fix an index $i \in [m]$. The cuboid corresponding to the set \mathcal{P}_i^2 has dimensions $[2i, 2i + 1] \times [0, kT + r_j] \times [t_1, t_2]$, where k denotes the class of rectangle r with respect to \mathcal{P}_i^2 . Note that there is one-to-one correspondence between rectangles $r \in \mathcal{R}'$ to hyper-cuboids $v_r \in M$. Let $\mathcal{H}^{-1}(v_r) = r$; that is, the rectangle $r \in \mathcal{R}'$ that corresponds to the hyper-cuboid $v_r \in M$. We say that v is picked to an extent of x'_r in the LP solution \vec{x}' to mean the extent to which the rectangle r is picked in the LP solution.

Fix a point $p \in U$. We say that p is contained in the hyper-cuboid $v \in M$, if it is contained in any of the m cuboids that constitute v . Now we are ready to define the demand e_p of a point p as $e_p = \left\lfloor \sum_{v:p \in v} x_{\mathcal{H}^{-1}(v)} \right\rfloor$.

To understand the above definition, let us consider a point $p \in U$ with coordinates $(2i + 1/2, kT + x, y)$. The cuboids that can contain p should have form $[2i, 2i + 1] \times [0, kT + r_j] \times [t_1, t_2]$, where $kT + x < r_j$ and $y \in [t_1, t_2]$. These cuboids correspond to the class k rectangles in \mathcal{R} . Therefore, demand of a point $p \in U$ with coordinates $(2i + 1/2, kT + x, y)$ is exactly equal to the number of class k rectangles that cover the point $p \in \mathcal{P}_i^2$ in the LP solution.

This immediately tells us that the LP solution \vec{x}' is a feasible fractional solution to the instance of GMCC problem produced by our reduction \mathcal{H} . Now, we show the opposite direction.

► **Lemma 14.** *If there is an integral solution to the instance of GMCC produced by our reduction \mathcal{H} of cost at most α times the LP cost, then there is an integral solution of the same cost to the instance of PR2C on the light instance.*

Proof. Suppose $S \subset M$ is a solution to GMCC problem. We show that there is a corresponding solution to PR2C problem of exactly the same cost. Our solution is simple. For every hyper-cuboid $v \in S$, we pick the corresponding rectangle r in our solution S' to PR2C. From our reduction \mathcal{H} , the costs of hyper-cuboids is same as the costs of rectangles, which implies that solution costs of the two problems are the same. It remains to show that S' is a feasible solution for the light instance.

Fix a point set \mathcal{P}_i , and consider any point $p \in \mathcal{P}_i$. Let $p = (x, y)$. In our reduction, there are $\log P$ different points $p'_1, p'_2, \dots, p'_{\log P}$, where the point p'_k has coordinates $(2i + 1/2, kT + x, y)$. The demand of $d_{p'_k}$ is exactly equal to the total number of class k rectangles that cover the point p in the LP solution. Recall that in our reduction \mathcal{H} , each rectangle $r \in \mathcal{R}$ produces exactly one cuboid at the level i (that is with first dimension $[2i, 2i + 1]$) in the corresponding hyper-cuboid $v_r \in S'$. By abusing notation, let us call denote this cuboid by $h(r, i)$.

Let class of point p be k^* . The total capacity of rectangles that cover point $p \in \mathcal{P}_i$ in our solution S' is $\sum_{k \leq k^*} \sum_{r \in S': r \in \text{cls}(k)} 2^k \cdot \tilde{c}_{\min}(i)$.

From the one-to-one correspondence between the rectangles and the cuboids, the term $\sum_{k \leq k^*} \sum_{r \in S': r \in \text{cls}(k)} 2^k \cdot \tilde{c}_{\min}(i)$ is exactly equal to the number of cuboids that cover the point p'_k , which is at least the demand $d_{p'_k}$. Now, in our reduction, the demand of the point p'_k (which has coordinates $(2i + 1/2, kT + x, y)$) is exactly equal to the rounded down value of the total number of class k rectangles that cover the point p in the LP solution. Therefore,

$$\sum_{k \leq k^*} \sum_{r \in S': r \in \text{cls}(k)} 2^k \cdot \tilde{c}_{\min}(i) \geq \sum_{k \leq k^*} 2^k \cdot \tilde{c}_{\min}(i) \cdot \left\lfloor \sum_{r \in S': r \in \text{cls}(k)} x'_r \right\rfloor.$$

Simplifying the righthand side of the above equation,

$$\begin{aligned} & \sum_{k \leq k^*} 2^k \cdot \tilde{c}_{\min}(i) \cdot \left\lfloor \sum_{r \in S': r \in \text{cls}(k)} x'_r \right\rfloor \geq \sum_{k \leq k^*} 2^k \cdot \tilde{c}_{\min}(i) \cdot \left(\left(\sum_{r \in S': r \in \text{cls}(k)} x'_r \right) - 1 \right) \\ & \geq \sum_{k \leq k^*} 2^k \cdot \tilde{c}_{\min}(i) \cdot \left(\sum_{r \in S': r \in \text{cls}(k)} x'_r \right) - 2^{k^*} \tilde{c}_{\min}(i) \geq 2d'(p) - d'(p) = d'(p), \end{aligned}$$

where the last inequality follows from the Eq.(9), and the fact that p is a light point. Therefore, S' is a feasible solution to GMCC. \blacktriangleleft

Thus, to complete our algorithm for the light instance, it remains to design an $O(\log m)$ approximation algorithm to GMCC. For that we once again plan to rely on the Theorem 8, which requires to us to bound the union complexity of objects in M .

► **Lemma 15.** *The union complexity of any q hyper-cuboids in M is at most $O(mq)$.*

Proof. Consider any $S \subseteq M$ of q hyper-cuboids. From our definition, each hyper-cuboid v is a set of m cuboids. For $i = 1, 2, \dots, m$, define \mathcal{L}_i as the set of all cuboids with first dimension $[2i, 2i + 1]$. Clearly, for any two indices i, i' , the cuboids in \mathcal{L}_i and $\mathcal{L}_{i'}$ do not intersect, as they are separated by a distance of 1 in the first dimension. Moreover, the cuboids in different \mathcal{L}_i have same form except that they are shifted in the first dimension. Thus, the union complexity of q hyper-cuboids in S is at most m times the union complexity of q cuboids in \mathcal{L}_i , for any $i \in [m]$.

Fix some i and consider the set of q cuboids in \mathcal{L}_i . Since all of the q cuboids have exactly the same side in the first dimension, $[2i, 2i + 1]$, the union complexity of q cuboids in \mathcal{L}_i is equal to the union complexity of the projection of the cuboids to the last two dimensions. This projection of the cuboids produces a set of axis parallel rectangles. Furthermore, these q axis parallel rectangles partition into $\log P$ sets, $Z_1, Z_2, \dots, Z_{\log P}$, such that no two rectangles from different sets intersect. This is true as these rectangles have coordinates $[0, kT]$ in the second dimension. In each set Z_k , the rectangles are abutting the Y -axis (or the axis of second dimension). Bansal and Pruhs [11] showed that the union of complexity such axis parallel rectangles abutting Y -axis is at most $O(|Z_k|)$. Therefore, the union complexity all the q axis parallel rectangles is at most $\sum_k |Z_k| = O(q)$. Thus we conclude that the union complexity q hyper-cuboids is equal to mq , which completes the proof. \blacktriangleleft

Thus, from Theorem 8, we get a $O(\log m)$ approximation algorithm for the light instance of PR2C.

D Proof of Lemma 11

Proof of Lemma 11. Suppose $S \subset \mathcal{R}'$ is a feasible solution of cost α satisfying the condition stated in the lemma. Then, we claim that the set of hyper-4 cuboids S' corresponding to the rectangles $r \in S$ is a feasible solution to HCCP. Consider a point $p' \in U$ in the instance of HCCP produced by our reduction. Suppose p' has coordinates $(2i + 1/2, x, y, \tilde{d}_p)$. Then, there must be a point $p \in \mathcal{P}_i^1$ with coordinates (x, y) and the demand \tilde{d}_p . Suppose $r \in S$ covers this point p , and has dimensions $r = [0, x] \times [y_1, y_2]$. Then, we claim that $v_r \in S'$ covers the point p' . This is true, since v_r contains a 4-cuboid with coordinates $h = [2i, 2i + 1] \times [0, x] \times [y_1, y_2] \times [0, \tilde{c}(i, r)]$. It is easy to verify that $p' \in h$ as $\tilde{c}(i, r) > \tilde{d}_p$, which follows from the condition of the lemma.

The opposite direction also follows from the one-to-one correspondence between the rectangles in \mathcal{R} and the hyper-4cuboids. Suppose S' is a feasible solution to the instance of HCCP problem. Then, it is easy to verify that picking the rectangles $r \in \mathcal{R}$ corresponding to the hyper-4cuboids $v_r \in S'$ defines a feasible a solution to $(\{\mathcal{P}_i^1\}_i, \mathcal{R}')$. ◀