The geometry of scheduling

Citation for published version (APA):

Bansal, N., & Pruhs, K. R. (2014). The geometry of scheduling. *SIAM Journal on Computing*, 43(5), 1684-1698.
<https://doi.org/10.1137/130911317>

DOI:

[10.1137/130911317](https://doi.org/10.1137/130911317)

Document status and date:

Published: 01/01/2014

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

THE GEOMETRY OF SCHEDULING*

NIKHIL BANSAL[†] AND KIRK PRUHS[‡]

Abstract. We consider the following general scheduling problem. The input consists of n jobs, each with an arbitrary release time, size, and monotone function specifying the cost incurred when the job is completed at a particular time. The objective is to find a preemptive schedule of minimum aggregate cost. This problem formulation is general enough to include many natural scheduling objectives, such as total weighted flow time, total weighted tardiness, and sum of flow time squared. We give an $O(\log \log P)$ approximation for this problem, where P is the ratio of the maximum to minimum job size. We also give an $O(1)$ approximation in the special case of identical release times. These results are obtained by reducing the scheduling problem to a geometric capacitated set cover problem in two dimensions.

Key words. scheduling, approximation, weighted flow

AMS subject classifications. 11Y16, 68W25

DOI. 10.1137/130911317

1. Introduction.

We consider the following general offline scheduling problem. *General scheduling problem* (GSP). The input consists of a collection of n jobs, and for each job j there is a positive integer release time r_j , a positive integer size p_j and a nondecreasing cost (or weight) function $w_j(t) \geq 0$ specifying a nonnegative cost for each time $t > r_j$. (We will specify later how these weight functions are represented.) A feasible solution is a preemptive schedule, which assigns to each job j time slots $[t, t + 1]$ (not necessarily consecutive and satisfying $t \geq r_j$), during which j is run. A job is completed once it has been run for p_j units of time. If job j completes at time t , then a cost of $w_j(t)$ is incurred for that job. The objective is to minimize the total cost, $\sum_{j=1}^n w_j(c_j)$, where c_j is the completion time of job j .

GSP generalizes several natural scheduling problems:

Weighted flow time. If $w_j(t) = w_j \cdot (t - r_j)$, where w_j is some fixed weight associated with job j , then the objective is total weighted flow time.

Flow time squared. If $w_j(t) = (t - r_j)^2$, then the objective is the sum of the squares of flow times.

Weighted tardiness. If $w_j(t) = w_j \max(0, t - d_j)$ for some deadline $d_j \geq r_j$, then the objective is total weighted tardiness.

In general, this problem formulation can model any cost objective function that is the sum of arbitrary nondecreasing cost functions of flow times for individual jobs. Flow time, which is the duration of time $c_j - r_j$ that a job is in the system, is one of the most natural and commonly used quality of service measures for a job in the computer systems literature. Many commonly used and commonly studied scheduling objectives are based on combining the flow times of the individual jobs.

*Received by the editors February 26, 2013; accepted for publication (in revised form) June 3, 2014; published electronically September 25, 2014. A preliminary version of this paper appeared in the *Proceedings of the 2010 IEEE Symposium on Foundations of Computer Science*.

<http://www.siam.org/journals/sicomp/43-5/91131.html>

[†]Eindhoven Institute of Technology, Eindhoven, Netherlands (bansal@gmail.com). This author was supported in part by NWO grant 639.022.211.

[‡]Computer Science Department, University of Pittsburgh, Pittsburgh, PA 15260 (kirk@cs.pitt.edu). This author was supported in part by an IBM faculty award and NSF grants CNS-0325353, IIS-0534531, CCF-0830558, CCF-1115575, CNS-1253218, and CCF-1421508.

Despite much interest, large gaps remain in our understanding for even basic flow time based scheduling objectives. For example, for weighted flow time, the best known approximation ratios achievable by polynomial-time algorithms are polylogarithmic. For weighted tardiness, and flow time squared, no nontrivial approximation ratios were previously known to be achievable. On the other hand, for all three of these problems, even the possibility of a polynomial-time approximation scheme (PTAS) has not been ruled out. We discuss the related previous work further in section 1.3.

1.1. Our results. We show the following results.

THEOREM 1.1. *There is a randomized polynomial-time $O(\log \log P)$ approximation algorithm for GSP, where P is the ratio of the maximum to minimum job size.*

THEOREM 1.2. *In the special case when all the jobs have identical release times, there is a deterministic polynomial-time 16-approximation algorithm for GSP.*

Representation of the weight functions. Our algorithms only require that there is an efficient procedure to answer the following type of queries about the weight function: For any job j and integer $q > 0$, what is the earliest time when the cost of completing j is at least q , i.e., what is the smallest t such that $w_j(t) \geq q$.

Clearly, any reasonable representation of the weight function that we are aware of satisfies such a property. In fact, one can weaken this requirement even further. For example, losing a factor 2 in the approximation ratio, we can assume that $w_j(t)$ is always a nonnegative integer power of 2, and hence it suffices to be able to answer these queries within an error of $2 - \epsilon$. We also allow $w_j(t)$ to take the value $+\infty$, which can model a hard deadline for j . Assuming such queries are allowed, the running time of our algorithm is polynomial in n and $\log W$. Here W is the maximum value (excluding the value $+\infty$) attained by any weight function.

1.2. Techniques. The key idea behind these results is to view the scheduling problem geometrically and cast it as a capacitated geometric set cover problem. We then use algorithmic techniques developed for geometric set cover problems and for capacitated covering problems. In particular, we show that GSP can be reduced (with only a loss of factor 4 in the approximation ratio) to a problem we call R2C, defined below. (Here R stands for rectangle, 2 for two dimensions, and C for capacitated.) We then prove Theorem 1.3 that there is a loglog factor approximation for R2C.

Definition of the R2C problem. The input contains a collection $\mathcal{P}$ of points in two-dimensional space, where each point $p \in \mathcal{P}$ is specified by its coordinates (x_p, y_p) and has an associated positive integer demand d_p . Furthermore, the input contains a collection $\mathcal{R}$ of axis-parallel rectangles, each of them abutting the y -axis, i.e., each rectangle $r \in \mathcal{R}$ has the form $(0, x_r) \times (y_r^1, y_r^2)$. In addition, each rectangle $r \in \mathcal{R}$ has an associated positive integer capacity c_r and positive integer weight w_r . The goal is to find a minimum weight subset of rectangles, such that for each point $p \in \mathcal{P}$, the total capacity of rectangles covering p is at least d_p . Foreshadowing slightly, the rectangle capacities in R2C will correspond to job sizes in our reduction; thus we also use P to denote the maximum ratio of rectangle capacities. The following is an exact integer programming formulation of the problem:

$$\begin{aligned}
 (1.1) \quad & \min \sum w_r x_r \\
 \text{s.t.} \quad & \sum_{r \in \mathcal{R}: p \in r} c_r x_r \geq d_p \quad \forall p \in \mathcal{P}, \\
 & x_r \in \{0, 1\} \quad \forall r \in \mathcal{R}.
 \end{aligned}$$

THEOREM 1.3. *There is a polynomial-time $O(\log \log P)$ approximation algorithm for R2C, where P is the ratio of the maximum to minimum rectangle capacity.*

To prove Theorem 1.3, we combine recent results on weighted geometric cover problems where the sets have low union complexity, together with approaches for handling capacitated covering problems. Specifically, we crucially use the fact that the rectangles in the R2C problem touch the y -axis, and hence the union complexity of the boundary of any k rectangles (i.e., the number of vertices and edges on the boundary of the union) is $O(k)$.

This low union complexity is very useful. If all the capacities and demands in the R2C instance are 1, i.e., if we consider the standard (uncapacitated) set cover version of R2C, then an $O(1)$ approximation follows from the result of Chan et al. [14], which is a refinement of the breakthrough result of Varadarajan [29] on weighted geometric set cover problems with low union complexity. Thus the reason we lose $O(\log \log P)$ the factor in Theorem 1.3 is actually due to the arbitrary capacities and demands in R2C.

To handle arbitrary capacities and demands we use a framework formalized by Chakrabarty, Grant, and Konemann [13] based on knapsack cover (KC) inequalities [12]. Specifically, [13] shows that for any capacitated covering problem, it suffices to design good linear programming (LP) based approximation algorithms for two types of *uncapacitated* problems derived from the original capacitated problem. In particular, an α upper bound on the integrality gap for the so-called priority cover version of the original problem and a β upper bound on the integrality gap for the multicover version together imply a $O(\alpha + \beta)$ integrality gap for the original problem. In the multicover version of R2C each point p has an arbitrary integer demand d_p specifying the number of distinct rectangles that must cover it. In the priority cover version of R2C each rectangle and each point has a priority, and each point has to be covered by at least one rectangle of higher priority than itself. Being uncapacitated, these priority and multicover problems are often easier to deal with.

We show that the priority version of R2C can be viewed as another geometric (uncapacitated) set cover problem, with the property that the complexity of the union of any k sets is $O(k \log P)$. Since we need this bound on the union complexity of the priority cover problem, we will reprove the results of Chakrabarty, Grant, and Konemann [13] and not use them as a black-box. By the result of Chan et al. [14], the $O(k \log P)$ union complexity implies an LP based $\alpha = O(\log \log P)$ approximation for the priority version of R2C. An $O(1)$ approximation for the multicover version of R2C follows from the result of Bansal and Pruhs [6], which is an extension of the result of Chan et al. [14] for weighted geometric set cover to the multicover setting.

We note that our solution of the R2C problem (and hence the GSP problem) is completely based on LP rounding. Thus one contribution of this work is to provide the first strong LP formulation for flow time related problems. Prior to this work, all known LP formulations for, say, weighted flow time had integrality gaps polynomially large in n . This LP is somewhat obscured as we present our results using geometric terminology. However, we will give an explicit description of this scheduling LP in section 4.

Identical release times. When all jobs have identical release times, the R2C instance that arises from our reduction from the GSP has a much simpler form. All the points to be covered lie on a line, and the rectangles are one-dimensional intervals. This is called the generalized caching problem in the literature, and a polynomial-time

4-approximation algorithm is known [7]. Together with our reduction from GSP to R2C, which incurs another factor 4 loss, this implies Theorem 1.2.

Cheung and Shmoys [19] proposed a primal dual algorithm for GSP with identical release times. Mestre and Verschae [26] reformulated this algorithm as a local ratio algorithm and showed that it yielded a 4-approximation. Finally, Höhn, Mestre, and Wiese [21] gave an $e + \epsilon$ approximation with quasi-polynomial running time.

Organization. The paper is organized as follows. In section 2 we give the reduction from GSP to R2C. We also consider here the case of identical release times. In section 3 we give the strengthened LP formulation of R2C based on KC inequalities and for completeness show how rounding this LP solution reduces to rounding the priority cover and multicover version of the problem. In section 3.1 we describe the approximation algorithm for the priority cover version of R2C. In section 3.2 we describe the approximation algorithm for the multicover versions of R2C. Finally, in section 4, we describe the underlying LP for the scheduling problem explicitly and make some final remarks.

1.3. Related results.

Scheduling. There has been extensive work on various completion time and flow time related objectives, in both the offline and online settings, and we refer the reader to [28] for a relatively recent survey. We discuss here some work on special cases of GSP. The most well-studied of these cases is perhaps the weighted flow time. The best known polynomial-time algorithms have an approximation ratio of $O(\log^2 P)$ [17], $O(\log W)$, and $O(\log nP)$ [2]. A quasi-PTAS with running time $n^{O_\epsilon(\log P \log W)}$ is also known [16]. In the special case when the weights are the reciprocal of job sizes, the objective is known as average stretch or slow-down, and a PTAS [9, 16] is known.

For weighted tardiness, an $(n - 1)$ approximation algorithm is known for identical release times [18], but nothing seems to be known for arbitrary release dates. A PTAS is possible with the additional restriction that there are only a constant number of deadlines [22] or if jobs have unit size [23]. For total flow time squared, no approximation guarantees are known, unless one uses resource augmentation [5].

Geometric set cover. The goal in geometric set cover problems is to improve the general $O(\log n)$ approximation bound for set cover by using the geometric structure. Until recently most of the research on geometric set cover problems focused only on the *unweighted* case. A key idea is the connection between set covers and ϵ -nets [10], where an ϵ -net is a subcollection of sets that covers all the points that are contained in at least an ϵ fraction of the input sets. For typical geometric problems, the existence of ϵ -nets of size at most $(1/\epsilon)g(1/\epsilon)$ implies an $O(g(OPT))$ -approximate solution for unweighted set cover [10]. Thus, proving better bounds on sizes of ϵ -nets (an active research of research is discrete geometry) directly gives improved guarantees for unweighted set cover. In another direction, Clarkson and Varadarajan [20] related the guarantee for unweighted set cover to the union complexity of sets. In particular, if the sets have union complexity $O(kh(k))$, which roughly means that the number of points on the boundary of the union of any collection of k sets is $O(kh(k))$, then there is an $O(h(n))$ approximation [20].¹ This was subsequently improved to $O(\log(h(n)))$ [29] and in certain cases extended to the unweighted multicover case [15].

However, none of these earlier results applies to the weighted case. The problem is that these algorithms are nonuniform in that they sample some sets with much

¹The notion of union complexity used by [20] was slightly different from the one mentioned here.

higher probability than that specified by the LP relaxation. In a breakthrough result, Varadarajan [30] designed a new quasi-uniform sampling technique that samples each set with probability approximately proportional to specified by the LP relaxation. [30], then used this quasi-uniform sampling technique to obtain a $2^{O(\log^* n)} \log(h(n))$ approximation for weighted geometric set cover problems with union complexity $O(kh(k))$. He also gave an improved $O(\log h(n))$ approximation when $h(n)$ grows (possibly quite mildly) with n . Chan et al. [14] refined this algorithm to obtain an $O(\log h(n))$ approximation (for all ranges of $h(n)$) and also extended the result to a more general setting where union complexity is replaced by shallow cell complexity. This result was extended further by Bansal and Pruhs [6] to the multicover setting.

KC inequalities. KC inequalities were developed by Carr et al. [12] for the KC problem. In this problem, we are given a knapsack with capacity B and items with capacities c_i and weight w_i . The goal is to find the minimum weight collection of items that covers the knapsack (i.e., with total capacity at least B). Perhaps surprisingly, the standard LP relaxation for this problem turns out to be arbitrarily bad. Carr et al. [12] strengthened the standard LP for the KC problem by adding exponentially inequalities and showed that it has an integrality gap of 2. Moreover, this LP can be solved almost exactly in time polynomial in n and $1/\epsilon$ to give an integrality gap of $1 + \epsilon$. They also give an explicit polynomial size LP with integrality gap $2 + \epsilon$. These inequalities have been very useful for various capacitated covering problems [1, 11, 24, 3, 13]. Recently, Chakrabarty, Grant, and Konemann [13] gave an elegant framework for using these inequalities, which allows one to solve capacitated covering problems in a black-box manner without even knowing what the KC inequalities are.

2. The reduction from GSP to R2C. Our goal in this section is to show the following result.

THEOREM 2.1. *A polynomial-time α -approximation algorithm for R2C implies a polynomial-time 4α approximation algorithm for GSP.*

We now give the reduction from GSP to R2C. Before giving the formal specification of the reduction, we give the background motivation. The first idea is the following. As the contribution of a job to the objective only depends on its completion time, instead of specifying a complete schedule, we might as well just specify a deadline c_j for each job j by which it must be completed. We call an assignment of deadlines *feasible* if there is a schedule in which all deadlines are met. Recall that if an assignment of deadlines is feasible, then scheduling the jobs in the earliest deadline first (EDF) order actually gives a valid schedule. In Lemma 2.3 we give a duality-based characterization for feasibility and then explain how to interpret the condition geometrically.

We need to momentarily digress to discuss our conventions when discussing time. An interval $I = [t_1, t_2]$ consists of time slots $[t_1, t_1 + 1], \dots, [t_2 - 1, t_2]$ and has length $|I| = t_2 - t_1$. When we refer to time t , we mean the beginning of slot $[t, t + 1]$. Specifically, a job j is completed by time t if it is completed in slot $[t - 1, t]$ or earlier, and if a job arrives at time t , then it arrives at beginning of $[t, t + 1]$ and can be executed during the slot $[t, t + 1]$.

DEFINITION 2.2. *For an interval $I = [t_1, t_2]$, let $X(I) := \{j : r_j \in I\}$ denote the set of jobs that arrive during I , i.e., $r_j \in \{t_1, \dots, t_2\}$. We define $\xi(I)$, the excess of I , as $\max(p(X(I)) - |I|, 0)$, where $p(X(I)) := \sum_{j \in X(I)} p_j$ is the total size of jobs in $X(I)$.*

LEMMA 2.3. *An assignment of deadlines c_j to jobs is feasible if and only if for every interval $I = [t_1, t_2]$, the jobs in $X(I)$ that are assigned a deadline after I have a total size of at least $\xi(I)$. That is,*

$$\sum_{j \in X(I): c(j) > t_2} p_j \geq \xi(I) \quad \forall I = [t_1, t_2].$$

Proof. Consider any interval $I = [t_1, t_2]$. As at most $|I| = t_2 - t_1$ amount of work can be done on jobs in $X(I)$ during the interval I , the jobs in $X(I)$ that finish during I can have a total size of at most $|I|$. Thus, the jobs in $X(I)$ that finish after I must have a total size of at least $\max(p(X(I)) - |I|, 0) = \xi(I)$.

For the converse, we show that if the assignment of deadlines is infeasible, then an inequality is violated for some interval I . Consider an infeasible assignment of deadlines and let c_j be the earliest deadline missed when jobs are executed in the EDF order. Let $[t_0 - 1, t_0]$ be the latest time slot before c_j when EDF was either idle or was working on some job with deadline strictly greater than c_j . Consider the interval $I = [t_0, c_j]$. By definition of t_0 , EDF always works on jobs with deadline $\leq c_j$ during I . Moreover, all these jobs arrive during I (as there are no such jobs available at $t_0 - 1$) and hence in lie $X(I)$. As EDF is always working on these jobs during I and still misses the deadline at c_j , it must be that $p(X(I)) > |I|$. $\square$

Geometric view of Lemma 2.3. Let us associate a point $p_I = (t_1, t_2)$ in two-dimensional space with each time interval $I = [t_1, t_2]$. We will view p_I as a witness that enforces that jobs in $X(I)$ finishing after I have total size at least $p(X(I)) - |I|$. So, we associate a demand $d(p_I)$ of $\xi(I)$ with I . Next, for each job j and each possible completion time c_j for it, we associate a rectangle $R_j(c_j) = [0, r_j] \times [r_j, c_j - 1]$. We assign $R_j(c_j)$ a cost of $w_j(c_j)$ and capacity of p_j .

We illustrate this view in Figure 1. On the left is an interval $I = (s, t)$. The job j arrives in I , i.e., $r_j \in I$, and is assigned completion time c_j outside interval I . On the right is the point (s, t) corresponding to the interval I and the rectangle $R_j = [0, r_j] \times [r_j, c_j - 1]$ corresponding to job j and completion time c_j . Lemma 2.4 notes that (s, t) must be contained in R_j .

LEMMA 2.4. *The point p_I lies in the rectangle $R_j(c_j)$ if and only if the job j lies in $X(I)$ and the completion time c_j lies outside (after) I . Moreover, if $p_I \in R_j(c_j)$, then $R_j(c_j)$ contributes exactly p_j toward satisfying the demand of p_I .*

Proof. If a point (s, t) lies in the rectangle $R_j(c_j) = [0, r_j] \times [r_j, c_j - 1]$, this means that $s \in [0, r_j]$ and $t \in [r_j, c_j - 1]$. This is equivalent to the conditions $r_j \in [s, t]$ and

FIG. 1. The figure on the left shows an interval $I = (s, t)$. The job j arrives in I , i.e., $r_j \in I$, and is assigned completion time c_j outside interval I . The figure on the right shows the point (s, t) corresponding to the interval I and the rectangle $R_j = [0, r_j] \times [r_j, c_j - 1]$ corresponding to job j and completion time c_j . Note that (s, t) is contained in R_j .

$c_j > t$, which is precisely the condition that $j \in X(I)$, where $I = [s, t]$ and has deadline c_j after t . $\square$

By Lemmas 2.4 and 2.3, GSP can equivalently be stated as follows: Find the minimum cost collection of rectangles satisfying the following two conditions:

Uniqueness: For each job j , exactly one rectangle of the form $R_j(c_j)$ (for some c_j) is chosen.

Covering: For each interval I , the demand $d_p(I) = \xi(I)$ of the corresponding point p_I is satisfied.

Note that in the formulation above, the uniqueness condition is critical; otherwise one may cheat by picking multiple rectangles belonging to the same job (as the demand of a point might be covered by two or more rectangles of the same job, in which case the connection to the scheduling problem breaks down). To get to the R2C problem, we now show how to remove this uniqueness condition.

To get around this problem, we use another trick (shown in Figure 2). By losing a factor of 2 in the approximation ratio, for each job j it suffices to consider only those times c_j when the cost $w_j(c_j)$ first reaches an integer power of 2. Call these times $c_j^0, c_j^1, \dots$. Now the crucial observation is that $R_j(c') \subset R_j(c)$ for $c' < c$. So, for each possible completion time c_j^i , we define a new modified rectangle $R'_j(i)$ as $R'_j(i) = R_j(c_j^i) \setminus R_j(c_j^{i-1})$ and give it weight $w(R_j(c))$ and capacity p_j . Now the resulting rectangles $R'_j(i)$ are disjoint, and yet any original rectangle $R_j(c_j^i)$ can be expressed as $R_j(c_j^i) = \cup_{i' \leq i} R'_j(i')$. As the costs of $R'_j(i)$ are geometrically increasing, the cost of $\cup_{i' \leq i} R'_j(i')$ is at most twice that of $R_j(c_j^i)$. As they are disjoint, using these modified rectangles R'_j instead of R_j will allow us to remove the uniqueness condition completely.

We now define the reduction formally.

Definition of the reduction from GSP to R2C. From an arbitrary instance $\mathcal{J}$ of GSP, we create an instance $\mathcal{J}'$ of R2C. Consider $\mathcal{J}$ and some job j . For each integer $k \geq 0$, let I_j^k denote the interval of times (possibly empty) such that $w_j(t) \in [2^k, 2^{k+1})$. Note that for any fixed j , the intervals I_j^k are disjoint and partition the interval $[r_j, \infty)$. Moreover, the number of intervals for any job is $O(\log W)$.

At the loss of factor at most 2 in the objective, we can assume that the deadline c_j for job j is at the right end point of some interval I_j^k . Let $\mathcal{T}$ denote the set of end

FIG. 2. The figure on the left shows the various rectangles $R_j(i)$ corresponding to the single job j . Note that their y-spans are overlapping. On the right are the modified rectangles $R'_j(i)$ with nonoverlapping y-spans.

points of intervals I_j^k for all jobs j and indices k . For each job j in $\mathcal{J}$ and $k \geq 0$, create a rectangle $R_j^k = [0, r_j] \times I_j^k$ in $\mathcal{J}'$ with capacity p_j and weight 2^{k+1} . Note that the rectangles $R_j^0, R_j^1, \dots$ corresponding to a job j are pairwise disjoint. Note also that for simplicity, we have changed our notation for rectangles from the motivational discussion above.

For each time interval $I = [t_1, t_2]$, where $t_1, t_2 \in \mathcal{T}$, create a point p_I in $\mathcal{J}'$ with demand $d_p = \xi(I)$.

We now discuss briefly the complexity of the reduction. Let m denote the number of points in the R2C instance. Clearly, $m = O((n + |\mathcal{T}|)^2)$, as the only relevant times one needs to consider while defining the points and rectangles are the release times of jobs and times in $\mathcal{T}$. As there are $O(\log W)$ intervals for each job j , $|\mathcal{T}| = O(n \log W)$. We now show in Lemmas 2.5 and 2.6 that this reduction is approximation preserving (within constant factors). These lemmas then immediately imply Theorem 2.1.

LEMMA 2.5. *If there is a feasible solution S to a GSP instance $\mathcal{J}$ with objective value v , then there is a feasible solution S' to the corresponding R2C instance $\mathcal{J}'$ with objective value at most $4v$.*

Proof. Consider solution S , and for each j , let $k(j)$ be the index of the interval $I_j^{k(j)}$ during which j completes, so the cost incurred by j in S is at least $2^{k(j)}$. Consider the solution S' obtained by choosing for each job j , the rectangles $R_j^0, \dots, R_j^{k(j)}$. Clearly, the cost contribution of j is $\sum_{i=0}^{k(j)} 2^{i+1} \leq 4 \cdot 2^{k(j)}$ and hence at most four times that in S .

It remains to show that S' is feasible, i.e., for every point $p_I \in S'$, the total capacity of rectangles covering p_I is at least $d(p_I) = p(X(I)) - |I|$. Suppose $p = (t_1, t_2)$ corresponds to the time interval $I = [t_1, t_2]$. As S is a feasible schedule, by Lemma 2.3, the total size of jobs in $X(I)$ that finish after I is at least $p(X(I)) - |I|$. By Lemma 2.4, for each job in $X(I)$ that finishes after I , there is some rectangle in $R_j^0, \dots, R_j^{k(j)}$ that contributes p_j toward satisfying the demand of p_I . Thus the demand of every point p_I is satisfied. $\square$

LEMMA 2.6. *If there is a feasible solution S' to the R2C instance $\mathcal{J}'$ with value v , then there is a feasible solution S to GSP instance $\mathcal{J}$ with value at most v .*

Proof. Note that for each job j , at least one rectangle R_j^i must be picked in S' . This is because if we consider the point p_I corresponding to the interval $I = [r_j, r_j]$, it has demand equal to $p(I) - |I| = p(X(r_j)) - 0 = p(X(r_j))$, the total size of jobs arriving on r_j . Since it can only be covered by jobs j in $X(r_j)$, and for any such job at most one rectangle R_j^i can contribute p_j toward the demand of p_I , we can conclude that one rectangle from each job in $X(r_j)$ must be used.

We construct the solution S as follows. For each job j , let $h(j)$ denote the largest index rectangle $R_j^{h(j)}$ that is chosen in S' . Set the deadline c_j for j as the right end point of $I_j^{h(j)}$. The cost of j in the schedule S is at most $2^{h(j)}$ and hence at most the cost of the rectangle $R_{h(j)}^j$ in $\mathcal{J}'$.

The schedule is feasible for the following reason. If point $p = p_I$ is covered by some rectangle R_j^i corresponding to job j , then by Lemma 2.4, $j \in X(I)$ and the deadline c_j for j is after I . As the demand $d(p_I)$ of each point p_I is satisfied in S' , the total size of jobs in $X(I)$ that are assigned deadline after I is at least $d(p_I)$ and hence by Lemma 2.3 the schedule is feasible. $\square$

2.1. Identical release times. In this subsection, we prove Theorem 1.2, that there is a deterministic polynomial-time 16-approximation algorithm for GSP in the special case of identical release times, that is, $r_j = 0$ for all j . Here, the reduction becomes much simpler. In particular, for each rectangle R_j^k , the first dimension $[0, r_j]$ becomes irrelevant and we obtain the following problem. For each job j and $k \geq 0$, there is an interval C_k^j corresponding to completion times with cost in the range $[2^k, 2^{k+1})$. This interval has capacity p_j and weight 2^k . All relevant points p_I corresponds to intervals of the form $[0, t]$ for $t \in \mathcal{T}$ and have demand $D - t$, where D is the total size of all the jobs. For each such interval $I = [0, t]$ (instead of a two-dimensional representation), we introduce a point $p_I = t$ with demand $d_I = D - t$. The goal is to find a minimum weight subcollection of intervals C_k^j that covers the demand.

This problem is a special case of the previously studied generalized caching problem, defined as follows. The input consists of arbitrary demands d_t at various time steps $t = 1, \dots, n$. In addition there is a collection of time intervals $\mathcal{I}$, where each interval $I \in \mathcal{I}$ has weight w_I , size c_I , and span $[s_I, t_I]$ with $s_I, t_I \in \{1, \dots, n\}$. The goal is to find a minimum weight subset of intervals that covers the demand. That is, find the minimum weight subset of intervals $S \subseteq \mathcal{I}$ such that

$$\sum_{I \in S: t \in [s_I, t_I]} c_I \geq d_t \quad \forall t \in \{1, \dots, n\}.$$

A deterministic polynomial-time 4-approximation algorithm, based on the local ratio technique, for generalized caching is given by Bar-Noy et al. [7]. This algorithm can equivalently be viewed as a primal dual algorithm applied to a linear program with KC inequalities [8]. Combining this result with Theorem 2.1 implies a polynomial-time 16-approximation algorithm for GSP in the case of identical release times.

3. Solving the R2C problem. In this section we focus on solving the R2C problem. We consider the natural LP formulation for R2C strengthened by KC inequalities and then show how to round it suitably. Using standard techniques, we show that the problem of rounding this LP reduces to finding a good rounding for two types of uncapacitated covering instances: a so-called priority cover instance, and several multicover instances. While we could directly use the framework of [13] here, we prefer to describe the reduction explicitly both for completeness and also since we will crucially need the fact that there are only $O(\log P)$ distinct priorities in the priority cover version, which is only implicit in [13]. Then in subsection 3.1, we give a rounding algorithm that produces a cover of the resulting priority cover instance with cost $O(\log \log P)$ times the LP cost, and in subsection 3.2 we give an algorithm that produces a cover of the resulting multicover instances with cost $O(1)$ times the LP cost. An $O(\log \log P)$ approximation for R2C is then obtained by picking a rectangle r in the final solution if it is included in the covers produced in any of the subinstances.

LP formulation. Consider the natural LP relaxation of the integer program for R2C given in line (1.1), obtained by relaxing $x_r \in \{0, 1\}$ to $x_r \in [0, 1]$. This LP has an arbitrarily large integrality gap, even when $\mathcal{P}$ consists of just a single point. (Observe that the R2C problem on a single point instance is precisely the KC problem [12].) Thus, we strengthen this LP by adding KC inequalities introduced in [12]. This gives the following linear program:

$$\begin{aligned}
& \min \sum_{r \in \mathcal{R}} w_r x_r \\
\text{s.t.} \quad & \sum_{r \in \mathcal{R} \setminus S: p \in r} \min \{c_r, \max(0, d_p - c(S))\} x_r \geq d_p - c(S) \quad \forall p \in \mathcal{P}, S \subseteq \mathcal{R}, \\
& x_r \in [0, 1] \quad \forall r \in \mathcal{R}.
\end{aligned}$$

Here $c(S)$ denotes the total capacity of rectangles in S . The constraints are valid for the following reason: For any p , and for any subset of rectangles S , even if all the rectangles in S are chosen, at least a demand of $d_p - c(S)$ must be covered by the remaining rectangles. Moreover, truncating the capacity of a rectangle from c_r to $d_p - c(S)$ (in the constraint for point p) does not affect the feasibility of an integral solution. Even though there are exponentially many constraints per point, for any $\epsilon > 0$ [12] give an efficient combinatorial algorithm to find a $(1 + \epsilon)$ -approximate solution. We note that the $(1 + \epsilon)$ factor loss is only in the cost, and in particular, all the constraints are satisfied exactly.

Residual solution. Let x be some $(1 + \epsilon)$ -approximate feasible solution to the LP above, and let OPT denote the objective value. We simplify x as follows. Let β be a small constant, and $\beta = 1/12$ suffices. Let S denote the set of rectangles for which $x_r \geq \beta$. We pick all the rectangles in S , i.e., set $x_r = 1$. Clearly, this cost of this set is at most $1/\beta$ times the LP solution.

For each point p , let $S_p = S \cap \{r : r \in \mathcal{R}, p \in r\}$ denote the set of rectangles in S that contain p . Let us consider the residual instance with rectangles restricted to $\mathcal{R} \setminus S$ and the demand of point p is $d_p - c(S_p)$. If $d_p - c(S_p) \leq 0$, then p is already covered by S and we discard it. Since the solution x satisfies all the KC inequalities and hence in particular for S_p , we have that

$$\sum_{r \in \mathcal{R} \setminus S_p: p \in r} \min \{c_r, d_p - c(S_p)\} x_r \geq d_p - c(S_p) \quad \forall p.$$

Henceforth, we will only use the fact that x satisfies these inequalities for the particular sets S_p .

Scale the solution x restricted to $\mathcal{R} \setminus S$ by $1/\beta$ times. Call this solution x' . Note that since $x_r \leq \beta$, it still holds that $x'_r \in [0, 1]$. Clearly, x' satisfies

$$\sum_{r \in \mathcal{R} \setminus S_p: p \in r} \min \{c_r, d_p - c(S_p)\} x'_r \geq \frac{d_p - c(S_p)}{\beta} \quad \forall p.$$

Let us define a new demand d'_p for p as $d_p - c(S_p)$ rounded up to the nearest integer power of 2. Similarly, define a new capacity c'_r of each rectangle r to be c_r rounded down to the nearest integer power of 2. The solution x' still satisfies

$$(3.1) \quad \sum_{r \in \mathcal{R} \setminus S_p: p \in r} \min \{c'_r, d'_p\} x'_r \geq \frac{d'_p}{4\beta} = 3d'_p \quad \forall p.$$

Decomposition into heavy and light points. We call r a class i rectangle if $c'_r = 2^i c'_{\min}$. Similarly, we call p a class i point if $d'_p = 2^i d'_{\min}$. (We remark that points could have negative class indices.) Note that the number of classes for rectangles is at most $O(\log P)$. We call a point p *heavy* if it is covered by rectangles with class at least as high as that of p in the LP solution. That is,

$$\sum_{r \in \mathcal{R}': c'_r \geq d'_p} \min(c'_r, d'_p) x'_r \geq d'_p,$$

or equivalently if

$$(3.2) \quad \sum_{r \in \mathcal{R}': c'_r \geq d'_p} x'_r \geq 1.$$

Otherwise we say that a point is *light*. By (3.1), we have that a light point satisfies

$$(3.3) \quad \sum_{r \in \mathcal{R}': c'_r < d'_p} c'_r x'_r \geq \left(\frac{1}{4\beta} - 1 \right) d'_p = 2d'_p.$$

We make the trivial note (for later use) that a point p with $d'_p \leq c'_{\min}$ is always heavy.

3.1. Covering heavy points. Covering the heavy points by larger class rectangles reduces to the following problem that we call R3U. Here R stands for rectangle, 3 for three dimensions, and U for uncapacitated.

Definition of problem R3U. The input consists of a collection $\mathcal{P}$ consisting of points $p = (p_x, p_y, p_z)$ in three-dimensional space and a collection of cuboids of the form $r = [0, x_r] \times [y_r^1, y_r^2] \times [0, z_r]$ and weight w_r . Note that these cuboids touch the $x - y$ plane and the $y - z$ plane. The goal is to find a minimum weight collection of cuboids that cover all points in $\mathcal{P}$.

Note that there are no demands and capacities in R3U.

LEMMA 3.1. *The problem of covering heavy points with rectangles of higher class is a special case of R3U.*

Proof. The reduction takes as input the instance $\mathcal{T}'$ for heavy points and the LP solution x' and creates an instance $\mathcal{A}$ of R3U. For each heavy point $p = (x, y) \in \mathcal{T}'$ with demand d'_p , there is a point (x, y, d'_p) in $\mathcal{A}$. For each rectangle $r = [0, x] \times [y_1, y_2]$ in $\mathcal{T}'$ with capacity c'_r , we define a cuboid $C_r = [0, x] \times [y_1, y_2] \times [0, c'_r]$ of weight w_r . Clearly, a point p in $\mathcal{A}$ can be covered by a cuboid C_r if and only if the corresponding point p in $\mathcal{T}'$ is covered by the corresponding rectangle r and the class of r is not smaller than that of p . $\square$

As the LP solution x' satisfies inequality (3.2) for heavy points, it gives a feasible LP solution to the R3U instance $\mathcal{A}$. We also note that the cuboids in $\mathcal{A}$ have at most $O(\log P)$ different heights, corresponding to the $O(\log P)$ distinct possible values for c'_r .

DEFINITION 3.2 (union complexity). *Given a collection X of n geometric objects, the union complexity of X is number of edges in the arrangement of the boundary of X . For three-dimensional objects, this is the total number of vertices, edges, and faces on the boundary of X .*

We would like to bound the union complexity of cuboids in $\mathcal{A}$. We begin with bounding the union complexity of the two-dimensional rectangles in R2C.

LEMMA 3.3. *For any collection of k rectangles of the type $[0, r] \times [s, t]$, the union complexity is $O(k)$.*

Proof. Each rectangle of the form $[0, r] \times [s, t]$ has a side touching the y -axis. Let us consider the union of an arbitrary subcollection of k such rectangles. The boundary of the union has horizontal faces and vertices. If two faces overlap each other, we break ties in some consistent way. Now the horizontal faces of any rectangle can contribute at most two faces to the boundary. In particular, if we consider some rectangle $R = [0, r] \times [s, t]$ and its bottom face (i.e., the segment joining $(0, s)$ to (r, s)), then only a subsegment joining (r', s) to (r, s) can appear in the union for some $r' < r$.

To bound the number of vertical faces on the boundary of the union, imagine looking at these vertical faces from $x = \infty$. For any two rectangles a and b , we note that the pattern $abab$ or $baba$ cannot appear. Thus the vertical faces form a Davenport–Schinzel sequence of order 2, which has size at most $2k - 1$ (see, for example, [25, chapter 7]).

The total faces is at most $O(k)$ and as the number of vertices is $O(1)$ times the number of faces, the result follows. $\square$

LEMMA 3.4. *The union complexity of any k cuboids in $\mathcal{R}$ is $O(k \log P)$.*

Proof. This directly follows from Lemma 3.3 and noting that the number of distinct heights is $O(\log P)$. In particular, since the heights are powers of 2, consider the slice of the cuboids between $z = 2^i$ and $z = 2^{i+1}$. This slice corresponds to a union of rectangles of the form $[0, r] \times [s, t]$. As there are $O(\log P)$ slides, the result follows. $\square$

We now use the following result of [6], which is an extension of the results of [29] and [14].

THEOREM 3.5 (see [6]). *Let I be an instance of an (uncapacitated) geometric set multicover problem on n points, such that the union complexity of every k sets is at most $kh(k)$ for all k . Then there is a polynomial-time $O(\log h(n))$ approximation for the problem, and moreover this approximation guarantee holds with respect to the optimum value of the standard LP relaxation for the problem.*

As x' is a feasible fractional solution to $\mathcal{A}$, and $h(k) = O(\log P)$ for $\mathcal{A}$ by Lemma 3.4, applying Theorem 3.5 allows us to conclude in Lemma 3.6 that we can obtain an $O(\log \log P)$ approximation for heavy points.

LEMMA 3.6. *The algorithm finds a cover for heavy points of value at most $O(\log \log P)$ times the cost of x' and hence at most $O(\log \log P)$ times the cost of the optimum R2C solution.*

3.2. Covering light points. We relate the problem of covering light points to the following R2M problem. Here R stands for rectangle, 2 for two dimensions, and M for multicover.

Definition of the R2M problem. The input consists of a collection $\mathcal{P}$ of two-dimensional points of the form $p = (x_p, y_p)$, each with an integer demand d_p , and a collection of axis-parallel rectangles r of the form $[0, x_r] \times [y_r^1, y_r^2]$, with cost w_r . The goal is to find a minimum cost subset $S \subset \mathcal{R}$ of rectangles such that each point $p \in \mathcal{P}$ is covered by at least d_p rectangles in S .

Reducing covering light points to instances of R2M. The reduction takes as input the instance $\mathcal{I}'$ for R2C, restricted to light points, and the LP solution x' satisfying inequality (3.3). As output, it creates $\log P$ instances of R2M, one instance B_ℓ for each capacity class ℓ . The reduction will ensure that an LP based α -approximation for R2M implies a cover for light points in $\mathcal{I}'$ with cost $O(\alpha)$ times the cost of the LP solution x' .

The instance B_ℓ . For each $\ell = 0, \dots, \log P$, the rectangles in B_ℓ are the class ℓ rectangles in $\mathcal{I}'$, i.e., those with capacity exactly $2^\ell c'_{\min}$. The points in B_ℓ are all the light points in $\mathcal{I}'$.

The cost of each rectangle in B_ℓ is the same as in $\mathcal{I}'$. The demand of a point p in the instance B_ℓ is defined as

$$d_p^\ell = \left\lceil \sum_{r: p \in r, r \in B_\ell} x'_r \right\rceil,$$

that is, the (rounded down) amount of class ℓ rectangles that cover p in the LP solution x' . In the instance B_ℓ , the goal is to cover each point $p \in B_\ell$ by at least d_p^ℓ distinct rectangles.

LEMMA 3.7. *The choice of demands d_p^ℓ ensures that the solution x' restricted to those rectangles in B_ℓ (i.e., the ones with capacity $2^\ell c'_{\min}$) is a valid solution to B_ℓ .*

Proof. The amount of p covered by x' restricted to those rectangles in B_ℓ is $\sum_{r:p \in r, r \in B_\ell} x(r)$, which is at least d_p^ℓ . $\square$

The following lemmas show that a good LP based approximation for R2M can be used separately for each instance B_ℓ to give a good solution to cover the light points.

LEMMA 3.8. *Suppose there is an LP based α -approximation for R2M. Let S_ℓ be any feasible solution to B_ℓ obtained by applying the LP based algorithm to the solution x' restricted to rectangles B_ℓ . Consider the union S of the rectangles picked in the solutions S_ℓ to the instances B_ℓ . Then S satisfies the demand of all the light points in $\mathcal{I}'$, and the cost of S is at most α times the cost of x' .*

Proof. As the cost of S_ℓ is at most α times the cost of x' restricted to rectangles in B_ℓ , the cost of S is at most α times that of x' . It remains to show that the solution S is feasible. Consider some point p and suppose it lies in class i in $\mathcal{I}'$, i.e., its demand $d'(p) = 2^i c'_{\min}$. We now wish to calculate the extent to which the demand of p is satisfied by $\bigcup_\ell S_\ell$. Intuitively, due to the rounding in the definition of d_p^ℓ , we lose at most one unit per class ℓ , which sums up to a total loss of at most $d'(p)$, which can be tolerated since in the fractional solution the point p is covered $2d'(p)$ times. More precisely, we have

$$\begin{aligned} \sum_{\ell < i} 2^\ell c'_{\min} d_p^\ell &= \sum_{\ell < i} 2^\ell c'_{\min} \left\lfloor \sum_{r:p \in r, r \in B_\ell} x'_r \right\rfloor \\ &\geq \sum_{\ell < i} 2^\ell c'_{\min} \left(\left(\sum_{r:p \in r, r \in B_\ell} x'_r \right) - 1 \right) \\ &\geq \left(\sum_{\ell < i} \sum_{r:p \in r, r \in B_\ell} x'_r c'_r \right) - 2^i c'_{\min} \\ &\geq 2d'(p) - d'(p) = d'(p), \end{aligned}$$

where the last inequality follows from inequality (3.3) and as $d'(p) = 2^i c'_{\min}$. $\square$

By Theorem 3.5 and Lemma 3.3, it follows that there is an LP based $O(1)$ approximation for R2M. Then Lemma 3.9 follows immediately by the application of Lemma 3.8.

LEMMA 3.9. *The algorithm finds a cover for light points of value at most $O(1)$ times the cost of solution x' and hence at most $O(1)$ times the optimum R2C solution.*

Lemmas 3.6 and 3.9 together give Theorem 1.3.

4. Concluding remarks. We showed how GSP can be viewed as a geometric covering problem, and we used this connection to give an $O(\log \log P)$ approximation for it and a 16-approximation for the case with identical release times. Given the lack of any inapproximability result for GSP, it seems reasonable to conjecture that an $O(1)$ approximation should exist. Recently, Bansal, Krishnaswamy, and Saha [4] showed that the geometric approach may be inherently lossy. Using the recent breakthrough constructions of Pach and Tardos [27] on lower bounds on size of ϵ -nets for geometric objects, they show that the LP relaxation for a general instance of the R3U problem

(which is a special case of the general R2C problem, when capacities of rectangles are well separated) has an integrality gap of $\Omega(\log \log P)$. Thus for the geometric approach to be of further use, some additional property in the reduction from GSP to R2C must be used.

Indeed, one can bypass the geometric approach altogether and give a direct LP relaxation for GSP based on KC inequalities. This was also our original approach to this problem, but our rounding techniques led us to the geometric approach (which is cleaner to present as several geometric results can be used as a black-box). The same LP was used in the recent works on the special case of identical release dates [19, 21, 26]. We conjecture that this LP has $O(1)$ integrality gap even for general release times.

Direct LP formulation of GSP. For each job j and time t we define a variable $x_{j,t}$ that is intended to be 1 if the job is unfinished at time t and is 0 otherwise. Note that one only needs to define this variable for *relevant* times t (i.e., release times, or when a weight $w_j(t)$ first reaches a power of 2), but it is convenient to assume that it is defined for each t . Furthermore we assume that $w_j(r_j) = 0$. This does not affect anything as each job j has $p_j > 0$ and hence is unfinished at r_j . As previously, an interval $I = [s, t]$ consists of the time slots $[s, s+1], \dots, [t-1, t]$ and $|I| = t - s$. $X(I)$ denotes the set of jobs that arrive during I (i.e., $r_j \in [s, t]$) and $\xi(I) = \max(0, p(X(I)) - |I|)$.

$$\begin{aligned} \min \sum_j \sum_{t > r_j} x_{j,t} (w_{j,t} - w_{j,t-1}) \\ \text{s.t.} \quad & x_{j,t} \leq x_{j,t-1} \quad \forall j, t > r_j, \\ \sum_{j \in X(I) \setminus S} \min(p_j, \xi(I) - p(S)) x_{j,t} &\geq \xi(I) - p(S) \\ &\forall I = [s, t], \forall S \subset X(I) \text{ with } p(S) \leq \xi(I), \\ x_{j,r_j} &= 1 \quad \forall j, \\ x_{j,t} &\geq 0 \quad \forall j, t \geq r_j. \end{aligned}$$

Observe that if a job completes at time t , then in the intended solution $x_{j,t} = 0$ and $x_{j,t'} = 1$ for $t' < t$, and hence the contribution to the objective is $\sum_{r_j < t' < t} (w_{j,t} - w_{j,t-1}) = w_j(t) - w_j(r_j) = w_j(t)$. The first set of constraints says that $x_{j,t}$ are monotonically nonincreasing, and the third set of constraints says that all the jobs are unfinished when they arrive. The second set of constraints is KC inequalities applied to the constraint that for each interval $I = [s, t]$ the jobs that finish after t must have a total size of at least $p(X(I))$.

Acknowledgments. We greatly thank Alexander Souza, Cliff Stein, Lap-Kei Lee, Ho-Leung Chan, and Pan Jiangwei for helpful discussions about this research.

REFERENCES

- [1] N. BANSAL, N. BUCHBINDER, AND J. NAOR, *Randomized competitive algorithms for generalized caching*, SIAM J. Comput., 41 (2012), pp. 391–414.
- [2] N. BANSAL AND K. DHAMDHERE, *Minimizing weighted flow time*, ACM Trans. Algorithms, 3 (2007).
- [3] N. BANSAL, A. GUPTA, AND R. KRISHNASWAMY, *A constant factor approximation algorithm for generalized min-sum set cover*, in Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, 2010, pp. 1539–1545.
- [4] N. BANSAL, R. KRISHNASWAMY, AND B. SAHA, *On capacitated set cover problems*, in Proceedings of APPROX-RANDOM, 2011, pp. 38–49.

- [5] N. BANSAL AND K. PRUHS, *Server scheduling to balance priorities, fairness, and average quality of service*, SIAM J. Comput., 39 (2010), pp. 3311–3335.
- [6] N. BANSAL AND K. PRUHS, *Weighted geometric set multi-cover via quasi-uniform sampling*, in Proceedings of the European Symposium on Algorithms, 2012, pp. 145–156.
- [7] A. BAR-NOY, R. BAR-YEHUDA, A. FREUND, J. NAOR, AND B. SCHIEBER, *A unified approach to approximating resource allocation and scheduling*, J. ACM, 48 (2001), pp. 1069–1090.
- [8] R. BAR-YEHUDA AND D. RAWITZ, *On the equivalence between the primal-dual schema and the local ratio technique*, SIAM J. Discrete Math., 19 (2005), pp. 762–797.
- [9] M. A. BENDER, S. MUTHUKRISHNAN, AND R. RAJARAMAN, *Approximation algorithms for average stretch scheduling*, J. Scheduling, 7 (2004), pp. 195–222.
- [10] H. BRÖNNIMANN AND M. T. GOODRICH, *Almost optimal set covers in finite VC-dimension*, Discrete Comput. Geom., 14 (1995), pp. 463–479.
- [11] T. CARNES AND D. B. SHMOYS, *Primal-dual schema for capacitated covering problems*, in Proceedings of the Conference on Integer Programming and Combinatorial Optimization, 2008, pp. 288–302.
- [12] R. D. CARR, L. FLEISCHER, V. J. LEUNG, AND C. A. PHILLIPS, *Strengthening integrality gaps for capacitated network design and covering problems*, in Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, 2000, pp. 106–115.
- [13] D. CHAKRABARTY, E. GRANT, AND J. KONEMANN, *On column restricted and priority integer covering programs*, in Proceedings of the Conference on Integer Programming and Combinatorial Optimization, 2010.
- [14] T. M. CHAN, E. GRANT, J. KÖNEMANN, AND M. SHARPE, *Weighted capacitated, priority, and geometric set cover via improved quasi-uniform sampling*, in Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, 2012, pp. 1576–1585.
- [15] C. CHEKURI, K. L. CLARKSON, AND S. HAR-PELED, *On the set multicover problem in geometric settings*, ACM Trans. Algorithms, 9 (2012).
- [16] C. CHEKURI AND S. KHANNA, *Approximation schemes for preemptive weighted flow time*, in Proceedings of the ACM Symposium on Theory of Computing, 2002, pp. 297–305.
- [17] C. CHEKURI, S. KHANNA, AND A. ZHU, *Algorithms for minimizing weighted flow time*, in Proceedings of the ACM Symposium on Theory of Computing, 2001, pp. 84–93.
- [18] T. C. E. CHENG, C. T. NG, J. J. YUAN, AND Z. H. LIU, *Single machine scheduling to minimize total weighted tardiness*, European J. Oper. Res., 165 (2005), pp. 423–443.
- [19] M. CHEUNG AND D. B. SHMOYS, *A primal-dual approximation algorithm for min-sum single-machine scheduling problems*, in Proceedings of APPROX-RANDOM, 2011, pp. 135–146.
- [20] K. L. CLARKSON AND K. R. VARADARAJAN, *Improved approximation algorithms for geometric set cover*, Discrete Comput. Geom., 37 (2007), pp. 43–58.
- [21] W. HÖHN, J. MESTRE, AND A. WIESE, *How Unsplittable-Flow-Covering Helps Scheduling with Job-Dependent Cost Functions*, arXiv:1403.1376 [cs.DS], 2014.
- [22] G. KARAKOSTAS, S. G. KOLLIPOULOS, AND J. WANG, *An FPTAS for the minimum total weighted tardiness problem with a fixed number of distinct due dates*, in Proceedings of the International Computing and Combinatorics Conference, 2009, pp. 238–248.
- [23] E. L. LAWLER, *A fully polynomial approximation scheme for the total tardiness problem*, Oper. Res. Lett., 1 (1982), pp. 207–208.
- [24] R. LEVI, A. LODI, AND M. SVIRIDENKO, *Approximation algorithms for the capacitated multi-item lot-sizing problem via flow-cover inequalities*, Math. Oper. Res., 33 (2008).
- [25] J. MATOUSEK, *Lectures on Discrete Geometry*, Springer, New York, 2002.
- [26] J. MESTRE AND J. VERSCHAE, *A 4-Approximation for Scheduling on a Single Machine with General Cost Function*, arXiv:1403.0298 [cs.DS], 2014.
- [27] J. PACH AND G. TARDOS, *Tight lower bounds for the size of epsilon-nets*, in Proceedings of the Symposium on Computational Geometry, 2011, pp. 458–463.
- [28] K. PRUHS, J. SGALL, AND E. TORNG, *Online scheduling*, in Handbook on Scheduling, CRC Press, Boca Raton, FL, 2004.
- [29] K. R. VARADARAJAN, *Epsilon nets and union complexity*, in Proceedings of the Symposium on Computational Geometry, 2009, pp. 11–16.
- [30] K. R. VARADARAJAN, *Weighted geometric set cover via quasi-uniform sampling*, in Proceedings of the ACM Symposium on Theory of Computing, 2010, pp. 641–648.