

Approximation Algorithms for Generalized Path Scheduling

by

Haozhou Pang

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science

Department of Computing Science
University of Alberta

© Haozhou Pang, 2020

Abstract

Scheduling problems where the machines can be represented as the edges of a network and each job needs to be processed by a sequence of machines that form a path in this network have been the subject of many researches (e.g. flow shop is a special case where the network as well as the sequence of machines for each job is a simple path). In this thesis, we consider one such problem, called Generalized Path Scheduling (GPS) problem, which can be defined as follows. Given a set of non-preemptive jobs J and identical machines M ($|J| = n$ and $|M| = m$). The machines are ordered on a path. Each job $j = \{P_j = \{s_j, t_j\}, p_j\}$ is defined by its processing time p_j and a sub-path P_j from machine with index s_j to t_j ($s_j, t_j \in M$ and $s_j \leq t_j$) specifying the order of machines it must go through. We assume each machine has a queue of infinite size where jobs can sit in the queue to resolve conflicts. Two objective functions, makespan and total completion time, are considered in this thesis.

In Chapter 2, we show that the GPS problem is NP-hard for both makespan and total completion time objectives. To complement the hardness results, we present several improved approximation for both objectives in Chapter 3. For the case of number of machines being sub-logarithmic in the number of jobs, we present a PTAS for both makespan and total completion time. This generalizes the result of Hall [11] for the classic problem of Flow shop. For general m , we present a $O(\frac{\log m}{\log \log m})$ -approximation algorithms for both objectives, which improve the previous best result of [4].

Acknowledgements

First and foremost, I would like to thank my supervisor, Mohammad R. Salavatipour. He has everything one could ask for in an advisor: brilliant, approachable and patient, methodical and good-tempered. He gives his students the maximal freedom to explore and find their own interests, and at the same time, he serves as a solid backing and a vast source of knowledge that constantly providing reassuring advice. This thesis would not have been possible without his help.

I wish to thank my committee members, Zachary Friggstad and Guohui Lin, for providing insightful suggestions that make this thesis a better one.

I would like to thank all members of the Theory group, with whom I have spent countless hours learning, discussing, and joking. Thank you all for making my time in the lab more enjoyable.

Last but not least, to my parents, Jianfeng Kong and Min Pang, and my parter, Jin. They provided much emotional support at times when I needed the most. This thesis is dedicated to them.

Contents

1	Introduction	1
1.1	Scheduling Fundamentals	2
1.2	Problem Formulation and Notations	3
1.3	Hardness Fundamentals and Approximation Algorithms	5
1.4	Literature Review	6
2	Hardness Results	9
2.1	Makespan objective	10
2.2	Min-sum objective	11
3	Approximation Algorithms	20
3.1	A PTAS for the makespan objective when m is sub-logarithmic . . .	20
3.1.1	The outline scheme	21
3.1.2	The PTAS	27
3.2	A $O(\frac{\log m}{\log \log m})$ -approximation for makespan objective	29
3.2.1	Instances with h terminal machines	30
3.3	A PTAS for min-sum GPS when m is sub-logarithmic	33
3.3.1	Segmented GPS	34
3.3.2	The PTAS	36
3.4	An $O(\frac{\log m}{\log \log m})$ -approximation for min-sum GPS	42
4	Conclusions	45

List of Figures

1.1	An instance of GPS that consists of three machines and two jobs. Job j_1 needs to be processed on M_1 , M_2 , and M_3 and has processing time $p_1 = 1$. j_2 needs to be processed on M_2 and M_3 and has processing time $p_2 = 3$. The smallest makespan is achieved by the schedule that j_1 get processed during $[0, 1]$, $[3, 4]$, $[6, 7]$, and j_2 get processed during $[0, 3]$, $[3, 6]$, so the makespan is 7. The smallest total completion time is achieved by the schedule that j_1 get processed during $[0, 1]$, $[1, 2]$, $[2, 3]$, and j_2 is processed during $[2, 5]$, $[5, 8]$, so the total completion time is 11.	4
2.1	An illustration of the instance of GPS reduced from a 3-PARTITION instance. The given machines are $M = \{M_1, M_2, \dots, M_{4t-2}\}$ and the given jobs are $J = \{T_1, \dots, T_t\} \cup \{U_1, \dots, U_{3t}\}$. U-type jobs have span 1 and only need to be processed on M_{2t} . T-type jobs have span $2t$.	10
2.2	The reduced instance of GPS is represented by (a). And (b) shows the timeline of machine M_l , where the shaded areas are the gaps created by the T-type jobs. (c) gives a closer look of each gap. The W-type jobs check the existence of a partition. The V-type and X-type jobs are used to fix T-type jobs.	12
3.1	the LP to assign small jobs. Recall that p_j is the processing time of job j , and P_j is the path of machines for job j .	23

3.2 In this example, $h = 3$, the dark squares represent the terminal machines in the first layer, which partition the machines into 3 segments. The squares with crossed lines represent the second layer terminal machines. The corresponding groups of jobs are defined as $G_1 = \{j_1, j_2\}$ and $G_2 = \{j_3, j_4, j_5\}$ 30

3.3 Partition the time line into intervals of size δ , which is defined w.r.t. B_1 . 35

3.4 the modified LP. π new constraints are added to ensure that at least $n_i^s = n_i - n_i^l$ small jobs are finished before time B_i 36

Chapter 1

Introduction

Scheduling problems are well-studied over the last several decades. One of the most classical version of it is as follows: Given a set J of n jobs and a set M of m machines. Each job j consists of a sequence of λ_j operations $O_{1,j}, O_{2,j}, \dots, O_{\lambda_j,j}$. The amount of time that job j takes to complete its operation $O_{i,j}$ on machine $M_i \in M$ is denoted as $p_{i,j}$. The goal of the problem is to find a feasible schedule that satisfies all constraints, while trying to optimize some objective function.

A feasible schedule should describe the times when each job starts being processed on each machine. Also, a feasible schedule should respect the following constraints: Each machine can only process at most one job at a time. Each job can only be processed by at most one machine at a time. More constraints can be added depending on specific interests. For example, one can enforce that operations of a job need to be processed in a specific order (known as precedence constraints), so an operation cannot be processed until all its preceding operations are finished.

Scheduling problems have drawn much attention because of their wide applications in many day-to-day situations. As an example, think of the given machines as routers and jobs as messages to be sent from one router to another through a specified path. A good scheduling algorithm can be applied here to send all messages efficiently. Also, as studied in [5], imagine the given machines form a star graph and jobs are supposed to start and end on leaves. This problem is closely related to the *biprocessor*

scheduling and *data migration* problem.

In Section 1.1 we introduce some basics about scheduling problems. Then, in Section 1.2, we formally state the problem studied in this thesis and the notations that are used. In Section 1.3, we give a brief introduction to approximation algorithms and NP-hardness. Next, we review related works in the literature and their connections to our problem in Section 1.4.

1.1 Scheduling Fundamentals

Many variants of the scheduling problem have arised in the past, and they can be characterized by the three-field notation $\alpha|\beta|\gamma$ introduced by [10], where α describes the machine enviroment, β describes the job type, and γ describes the optimalilty criteria.

Machine enviroments: Problems can be characterized by the machines. For example, when the processing time of job j on machine i , p_{ij} , is independent of machines, i.e. $p_{ij} = p_j \forall i$, then we say the machines are *identical*. If $p_{ij} = \frac{p_j}{s_i}$, where s_i is the speed of machine i , then we say the machines are *related*. Otherwise, they are *unrelated*.

Job types: Jobs can be either preemptive or non-preemptive. In the non-preemptive setting, suspension of a job in the middle of processing is not allowed. That is, once a machine has started processing a job, it must continue running until the current job is finished. If all jobs have the same processing time, then we have the *unit processing time* case.

If the operations of a job can be processed in arbitrary order, then we get the open shop problem. Conversely, if operations of a job has to be processed in some sopecific order, then we get the job shop problem. A special case of job shop called the acyclic job shop problem requires each job has at most one operation on each machine.

Jobs can also be made associated with *release times* r_j and *deadlines* d_j . They enforce that job j cannot start being processed before r_j and has to be finished before

d_j . We can remove those two constraints by saying $r_j = 0$ and $d_j = \infty$, for all j .

Optimality criteria: The two common objective functions that we consider in this thesis, which have been studied extensively in the literature, are the makespan and total completion time. Let C_j be the completion time of job j (the time when j finishes its last operation) in a given schedule. The makespan of the schedule is defined as $C_{max} := \max\{C_1, \dots, C_n\}$, and the total completion time is defined as $\sum_{j=1}^n C_j$. The latter is also referred as the min-sum objective in this thesis.

Other usual objectives include but not limit to min-max or min-sum version of lateness $L_j = C_j - d_j$, earliness $E_j = \max\{0, d_j - C_j\}$, and tardiness $T_j = \max\{0, C_j - d_j\}$

1.2 Problem Formulation and Notations

The problems studied in this thesis have been studied in [5, 13]. Authors of [5] point out that getting a constant factor approximation for acyclic flow shop with identical machines where the machines form a path and each job is assigned a subpath of machines to be processed on is still an open problem. We call this problem GENERALIZED PATH SCHEDULING (GPS) problem, it is also known as *message routing along a directed path* in [13]. It can be formally described as:

Definition 1 (Generalized Path Scheduling) *Given a set of non-preemptive jobs J (available at time 0) and identical machines M ($|J| = n$ and $|M| = m$). The underlying network of machines is a path. Each job $j = \{P_j = \{s_j, t_j\}, p_j\}$ is defined by its processing time p_j and a sub-path P_j from s_j to t_j ($s_j, t_j \in M$, and $s_j \leq t_j$) specifying the order of machines it must go through. Each machine has a queue of infinite size where jobs can sit in the queue to resolve conflicts, jobs sitting in the queue do not occupy the machine. We consider minimizing makespan (largest completion time) and/or total completion time (also called min-sum).*




Figure 1.1: An instance of GPS that consists of three machines and two jobs. Job j_1 needs to be processed on M_1 , M_2 , and M_3 and has processing time $p_1 = 1$. j_2 needs to be processed on M_2 and M_3 and has processing time $p_2 = 3$. The smallest makespan is achieved by the schedule that j_1 get processed during $[0, 1]$, $[3, 4]$, $[6, 7]$, and j_2 get processed during $[0, 3]$, $[3, 6]$, so the makespan is 7. The smallest total completion time is achieved by the schedule that j_1 get processed during $[0, 1]$, $[1, 2]$, $[2, 3]$, and j_2 is processed during $[2, 5]$, $[5, 8]$, so the total completion time is 11.

The *span* of a job j , $\lambda_j = t_j - s_j + 1$, is the number of machines on the path of j . The *length* of job j , $L_j = p_j \cdot \lambda_j$, is the minimum total time that job j needs to complete. We say job j is *delayed* on machine i if the start time of j on i is strictly greater than its arrival time. The completion time C_j of job j is then equal to its *length* plus the total amount of time it has been delayed. The GPS problem is closely related to the FLOWSHOP SCHEDULING problem, except in GPS, jobs can start/end at any machine. In Chapter 2, we show that the GPS problem is NP-hard for both makespan and min-sum objectives. So we resign to good approximation algorithms that run in polynomial time.

For the makespan objective, let C be the largest *congestion* over all machines (the maximum total running time of jobs that use machine i over all machines) and D be the maximum *length* over all jobs. Then clearly, both C and D are lower bounds for the makespan of the optimal schedule. In [13], Koch et al. give another lower bound on the minimum completion time on machines. Let $k, l \in M$ with $k \leq l$ be two machines and $j \in J$ be a job that traverses both of these machines. Let $C^{k \rightarrow l}$ be the total size of jobs passing k and l (their common congestion). Then the minimum completion time of machine l , denoted by MAK_l , is at least $C^{k \rightarrow l} + p_j \cdot (l - k)$.

1.3 Hardness Fundamentals and Approximation Algorithms

The definitions here are from [25].

NP: A Turing machine M is a polynomial-time verifier if for all input pairs $(x, y) \in \{0, 1\}^*$ it reads, M halts after $p(|x|)$ steps where p is a polynomial function. A language L is decidable by M if for all $x \in \{0, 1\}^*$,

- if $x \in L$, there exists $y \in \{0, 1\}^{p(|x|)}$ such that M accepts (x, y) .
- if $x \notin L$, for every $y \in \{0, 1\}^*$, M rejects (x, y) .

The class **NP** is all languages decided by a polynomial-time verifier.

Reduction: Let L and L' be two languages in **NP**, we say L reduces to L' if there exists a poly-time computable function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ such that $x \in L$ if and only if $f(x) \in L'$. A language L is **NP-hard** if for all $L' \in \mathbf{NP}$, L' reduces to L . A language is **NP-complete** if it is in the class **NP** and is **NP-hard**.

NP optimization problems: An NP optimization (minimization or maximization) problem consists of the following components

- A set of all valid instances D_{Π} . Each valid instance $I \in D_{\Pi}$ is recognizable in polynomial time (in terms of the size of the input $|I|$).
- A set of feasible solutions $S_{\Pi}(I)$ for each $I \in D_{\Pi}$, and the size of each solution $s \in S_{\Pi}(I)$, $|s|$, is polynomial in $|I|$.
- A poly-time computable objective function f that assigns a non-negative rational value to each feasible solution s of instance I . For maximization (minimization) problem, our goal is to find a feasible solution s that has the largest (smallest) objective value. Such solution is called the optimal solution for I , we denote its value by $OPT(I)$ or simply OPT if the context is clear.

Approximation algorithms: However, many NP optimization problems are **NP**-hard, meaning that it's not possible to find the optimal solution for all instances in polynomial time with the assumption $\mathbf{P} \neq \mathbf{NP}$. Therefore, we try to find solutions that are nearly good (with provable guarantee) in polynomial time. More precisely, we say A is an α -approximation algorithm for problem Π if A runs in polynomial time and for any instance $I \in D_\Pi$, it finds a feasible solution s such that $f(s, I) \leq \alpha \cdot OPT(I)$ if Π is a minimization problem; or $f(s, I) \geq \frac{OPT}{\alpha}$ if Π is a maximization problem. α is called *approximation ratio* or *performance ratio*.

For an **NP**-hard optimization problem, the best approximation algorithm that one can hope for is a PTAS (Polynomial Time Approximation Scheme) (assuming $\mathbf{P} \neq \mathbf{NP}$), which is an algorithm A that for any instance I and any fixed value $\epsilon > 0$, it finds a solution s in polynomial time such that $f(s, I) \leq (1 + \epsilon) \cdot OPT(I)$. The approximation ratio can be made arbitrarily close to 1 with a tradeoff of runtime (but still polynomially bounded). If runtime of A is polynomially bounded by both $|I|$ and $1/\epsilon$, we say A is a Fully Polynomial Time Approximation Scheme.

1.4 Literature Review

One of the most general version of scheduling problem is the job shop scheduling with unrelated machines. The first polynomial time approximation algorithm for this problem is an $O(\frac{\log^2(m\lambda_{max})}{\log \log(m\lambda_{max})})$ -approximation (λ_{max} is the maximum span) for the makespan objective, given by [21]. Later, [9] improves the result by a $O(\log \log(m\lambda_{max}))$ factor, this is also the best known result for this problem. If the amount of time that every job takes to be processed on any machine is the same, then we get the packet routing problem when machines are edges of a graph. Leighton et al. [16, 17] shows that there always exist a schedule of length $O(lb)$, where $lb = \max\{C, D\}$ is the trivial congestion/dilation lower bound for makespan objective. Later, the authors in [12] present a constructive algorithm that finds a schedule of length at most $8.84(C + D)$. However, the algorithms for the unit-processing time case seem hard to be adapted

to the general processing time case, because they use the idea of iteratively partitioning the schedule into frames (time intervals), and view each frame as a separate scheduling problem where the origin of a packet is its location at the beginning of the frame, and its destination is its location at the end of the frame. This could be done because all operations have size 1 so no operations would straddle between two frames. However, for the general processing time case, the frame size cannot be smaller than p_{max} .

Li et al. [18] show that if there is a α -approximation w.r.t. the lower bound $lb = \max\{C, D\}$ for the makespan objective, then there is a $2e\alpha$ -approximation algorithm for the min-sum objective (e is the euler's number). This provides a framework of converting the makespan objective to the min-sum objective without affecting the approximation ratio asymptotically.

Acyclic job shop is a special case of the Job Shop problem, where each job can have at most one operation on each machine. Feige et al. [3] give an $O(\log lb \log \log lb)$ -approximation for the makespan objective of this problem. They also show the upper bound is nearly tight by proving the existence of instances of shortest makespan $\Omega(\frac{lb \log lb}{\log \log lb})$ even when machines are identical. The GPS problem considered in this thesis is a special case of the acyclic job shop, where the machines are identical and form a path.

The best known result for the GPS problem is due to [5], they present an $O(\min\{\log n \lambda_{max}, \log p_{max}\})$ -approximation for (the more general problem of) acyclic job shop with identical machines, under both the makespan and min-sum objective. However, many special cases of GPS problem can actually be solved exactly in polynomial time or have an $O(1)$ -approximation algorithms. For example, if the network of the machines form a rooted tree and all the job paths have to go through the root of the tree, then the problem becomes the *junction tree* problem studied in [5], for which they present a 4-approximation for makespan and $8e$ -approximation for min-sum. For the special case where all jobs have the same processing time, [1, 15] show that the

greedy algorithm such that each machine processes the job with the largest number of remaining machines (furthest-to-go) gives the optimal makespan. Conversely, [1] shows shortest-to-go gives optimal min-sum for unit-processing time case. Moreover, authors of [13] show that furthest-to-go algorithm computes the optimal makespan on non-nested instances for general processing times; non-nested means there does not exist two jobs j and j' such that $s_j < s_{j'}$ and $t_j > t_{j'}$ (recall that $s_i(t_j)$ is the index of the first (last) machine of job j). If all jobs need to be processed on all (identical) machines from left to right, then the problem becomes the proportionate flow shop. It is straightforward that any fixed priority rule would give optimal for the makespan objective; for the (weighted) min-sum objective, [20] gives an exact algorithm that runs in $O(n^2)$ time. Bi-directional version of the problem, where there are jobs moving from left to right and right to left, can be dealt with by interleaving the uni-directional algorithms

If $m = O(1)$, then many scheduling problems admit better approximation ratios. For example, [19] gives a PTAS for the open shop makespan minimization problem, [21] gives a $(2 + \epsilon)$ -approximation algorithm for job shop, and [11] gives a PTAS for the flow shop. Note that all the results discussed are based on the fact that m is fixed. Hall [11] introduces the notion of *outline scheme*, which is used in a couple of our algorithms in a fundamental way in this thesis.

The scheduling problems where networks of machines have other specific structures have been the subjects of many researches. For example, when the machines form a grid, [14] shows that by applying the furthest-to-go algorithm vertically and horizontally they get a 3-approximation for the makespan for unit processing time case. Same approximation ratio applies to the rooted tree network, where jobs can either go vertically upward, vertically downward, or upward-downward. When the network of machines is a star and jobs start/end at leaves, [5] gives a 1.796-approximation for the min-sum objective, and a 7.279-approximation for the general processing time case.

Chapter 2

Hardness Results

In this chapter, we present some hardness results for the GENERALIZED PATH SCHEDULING problem for both *makespan* and *min-sum* objectives. The most important part of an NP-hardness proof is the ‘reduction’, we need to show that a known NP-hard problem A can be *reduced* to our problem B . The reduction is defined as given an instance a of A we can construct an instance b of B s.t. a has a solution if and only if b has a solution. Also, both the input size of b and the time for constructing b have to be polynomially bounded by the input size of a . The ‘known problem’ that we use to prove the NP-hardness of our problem is the 3-PARTITION problem, which is quite commonly used to prove NP-hardness for different variants of scheduling problems. See [26] for an example. Here we define the 3-PARTITION problem:

Definition 2 (3-Partition) *Given $3t$ integers a_1, a_2, \dots, a_{3t} and an integer b , can we partition the $3t$ integers into t triples such that each triple has sum equal to b ?*

Note that the 3-PARTITION problem is strongly NP-hard, and assuming that each a_i satisfies $b/4 < a_i < b/2$ does not affect the NP-hardness [8]. The general idea of the reduction for both objectives is that we view the $3t$ integers as $3t$ jobs whose processing times are equal to the values of the integers. Some other jobs are introduced to create so called *gaps* on the timeline of machines so that a pre-defined bound D can be achieved if and only if we can perfectly fill the *gaps* with the $3t$ jobs (i.e a

solution to the 3-PARTITION instance). Details of reduction are included in Sections 2.1 and 2.2.

2.1 Makespan objective

Theorem 3 *The GENERALIZED PATH SCHEDULING makespan minimization problem is NP-hard.*

Proof. Suppose we are given a 3-PARTITION instance a_1, \dots, a_{3t} , and b . We construct the following instance of GPS that uses $n = 4t$ jobs and $m = 4t - 2$ machines. The processing times and paths of jobs are specified as follows.

We define two classes of jobs: the T-type jobs and the U-type jobs where:

$$T = \{T_i | 1 \leq i \leq t\},$$

$$U = \{U_i | 1 \leq i \leq 3t\},$$

T_i has proc. time $p_{T_i} = b$ and path $P_{T_i} = \{M_{2(t-i)+1}, M_{4t-2i}\}$;

U_i has proc. time $p_{U_i} = a_i$ and path $P_{U_i} = \{M_{2t}, M_{2t}\}$.




Figure 2.1: An illustration of the instance of GPS reduced from a 3-PARTITION instance. The given machines are $M = \{M_1, M_2, \dots, M_{4t-2}\}$ and the given jobs are $J = \{T_1, \dots, T_t\} \cup \{U_1, \dots, U_{3t}\}$. U-type jobs have span 1 and only need to be processed on M_{2t} . T-type jobs have span $2t$.

And we define the bound $D = 2bt$. We argue that a schedule of makespan at most D can be obtained if and only if the 3-PARTITION instance has a solution.

First suppose there is a good partition into triples of the desired form. Consider the T-type jobs on machine M_{2t} . Job T_i will arrive at machine M_{2t} at time $(2i - 1)b$. Observe that every T-type job has *length* $2bt$, so none of them can be delayed, otherwise, the bound D is exceeded. More precisely, job T_i must use machine M_{2t} during time $[(2i - 1)b, 2ib]$. Therefore, there is a *gap* of size b between T_i and T_{i+1} on machine M_{2t} , and we can use the t triples of U-type jobs to fill the gaps. And the makespan of the schedule is exactly D , as desired.

Conversely, Suppose there does not exist a good partition, then there is a U-type job u that cannot be fitted into one of the gaps. Job u cannot delay T-type jobs so u cannot start processing until time $2bt$ so the makespan is at least $D + p_u > D$.

■

2.2 Min-sum objective

In this section, we show that the GPS min-sum minimization problem is also NP-hard. The NP-hardness proof generalizes the idea in [7], which shows the NP-hardness of the *2-machine flowshop* problem. Similar to the proof of Theorem 3, we reduce the 3-PARTITION problem to the GPS min-sum minimization problem by using a class of jobs to create *gaps* and using some other jobs to fill the *gaps* and check the existence of a solution to the 3-PARTITION instance. The analysis is more involved.

Theorem 4 *The GENERALIZED PATH SCHEDULING min-sum minimization problem is NP-hard.*

Proof. For a given instance of 3-PARTITION a_1, \dots, a_{3t} and b , we define the parameters that are used throughout the proof as:

$$\begin{aligned} u &= 3tb + 1 \\ v &= u + 3tb + tu + t(t - 1)u(b + 1)/2 \\ g &= uv + b + 1 \\ x &= 2(t + 2)g + v \end{aligned}$$

We then construct the corresponding instance of the GPS problem (Fig. 2.2). The machines $M_1, M_2, \dots, M_{tg+1}$ form a path, and we introduce four classes of jobs:

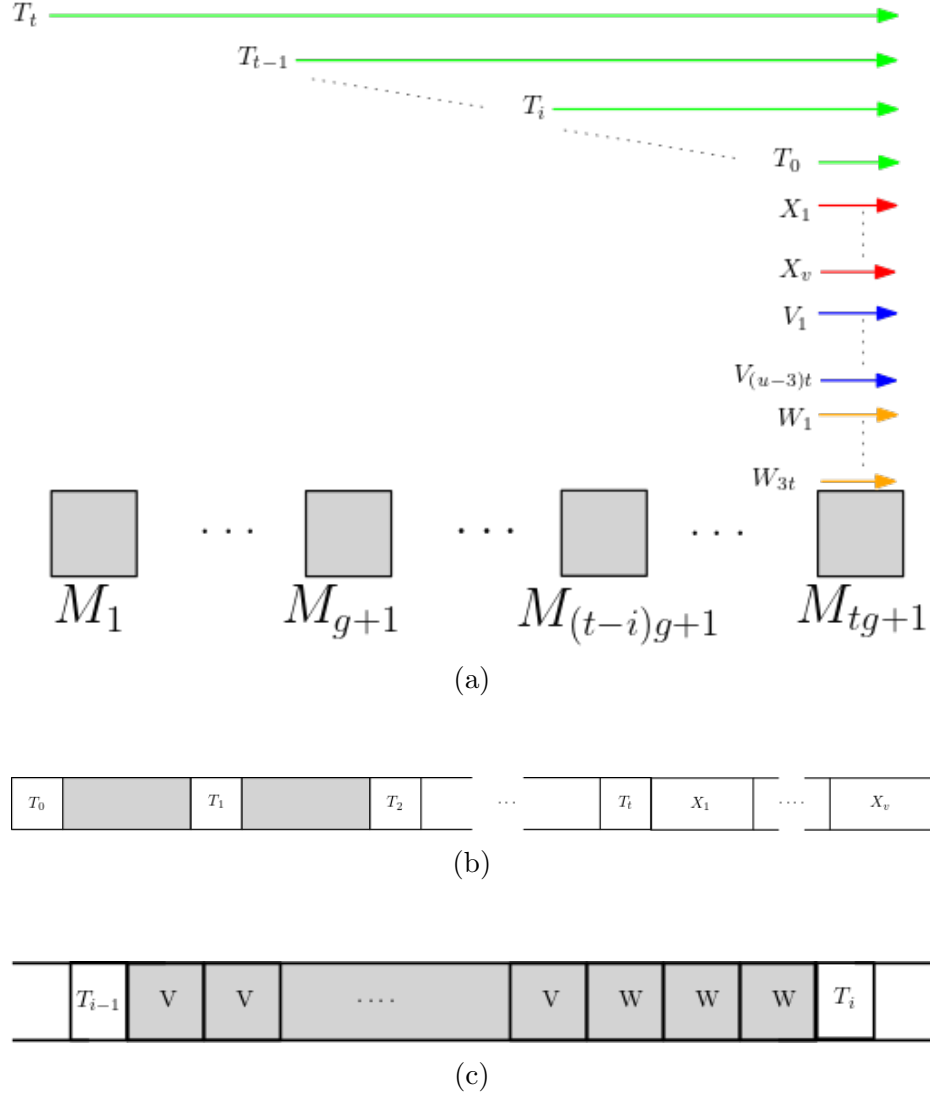


Figure 2.2: The reduced instance of GPS is represented by (a). And (b) shows the timeline of machine M_t , where the shaded areas are the gaps created by the T-type jobs. (c) gives a closer look of each gap. The W-type jobs check the existence of a partition. The V-type and X-type jobs are used to fix T-type jobs.

$$\begin{array}{ll}
T = \{T_i | 0 \leq i \leq t\} & |T| = t + 1 \\
X = \{X_i | 1 \leq i \leq v\} & |X| = v \\
V = \{V_i | 1 \leq i \leq (u - 3)t\} & |V| = (u - 3)t \\
W = \{W_i | 1 \leq i \leq 3t\} & |W| = 3t
\end{array}$$

The T-type jobs are used to create *gaps*. For all $0 \leq i \leq t$, job T_i has processing time $p_{T_i} = 1$ and path $P_{T_i} = \{M_{(t-i)g+1}, M_{tg+1}\}$. Observe that if we ignore the presence of other jobs, those $t + 1$ many T-type jobs create t *gaps* of size $g - 1$ on machine M_{tg+1} . Those *gaps* will be used to test the existence of a solution to the 3-PARTITION instance. For the ease of notation, we denote the last machine M_{tg+1} by M_l .

However, for the *gaps* to be fixed, we need to ensure that every T-type job $T_i, 0 \leq i \leq t$, is scheduled on machine M_l immediately as soon as it becomes ready. To achieve this, we introduce another class of jobs, the X-type jobs:

For all $1 \leq i \leq v$, job X_i has processing time $p_{X_i} = x$ and path $P_{X_i} = \{M_l, M_l\}$. Intuitively, the X-type jobs only run on the last machine and their processing time x is chosen to be large enough so that any reasonable schedule would process them at the end. Moreover, the value of v is chosen such that if job T_i is delayed even for one unit of time, the total completion time will exceed the pre-defined bound D .

The V-type jobs and W-type jobs are used to fill the *gaps*. For all $0 \leq i \leq (u - 3)t$, job V_i has processing time $p_{V_i} = v$ and path $P_{V_i} = \{M_l, M_l\}$; and for all $0 \leq i \leq 3t$, job W_i has processing time $p_{W_i} = v + a_i$ and path $P_{W_i} = \{M_l, M_l\}$, where the a_i 's come from the 3-Partition instance. The purpose of V-type jobs and W-type jobs is that if there exists a good partition of integers into triples, then we can fit three W-type jobs (representing the integers) and $(u - 3)$ V-type jobs in the gap created by T_i and T_{i+1} (the gap size is $g - 1 = uv + b$, sum of W-type jobs is $3v + b$, and the sum of V-type jobs is $(u - 3)v$). Observe that we have in total u jobs inside each *gap*,

the value of u is chosen to ensure that each job T_i , for $0 \leq i \leq t - 1$, is scheduled immediately after it becomes ready on M_l . Therefore, a delay of even one unit of time on those u many jobs would result in exceeding the pre-defined bound D on total completion time.

Note that the number of jobs we have is $(u + 1)t + v + 1$, and the number of machines is $tg + 1$ (u, v, g are all $\text{poly}(tb)$ based on their definitions). Therefore, the size of the instance, though large, is still polynomially bounded by tb .

Combine the V-type jobs and W-type jobs as U-type jobs, i.e. $U = V \cup W$. We define the bound on the total completion time by $D = \hat{T} + \hat{X} + \hat{U}$, where

$$\hat{T} = \sum_{i=0}^t (ig + 1) = t + 1 + \frac{t(t + 1)g}{2}$$

$$\hat{X} = \sum_{i=1}^v (tg + 1 + ix) = v(tg + 1) + \frac{v(v + 1)x}{2}$$

$$\begin{aligned} \hat{U} &= 3tb + \sum_{i=0}^{t-1} \left[\sum_{j=1}^u (jv + ig + 1) \right] \\ &= 3tb + tu + \frac{t(t - 1)u(b + 1)}{2} + \frac{tu(tu + 1)v}{2} \end{aligned}$$

Then it only remains to show that there exists a partition of integers into triples with sum b if and only if there exists a schedule of total completion time at most D .

Suppose there is a good partition A_1, A_2, \dots, A_t such that $A_i = \{a_{i_1}, a_{i_2}, a_{i_3}\}$ and $\sum_{j=1}^3 a_{i_j} = b$ for all i . Consider the following natural schedule: all T-type jobs start running from their starting machine towards M_l since time 0 and never wait for any other jobs. All X-type jobs start running one by one on M_l since time $tg + 1$. Consider machine M_l . There are t gaps of size $g - 1 = uv + b$ before time $tg + 1$ between T-type jobs. For each gap, we schedule $(u - 3)$ V-type jobs and 3 W-type jobs (taken from

corresponding triple), which takes time exactly $uv + b$. The feasibility of the schedule should be straightforward to verify, because a good partition would suggest how to fill those t gaps perfectly without conflicts.

The next step is to calculate the total completion time of the proposed schedule. The sum of completion times of U-type jobs is:

$$\begin{aligned}
\sum_{j \in U} C_j &= \sum_{i=0}^{t-1} \left[\sum_{j=1}^u (jv + ig + 1) + 3a_{i_1} + 2a_{i_2} + a_{i_3} \right] \\
&< \sum_{i=0}^{t-1} \left[\sum_{j=1}^u (jv + ig + 1) \right] + 3 \sum_{i=0}^{3t} a_i \\
&= \hat{U}
\end{aligned} \tag{2.1}$$

The first equality of (2.1) is because there are u many U-type jobs ($u - 3$ many V-type jobs and three W-type jobs) between the gap created by T_i and T_{i+1} . Those $u - 3$ many V-type jobs start running one by one since time $(ig + 1)$, and the three W-type jobs are scheduled at the end. The order of the W-type jobs does not matter, without loss of generality, we assume the job representing a_{i_1} is scheduled first, followed by a_{i_2} and a_{i_3} , that's why we have the additive term $3a_{i_1} + 2a_{i_2} + a_{i_3}$ in the equality.

Also, note that the sum of completion times of T-type jobs is exactly \hat{T} , and the sum of completion times of X-type jobs is exactly \hat{X} . Therefore, the total completion time of the proposed schedule is less than $\hat{T} + \hat{X} + \hat{U} = D$, as wanted.

Conversely, suppose there is a schedule with total completion time upper bounded by D . Then we show that there must exist a good partition. The basic idea is we show the properties that a good schedule must have, then we are able to conclude that if a good schedule exists, it must be like the one described in Fig. 2.2(b). Then the desired partition will be given by the W-type jobs that are grouped in each gap.

For a job j , let $S(j)$ denote the starting time of j on machine M_l . Without loss of generality, we can make the following assumptions on the properties that a good schedule should have:

Property 1: *T-type jobs are not delayed on machines M_1, M_2, \dots, M_{tg} .* This is because all T-type jobs have the same processing time and distinct starting machines. So assuming all T-type jobs start running at time 0, they will never conflict with any other jobs until the last machine M_l . This implies that T_i becomes ready on M_l at time ig , for $0 \leq i \leq t$.

Property 2: $S(T_i) < S(T_{i+1})$, for $0 \leq i \leq t$. This is because all T-type jobs have the same processing time. So if $S(T_j) > S(T_{j+1})$ for some j . We can just swap the order of T_j and T_{j+1} without changing the value of the solution.

Property 3: *Similarly, we can order the X-type jobs such that $S(X_i) < S(X_{i+1})$, $1 \leq i \leq v - 1$.* Because all X-type jobs are identical.

Property 4: *Let Σ be a good schedule having all previous properties, then all X-type jobs must start running at or after time $tg + 1$ in Σ . i.e., $S(X_i) \geq tg + 1$, $1 \leq i \leq v$.*

Suppose Σ is a minimum total completion time schedule satisfying all previous properties, but there is a job X_i such that $S(X_i) < tg + 1$ in Σ . Consider the first such job from X and so it must be X_1 . Moreover, by the time X_1 finishes, all jobs are ready to be scheduled on M_l . Therefore, in order to minimize the total completion time, we must schedule those jobs by processing the faster jobs first. Next, we show that if such a job X_1 exists, then Σ cannot be a good schedule.

Note that the completion time of X_1 is no earlier than x , and all other X-type jobs must be scheduled at the end. So we can obtain a lower bound $lb(X)$ on the total completion time of X-type jobs:

$$lb(X) = x + \sum_{i=1}^{v-1} (x + tg + 1 + ix) = \hat{X} - tg - 1$$

Since $S(X_1) < tg + 1$, jobs T_t has to wait until X_1 finishes on M_l . Also, the total completion times of jobs T_0, \dots, T_{t-1} is at least $\sum_{i=0}^{t-1} (ig + 1)$. So we obtain a lower

bound $lb(T)$ on the total completion time of T-type jobs:

$$lb(T) = \sum_{i=0}^{t-1} (ig + 1) + x = \hat{T} + (t + 4)g + v - 1$$

Also, a trivial lower bound $lb(U)$ on the total completion time of U-type jobs can be obtained by viewing them all as V-type jobs (ignore the a_i term of W-type jobs).

Then we can schedule them one by one:

$$\begin{aligned} lb(U) &= \sum_{i=1}^{tu} (iv) = \frac{tu(tu + 1)v}{2} \\ &= \hat{U} - \underbrace{\left[3tb + tu + \frac{t(t-1)u(b+1)}{2} \right]}_{\delta} \end{aligned}$$

Therefore, the total completion time of Σ is at least: $lb(T) + lb(X) + lb(U) = \hat{T} + \hat{X} + \hat{U} + v - \delta + 4g - 2$. Since $v > \delta$ and $4g > 2$, the total completion time exceeds the bound D , so Σ cannot be a good schedule.

Property 5: *Suppose Σ is a good scheduling with minimum total completion time.*

Then ① $S(X_1) = tg + 1$ and ② M_l is not idle before time $tg + 1$.

We first show ①. From property 4 we know that $S(X_1) \geq tg + 1$. For the sake of contradiction, suppose $S(X_1) > tg + 1$, then the total completion time of X-type jobs is at least:

$$\sum_{i=1}^v (tg + 2 + ix) = \hat{X} + v$$

The total completion time of T-type jobs is at least \hat{T} , and the lower bound $lb(U) = \hat{U} - \delta$ (from Property 4) still holds. Adding them up, since $v > \delta$, the total completion time of Σ is greater than D . Hence Σ cannot be a good schedule.

For ②, observe that the sum of processing times on M_l over all T-type and U-type jobs is equal to $t + 1 + (u - 3)tv + tb + 3tv$, which is exactly $tg + 1$. From ① we know that $S(X_1) = tg + 1$. Therefore, if M_l is idle before time $tg + 1$, there will be

a job j that is faster than X_1 but get processed after X_1 . So if we swap the order of j and X_1 on M_l , we get a feasible schedule with smaller total completion time. This contradicts with the optimality of Σ .

Property 6: *In a good schedule Σ , on machine M_l , every T-type job T_i is preceded by exactly iu many U-type jobs, $0 \leq i \leq t$. i.e., there are exactly u many U-type jobs between T_i and T_{i+1} .*

For the sake of contradiction, suppose job T_i is preceded by less than iu U-type jobs. Then the total processing times of jobs that start on M_l before T_i is at most $i + (iu - 1)v + tb$, which is less than ig . From Property 1 we know that T_i becomes ready on M_l at time ig . So machine M_l is idle before time ig , so before $tg + 1$. This contradicts with Property 5.

Conversely, suppose job T_i is preceded by more than iu U-type jobs. Then the total processing times of jobs that start on M_l before T_i is at least $i + (iu + 1)v = ig - ib + v$. Therefore, $S(T_i) \geq ig - ib + v$, and the total completion time of T-type jobs is at least:

$$\sum_{i=0}^t (ig + 1) - ib + v = \hat{T} - ib + v$$

By Properties 4 and 5, the total completion time of X-type jobs is at least \hat{X} . The lower bound $lb(U) = \hat{U} - \delta$ from Property 4 still holds here. Therefore, the total completion time of Σ is at least

$$\hat{T} + \hat{X} + \hat{U} - ib + v - \delta$$

Recall that $v = u + 3tb + tu + t(t - 1)u(b + 1)/2 = u + \delta$. Since $u > ib$, this value exceeds D , so Σ cannot be a good schedule.

Property 7: *In a good schedule Σ , $S(T_i) = ig$, for $0 \leq i \leq t$.*

From Property 1, $S(T_i) \geq ig$. For the sake of contradiction, suppose $S(T_i) > ig$ for some i . By Property 6, there are exactly u many U-type jobs between T_i and T_{i+1} . So the total completion time of those u jobs is at least $\sum_{j=1}^u (ig + 2 + jv)$. Therefore we obtain another lower bound on the total completion time of U-type jobs:

$$\begin{aligned}
lb(U) &= \sum_{i=0}^{t-1} \left[\sum_{j=1}^u (jv + ig + 1) \right] + u \\
&= \hat{U} - 3tb + u \\
&= \hat{U} + 1
\end{aligned} \tag{2.2}$$

The additive u in the first equation of (2.2) is because all the u jobs between T_i and T_{i+1} are additionally delayed by at least one unit of time. Combining with the lower bounds on T-type and X-type jobs, the total completion time of Σ is at least $\hat{T} + \hat{X} + \hat{U} + 1 > D$, so Σ is not a good schedule.

After showing the properties that a good schedule must have, we are ready to prove Theorem 4. Suppose there exists a good schedule Σ , then in Σ , job T_i starts processing on machine M_l at time ig , so there is a *gap* of size $g - 1 = uv + b$ between T_i and T_{i+1} . Also, there are exactly u many U-type jobs between T_i and T_{i+1} , recall that V-type jobs have processing time v and W-type jobs have processing time $v + a_i$, where a_i is between $\frac{b}{4}$ and $\frac{b}{2}$, so there must be exactly three W-type jobs inside each *gap*. These disjoint triples of W-type jobs form the desired partition. This completes the NP-hardness proof and concludes this chapter.

■

Chapter 3

Approximation Algorithms

Showing that the GPS problem is strongly NP-hard for both *makespan* and *min-sum* objectives implies that it is not possible to solve the problem exactly in polynomial time unless $P = NP$. Instead we look at solving the problem approximately, and efficiently. Several approximation algorithms are presented in this chapter.

In Section 3.1, we show that if the number of machines m is sub-logarithmic in the number of jobs n , more precisely $m = O(\frac{\log^{1/6} n}{\log \log n})$, then we have a $(1+\epsilon)$ -approximation (PTAS) for the *makespan* objective. Then in Section 3.2, we use this as a subroutine to get an $O(\frac{\log m}{\log \log m})$ -approximation for the general setting, which improves the previous best result by a double logarithmic factor. Next, in Section 3.3, we discuss the approximation algorithm for *min-sum* objective when $m = O(\frac{\log^{1/6} n}{\log \log n})$, for which we introduce a new variant of the GPS problem, called the *segmented* GPS problem and we give a $(1 + \epsilon)$ -approximation for it. This problem is important as we use it to convert the *makespan* objective to *min-sum* objective, i.e., the *min-sum* version of GPS can be reduced to this problem by small loss. Finally, in Section 3.4, we present an $O(\frac{\log m}{\log \log m})$ for min-sum GPS for general m .

3.1 A PTAS for the makespan objective when m is sub-logarithmic

The main theorem that we prove in this section is the following:

Theorem 5 *There is a PTAS for GPS with makespan objective when $m = O(\frac{\log^{1/6} n}{\log \log n})$.*

The algorithm is built based on *outline scheme* and linear program. Similar ideas have been used in [11, 19] to design PTAS for the flow shop and open shop problems when $m = O(1)$. The general framework is to label the jobs as *big* and *small* based on their processing times. It can be shown that the number of big jobs cannot be too large so we can guess (enumerate) their schedule on the machines with a good accuracy. For small jobs, we schedule them by using a Linear Programming (LP) relaxation and rounding with small error.

3.1.1 The outline scheme

Definition 6 *An outline scheme partitions all feasible solutions into classes (outlines), such that solutions that get grouped together share some common characteristics.*

The outline scheme should suggest a natural way to obtain a good schedule. Our goal is to show that: ① The number of *outlines* is polynomially bounded. ② For each *outline*, we can generate a schedule such that the makespan of the schedule is approximately $(1 + \epsilon)$ as good as the optimal schedule in this *outline*. Since the optimal schedule must be contained in one of those *outlines*, by enumerating all of them, we are guaranteed to find a nearly good schedule.

Suppose we have an upper bound T on the length of the optimal schedule T^* . Such an upper bound can be obtained by using a naive algorithm that simply processes operations starting from M_1 and move to the next machine if all operations on the current machine are finished. Therefore $T = mT^*$ is always a valid upper bound. Then we partition the time line from 0 to T into κ intervals of size $\delta = \frac{T}{\kappa}$, and we refer the interval $[(k-1)\delta, k\delta)$ as the k -th δ -interval, $1 \leq k \leq \kappa$. Values of δ and κ are to be determined.

Also, we classify the jobs into *big* and *small* jobs. The *big* jobs are those with

processing time $\geq \gamma$, and *small* jobs are those with processing time $< \gamma$. The value of γ will be specified later. Then we are ready to formally define the *outline scheme*.

Each *outline* consists of:

- The δ -interval in which each operation of a *big* job begins.
- For each machine and δ -interval, the approximate (rounded up to the nearest multiple of γ) amount of time allocated to the operations associated with *small* jobs that begin in that δ -interval.

Therefore, the *outline* specifies which δ -interval each operation of each *big* job should begin in, and how much small-jobs-time is allocated for each δ -interval on each machine. The reason that we label jobs as *big* and *small* is because we cannot afford to guess too much detail on every job. Instead, for the *small* jobs, whose order of scheduling do not impact the overall length significantly, we can schedule them approximately by using an LP. How many *outlines* do we need to guess? Suppose the number of *big* jobs is L , then the number of possible assignments of *big*-job operations to δ -intervals is at most κ^{mL} . And observe that the number of possible assignments of small-jobs-time to intervals is at most $(\frac{\delta}{\gamma} + 1)^{m\kappa}$. Therefore, the number of outlines is bounded by:

$$\kappa^{mL}(\delta/\gamma + 1)^{m\kappa}.$$

For a given outline, we introduce a Linear Programming to determine the assignment of small-job operations to δ -intervals. Let $J_1, J_2, \dots, J_{n'}$ be the small jobs, and job J_i is to be processed on machines $M_{i_1}, \dots, M_{i_{\lambda_i}}$ (recall $\lambda_i \leq m$ is the span of job J_i) in the specified order. Then we construct an LP with the following variables:

$$x_{j,(t_1,t_2,\dots,t_{\lambda_j})}, j = 1, \dots, n', 1 \leq t_1 \leq t_2 \leq \dots \leq t_{\lambda_j} \leq \kappa$$

$$\begin{aligned}
& \sum_{t_1, t_2, \dots, t_{\lambda_j}} x_{j, (t_1, t_2, \dots, t_{\lambda_j})} = 1, j = 1, \dots, n', \\
& \sum_{\{j | M_1 \in P_j\}} p_j x_{j, (\dots, k, \dots)} \leq \alpha_1^k, k = 1, \dots, \kappa, \\
& \sum_{\{j | M_2 \in P_j\}} p_j x_{j, (\dots, k, \dots)} \leq \alpha_2^k, k = 1, \dots, \kappa, \\
& \vdots \\
& \sum_{\{j | M_m \in P_j\}} p_j x_{j, (\dots, k, \dots)} \leq \alpha_m^k, k = 1, \dots, \kappa, \\
& x \geq 0.
\end{aligned}$$

Figure 3.1: the LP to assign small jobs. Recall that p_j is the processing time of job j , and P_j is the path of machines for job j .

where $x_{j, (t_1, t_2, \dots, t_{\lambda_j})} = 1$ means that job J_j is assigned to δ -interval t_1 on machine M_{j_1} , t_2 on machine M_{j_2} , and so on. We use $\alpha_1^k, \alpha_2^k, \dots, \alpha_m^k$ to denote the amount of time (for *small* jobs) assigned (by *outline*) to the k -th δ -interval on machines M_1, M_2, \dots, M_m , respectively. We want to find a basic feasible solution against the constraints in Fig. 3.1.

The first and last constraints ensure that the operations of all *small* jobs are assigned to some δ -intervals, and all constraints in the middle ensure that the small-job-time in the solution in each interval on each machine is no more than the value described by the outline.

Observe that the LP has $n' + m\kappa$ constraints and at most $n'\kappa^m$ variables. A basic feasible solution of this LP is guaranteed to have at most $n' + m\kappa$ positive variables. Also, each job must have at least one positive variable associated with it. This is because of the first constraint of the LP. Thus, A job that receives fractional assignment must have at least one more positive variable. Combining with the fact that the bfs has at most $n' + m\kappa$ positive variables, we know that such a solution can have at most $m\kappa$ jobs that actually receive fractional assignments and the remaining

small jobs will have unique integral assignment to δ -intervals. Let's just ignore the small jobs that received fractional assignments. They will be appended to the end of the schedule with a cost of at most $(m\kappa + m - 1)\gamma$.

For the remaining jobs (*big* jobs + *small* jobs with integral assignments), we describe a two-step algorithm to construct a schedule based on their assignments to δ -intervals.

In the first step, we 'greedily' schedule each machine independently. For a machine M_i , we order the operations assigned to each δ -interval such that the longest operation is the last. More precisely, let I be the set of indices k such that there are some operations assigned to the k th δ -interval in the first step schedule. Then we schedule the operations in the order of their indices in I where operations in k th δ -interval start at time σ_k and end at time τ_k , where

$$\sigma_k = \max\{(k-1)(\delta + \gamma), \max_{1 \leq h \leq k-1, h \in I} \{\tau_h\}\},$$

τ_k becomes well-defined once we defined σ_k . Another way to view the first step schedule is that: for machine M_i , all operations assigned in the first δ -interval get scheduled first as a block with no idle time in between, followed by the operations in the second δ -interval, and so on. Within each block, we schedule the largest operation the last. Therefore, jobs in each block do not overlap, and they are not scheduled before the specified starting time.

Let Σ be the optimal schedule in a fixed the outline, and let \tilde{T} be its length. We focus on a specific machine M_i . Let s_k and t_k denote the start time and (end time) of the first and (last) operations during the k th δ -interval in Σ . Then:

Lemma 7 *for all k , $\sigma_k \leq s_k + (k-1)\gamma$, and $\tau_k < t_k + k\gamma$*

Proof. First observe that for any $k \in I$, $\tau_k - \sigma_k < t_k - s_k + \gamma$. This is because jobs in each block get scheduled with no idle time in the first step schedule, and the additional γ comes from rounding up the small-job-time to the nearest multiple of γ .

We prove the first inequality in the lemma by induction. With the smallest $k \in I$, we have $\sigma_k = (k - 1)(\delta + \gamma)$, which is at most $s_k + (k - 1)\gamma$. Then, inductively:

$$\begin{aligned}
\sigma_k &= \max\{(k - 1)(\delta + \gamma), \max_{1 \leq h \leq k-1, h \in I} \{\tau_h\}\} \\
&= \max\{(k - 1)(\delta + \gamma), \max_{1 \leq h \leq k-1, h \in I} \{\sigma_h + (\tau_h - \sigma_h)\}\} \\
&\leq \max\{(k - 1)(\delta + \gamma), \max_{1 \leq h \leq k-1, h \in I} \{s_h + (h - 1)\gamma + (t_h - s_h + \gamma)\}\} \\
&\leq (k - 1)\gamma + \max\{(k - 1)\delta, \max_{1 \leq h \leq k-1, h \in I} \{s_h + (t_h - s_h)\}\} \\
&\leq s_k + (k - 1)\gamma
\end{aligned}$$

Since $\tau_k - \sigma_k < t_k - s_k + \gamma$, we have $\tau_k < t_k + k\gamma$, which completes the proof. ■

Corollary 8 *Lemma 7 implies that the makespan of the first step schedule is at most $\tilde{T} + \kappa\gamma$.*

However, notice that the first-step schedule is very likely an infeasible schedule, because we only focus on each machine individually, so the operations of a job might get processed on different machines at the same time (overlaps). The second step of the algorithm is to remove the potential overlaps of operations by delaying the operations on M_i by $2(i - 1)(\delta + \gamma)$ units of time, for $i = 2, \dots, m$. We will eventually show that the schedule after injecting delays on every machine will be feasible, but before that we need to prove the following lemma first:

Lemma 9 *Consider operations inside an arbitrary k th δ -interval on a arbitrary machine in the first step schedule. (1) Each large operation starts processing during $[(k - 1)(\delta + \gamma), k(\delta + \gamma))$. (2) And each small operation starts and finishes processing during $[(k - 1)(\delta + \gamma), k(\delta + \gamma) + \gamma)$.*

Proof. It is clear that all these operations start at or after $(k - 1)(\delta + \gamma)$. Then it remains to show that they don't start (end) too late. We consider the following

two cases. First, suppose there exists an large operation. Observe that all large operations are scheduled in the same δ -interval in Σ (recall Σ is the optimal schedule in the *outline*) as well. Let O_j be the operation that was scheduled the last in our algorithm (it is the largest), then it suffices to show O_j starts before time $k(\delta + \gamma)$. From lemma 7, we know $\tau_k < t_k + k\gamma$, also there is some operation $O_{j'}(p_j \geq p_{j'})$ is scheduled the last and completes at t_k in Σ . Also, $O_{j'}$ starts before time $k\delta$. Therefore, the last operation of this interval starts at time $\tau_k - p_j < t_k + k\gamma - p_{j'} < k(\delta + \gamma)$.

The other case is when all operations are small ($< \gamma$). In this case, $t_k < k\delta + \gamma$, so by lemma 7: $\tau_k < t_k + k\gamma < k(\delta + \gamma) + \gamma$. This completes the proof. ■

Now we are ready to prove that the resulting schedule after step 2 is feasible.

Lemma 10 *After delaying the operations on M_i by $2(i - 1)(\delta + \gamma)$ units, for $i = 2, \dots, m$, the resulting schedule is feasible.*

Proof. The schedule we obtained from first step is conflict-free, but it is still likely infeasible because there might be a job starting on a machine before its previous operation finishes on the previous machine (the job starts before it becomes available). In step two, we delay operations in M_2 by $2(\delta + \gamma)$, jobs in M_3 by $4(\delta + \gamma)$, and so on. So the makespan of the schedule increases by at most $2(m - 1)(\delta + \gamma)$. And we show that the schedule after injecting delays is feasible.

Consider an arbitrary job j , and two consecutive operations of j on machine M_i and M_{i+1} , call them $O_{j,i}$, and $O_{j,i+1}$. It suffices to prove that these two operations are scheduled in order and do not overlap.

First consider the case when j is a *big* job. Suppose $O_{j,i}$ is assigned to k th δ -interval and $O_{j,i+1}$ is assigned to the l th δ -interval ($l \geq k$). By Lemma 9, the difference of their starting time is at least $(l - k - 1)(\delta + \gamma)$, i.e. in the worst case $O_{j,i+1}$ starts on M_{i+1} $(\delta + \gamma)$ units before $O_{j,i}$ starts on M_i in the first step schedule. Note that operations on M_{i+1} are delayed by $2(\delta + \gamma)$ more units relative to operations on M_i

in step two, so once the delays have been injected, $O_{j,i}$ and $O_{j,i+1}$ will be scheduled in order and do not overlap.

Another case is when j is a small job, and we still use $O_{j,i}$ and $O_{j,i+1}$ to denote the two consecutive operations of j . And suppose $O_{j,i}$ is assigned to k th δ -interval and $O_{j,i+1}$ is assigned to the l th δ -interval ($l \geq k$). Again by Lemma 9, after the delays have been injected, $O_{j,i}$ will complete before $(k + 2(i - 1))(\delta + \gamma) + \gamma$, and $O_{j,i+1}$ will start at or after time $(2i + k - 1)(\delta + \gamma)$. Since $(2i + k - 1)(\delta + \gamma) > (k + 2(i - 1))(\delta + \gamma) + \gamma$, $O_{j,i}$ and $O_{j,i+1}$ will be scheduled in order and do not overlap. ■

Combining previous lemmas, we obtain the following theorem:

Theorem 11 *For a given δ , and γ and an outline with an associated optimal schedule of length \tilde{T} , we can generate a feasible schedule of length:*

$$\tilde{T} + \kappa\gamma + 2(m - 1)(\gamma + \delta) + (m\kappa + m - 1)\gamma$$

Proof. The additive $\kappa\gamma$ follows from Corollary 8, second term $2(m - 1)(\gamma + \delta)$ follows from Lemma 10, and $(m\kappa + m - 1)\gamma$ comes from the fractional *small* jobs that get appended at the end. ■

3.1.2 The PTAS

Now, we are ready to complete the proof of Theorem 5. Let $\delta = \frac{T\epsilon}{u}$, and $\gamma = \frac{T\epsilon^2}{uv}$. The value of u and v will be specified later. So the additive error from Theorem 11 becomes:

$$\begin{aligned} & \kappa\gamma + 2(m - 1)(\gamma + \delta) + (m\kappa + m - 1)\gamma \\ &= \frac{u}{\epsilon} \cdot \frac{T\epsilon^2}{uv} + 2(m - 1)\left(\frac{T\epsilon(v + \epsilon)}{uv}\right) + \left(\frac{mu}{\epsilon} + m - 1\right)\frac{T\epsilon^2}{uv} \\ &= \left[\frac{3m\epsilon + 2mv + (m + 1)u - 3\epsilon - 2v}{uv} \right] \epsilon T \end{aligned}$$

Moreover, recall that T is the upper bound of \tilde{T} , suppose $T = \beta\tilde{T}$, such β is at most m . Let L be the number of large jobs, then L is at most $\frac{m\tilde{T}}{\gamma} = \frac{muv}{\beta\epsilon^2}$, therefore the total possible number of assignments of large operations to δ -intervals is at most κ^{mL} .

The number of small-job-time that we assign to each δ -interval is $\frac{\delta}{\gamma} + 1 = \frac{v}{\epsilon} + 1$ (because we round it up to multiple of γ), so the number of possible assignments of small-job-time to each machine during each interval is at most $(\frac{v}{\epsilon} + 1)^{m\kappa}$. Therefore, the total number of outlines is at most:

$$\left(\frac{u}{\epsilon}\right)^{m^2uv/\epsilon^2} \left(\frac{v}{\epsilon} + 1\right)^{mu/\epsilon}$$

For sufficiently large u, v , for example, $u = 4(m-1)\beta$ and $v = 2(m+1)\beta + (\frac{3}{2})\epsilon$, then the additive error becomes:

$$\begin{aligned} \text{Error} &= \left[\frac{3m\epsilon + 4m(m+1)\beta + 3m\epsilon + 4(m+1)(m-1)\beta - 3\epsilon - 4(m+1)\beta - 3\epsilon}{8(m+1)(m-1)\beta^2 + 6(m-1)\beta\epsilon} \right] \cdot \epsilon\beta\tilde{T} \\ &= \left[\frac{8(m+1)(m-1)\beta + 6(m-1)\epsilon}{8(m+1)(m-1)\beta^2 + 6(m-1)\epsilon\beta} \right] \cdot \epsilon\beta\tilde{T} \\ &= \epsilon\tilde{T} \end{aligned}$$

Therefore, we can guarantee the additive error is at most $\epsilon\tilde{T}$

Runtime: The runtime is given by the number of *outlines* that we need to consider multiplied by the time needed to solve an individual *outline*. First note that the number of *outlines* is $O(m/\epsilon)^{O(m^6/\epsilon^2)}$. For each *outline*, we need to find a basic feasible solution for the associated LP; recall that the LP has $n + \frac{mu}{\epsilon}$ constraints and at most $\frac{nu^m}{\epsilon^m}$ variables. By using the LP solver from [24], we can solve the LP in time $O(N^{3.5})$, where N is the input size. Therefore the total runtime is $O(N^{3.5}(m/\epsilon)^{O(m^6/\epsilon^2)})$. Suppose $m = O(\frac{\log^{1/6} n}{\log \log n})$, the total runtime becomes:

$$\begin{aligned}
O(N^{3.5}m^{(m^6)}) &= O\left(N^{3.5}\left(\frac{\log^{1/6}n}{\log\log n}\right)^{\log n/\log^6\log n}\right) \\
&= O\left(N^{3.5}2^{\frac{\log n}{\log^6\log n}(\log\log^{1/6}n)}\right) \\
&= O(N^{3.5}n)
\end{aligned}$$

Therefore, the total runtime is polynomially bounded.

Observation: If we define *big* jobs to be those with maximum processing time (over their spans of machines) at least γ , and *small* jobs are those with maximum processing time $< \gamma$, then the PTAS actually holds even on unrelated machines setting (when m is sublogarithmic), which generalizes the result of Hall [11] for the classical problem of flow shop.

3.2 A $O\left(\frac{\log m}{\log\log m}\right)$ -approximation for makespan objective

In this section, we prove

Theorem 12 *For GPS with makespan objective, there is an $O\left(\frac{\log m}{\log\log m}\right)$ -approximation.*

Consider the special case of the GPS problem where there are h machines, called *terminal* machines, such that any job j has to start/end at one of the h machines. This special case by itself is quite interesting because if one can have a good approximation for this special case, then the algorithm for this special case can be adapted to solve the general problem, as follows.

For a given instance, we select h machines that equally partition the machines into h segments (as in Fig. 3.2). For all jobs whose span crosses these h machines (i.e. uses machines in two segments), we group them into group G_1 . So all the jobs in



Figure 3.2: In this example, $h = 3$, the dark squares represent the terminal machines in the first layer, which partition the machines into 3 segments. The squares with crossed lines represent the second layer terminal machines. The corresponding groups of jobs are defined as $G_1 = \{j_1, j_2\}$ and $G_2 = \{j_3, j_4, j_5\}$.

$J - G_1$ have their spans entirely within one segment and those that fall into different segments, can be scheduled independently (as their paths do not overlap). For each segment, we again select h machines (to partition that segment into equal size sub-segments) and all jobs in $J - G_1$ that pass any one of the secondary terminal machines (h^2 many) form group G_2 . And we do this recursively. Eventually, we have $O(\log_h m)$ groups. Also, by losing a constant factor, we can assume that jobs among a group have to start/finish at one of the h terminal machines. This is because the instance where there is a machine that is used by all jobs is a special case of the *junction tree* problem studied in [5]. Therefore we can use their two-stage algorithm and in $\leq 2OPT$ time send all jobs to their first terminal machines; once all jobs reach their last terminal machines on their paths, it takes at most $2OPT$ time to deliver them to the final destinations. Suppose we have an ρ -approximation for a single group of jobs, then if we schedule all groups sequentially, we obtain an $O(\log_h m \cdot \rho)$ -approximation for the general case. Therefore, the problem becomes how large can we push the value of h to? and what's the best value for ρ we can get for an individual group?

3.2.1 Instances with h terminal machines

In this section, we show how one can extend the idea of the PTAS in Section 3.1 to solve the instances with h terminal machines. Similarly, suppose we have an upper bound T on the *makespan* of the optimal schedule, and we partition the time line from 0 to T into κ intervals of size $\delta = \frac{T}{\kappa}$.

The definition of a job being *big* or *small* is slightly different. Suppose the h terminal machines partition the machines into h equal-size segments (except the last

one). A job j is *big* if the time it takes to travel a segment is $\geq \gamma$ (i.e, $p_j \times \text{segment size} \geq \gamma$). Otherwise, we say the job is *small*.

Each *outline* should specify:

- The δ interval in which a *big* job starts running on a terminal machine.
- For each terminal machine and δ -interval, how much time is allocated to small jobs that begins in that δ -interval, rounded up to the nearest multiple of γ .

Suppose the number of *big* jobs is L , the number of guesses of the starting time interval of all *big* jobs is at most κ^{hL} ; also, the number of possible assignments of small-job-time to δ -intervals is at most $(\delta/\gamma + 1)^{h\kappa}$. So the number of outlines is at most:

$$\kappa^{hL}(\delta/\gamma + 1)^{h\kappa}$$

For the small jobs, we again construct an LP as in Section 3.1 and find a basic feasible solution of it. Then we have at most $h\kappa$ jobs that actually receive fractional assignments, we can ignore them for now and append them at the end of the schedule with a cost at most $(h\kappa + h - 1)\gamma$.

For all *big* jobs and *small* jobs with integral assignments, we schedule them according to their assignments to δ -intervals in two steps. In the first step, we schedule each segment independently. For a fixed segment, let M_i be the first machine (a terminal machine) of this segment. We order the jobs assigned to each δ -interval according to their processing times and send them based on *faster first*. Let σ_k be the time when the first job in the k th δ -interval begins on the first machine of the segment, and let τ_k be the time when the last job in the k th δ -interval finishes on the last machine of the segment.

Let Σ be the optimal schedule in the *outline* that we are focusing on, say the makespan of Σ is \tilde{T} . Similarly, we define s_k and (t_k) to be the start time (end time)

of the first (last) job during the k th δ -interval of Σ on the same segment. Then we show the following:

Lemma 13 *For all k , $\sigma_k \leq s_k + (k - 1)\delta$, and $\tau_k < t_k + k\gamma$*

Proof. One important observation is that $\forall k, \tau_k - \sigma_k < t_k - s_k + \gamma$. This is given by the fact that *faster first* algorithm gives optimal solution for flowshop with identical machines. That is, if we consider the jobs that traverse a specific segment during a specific δ -interval, view it as an instance of flowshop with identical machines. The schedule defined by *faster first* is no worse than the optimal schedule. The additive γ comes from the rounding error. The rest of proof is analogous to the proof of Lemma 7, hence omitted here.

■

Then we can conclude the *makespan* of the schedule obtained from the first step is at most $\tilde{T} + \kappa\gamma$. However, this schedule is likely to be infeasible because we schedule each segment independently without caring about their consistency. The second step is to inject delays to jobs so that the jobs in the resulting schedule are processed in order. We delay operations on the second terminal machine by $2(\delta + \gamma)$, delay operations on the third terminal machine by $4(\delta + \gamma)$, and so on. So eventually, if we consider two adjacent terminal machines, operations on the later one are delayed by $2(\delta + \gamma)$ units relative to the previous terminal machine. And we need to show:

Lemma 14 *After delaying the operations on the i th terminal machine M_{h_i} by $2(i - 1)(\delta + \gamma)$ units of time, for $i = 2, \dots, h$. The resulting schedule is feasible.*

Proof. Consider an arbitrary job j and its operations traveling two adjacent terminal machines M_{h_i} and $M_{h_{i+1}}$. Observe that all operations of job j in the segment starting with machine M_{h_i} must be scheduled in order because we schedule them based on *faster first*. So it remains to show that after injecting the delays, j doesn't start

on $M_{h_{i+1}}$ before all previous operations are finished. The rest of proof is analogous to the proof of Lemma 10.

■

Therefore, the final schedule is of length at most $\tilde{T} + \kappa\gamma + 2(h-1)(\delta + \gamma) + (h\kappa + h-1)\gamma$. For sufficiently small $\delta = O(\frac{T\epsilon}{h})$ and $\gamma = O(\frac{T\epsilon^2}{h^2})$, the additive error becomes $\epsilon\tilde{T}$. Moreover, the total number of *outlines* is $O(h^{(h^6)})$. Suppose h is sub-logarithmic, say $h = \frac{\log^{1/6} m}{\log \log m}$, then the runtime becomes polynomially bounded. This implies a $O(\log_h m \cdot \rho) = O(\frac{\log m}{\log \log m})$ -approximation for the general problem, which completes the proof of Theorem 12.

3.3 A PTAS for min-sum GPS when m is sub-logarithmic

Recall that the completion time C_j of a job j is the time when the last operation of j finishes processing. The *makespan* objective aims at minimizing the completion time of the whole system. On the other hand, the *min-sum* objective aims at minimizing the average completion time over all jobs.

In this section, we study the approximability of the *min-sum* objective. The idea of designing approximation algorithms for the min-sum objective by using algorithms for the min-max (makespan) variants have been used extensively for various problems such as scheduling and vehicle routing problems (to name a few see [2, 5, 18, 22]). Here we borrow ideas from [22], which designs a PTAS for the minimum latency traveling repairman problem on Euclidean metrics by reducing it to a variant of min-max version of it. This technique is used to design algorithms for many other problems, see [6, 23] for an example. First in subsection 3.3.1, we introduce a variant of the GPS problem, called the *segmented* GPS problem and we give a PTAS for it when $m = O(\log^{1/6} n / \log \log n)$. Then in subsection 3.3.2, we use it as a subroutine to prove the following:

Theorem 15 *There is a PTAS for GPS with min-sum objective when $m = O\left(\frac{\log^{1/6} n}{\log \log n}\right)$.*

3.3.1 Segmented GPS

First we define an interesting variant of the GPS problem called the *segmented* GPS as follows.

Definition 16 (segmented GPS) *An instance of segmented GPS is given by a set of m identical machines that form a path, and also a set of n jobs each needs to be processed on a sub-path. Also, for some constant π , given bounds $B_1 \leq B_2 \leq \dots \leq B_\pi$ such that $B_i/B_{i-1} = \eta$ where η is a constant, and given numbers $n_1 \leq n_2 \leq \dots \leq n_\pi = n$. A feasible solution is a schedule such that at least n_i jobs are finished within the first B_i units of time for all $i \in \{1, \dots, \pi\}$, and the length of the schedule is at most B_π . We say an algorithm gives an α -approximation if for any feasible instance it finds a schedule that finishes at least n_i jobs within αB_i units of times.*

In this subsection, we focus on approximating the *segmented* GPS problem, the consequences of it will be discussed later in next subsection.

Theorem 17 *There is an $(1 + \epsilon)$ approximation for the segmented GPS when $m = O\left(\frac{\log^{1/6} n}{\log \log n}\right)$.*

The algorithm also adapt the idea of *outline scheme*. Intuitively, we again partition the time line from 0 to B_π into polynomially many δ -intervals, and we use the notion of *big* and *small* to classify jobs so that we can afford to fully guess the assignment of *big* jobs to δ -intervals. For *small* jobs, we *guess* approximately the amount of time that is allocated for them on each δ -interval and each machine, and we then assign *small* jobs by an LP.

However, we define δ w.r.t. B_1 instead of B_π . This gives us better precision so that the additive error at the end depends on B_1 , so is relatively small. Also, at the same time, the number of δ intervals doesn't blow up, the number is at most η^π (which is

a constant since both η, π are constants) times what we used to have. Moreover, for the same reason, we define γ w.r.t. B_1 . The number of large jobs L is also η^π (which is a constant) times what we used to have.




Figure 3.3: Partiton the time line into intervals of size δ , which is defined w.r.t. B_1 .

More precisely, Let $\delta = \frac{B_1 \epsilon}{u}$ and $\gamma = \frac{B_1 \epsilon^2}{uv}$, as before, $u = O(m)$ and $v = O(m)$ are linear functions of m to be specified later. The number of δ -intervals is $\kappa = \frac{\eta^\pi B_1}{\delta} = \frac{u \eta^\pi}{\epsilon}$. An *outline* specifies the δ -interval in which an operation of a *big* jobs begins in, and amount of time that is allocated to *small* job in each interval on each machine, rounded to nearest multiple of γ . Therefore, the total number of *outlines* is $\kappa^{mL} (\frac{\delta}{\gamma} + 1)^{m\kappa} = O(m^{(m^6)})$, which is polynomially bounded as $m = O(\frac{\log^{1/6} n}{\log \log n})$.

Assume we know the assignment of *big* jobs. As before, we process the operations assigned to each δ -intervals such that the longest operation is the last. Then, for each bound B_i , we know the number of large jobs n_i^l that are finished before B_i . Therefore, when we assign *small* jobs, we modify the LP in Figure. 3.1 by adding π extra constraints to ensure that x values that fall in the first B_i units of time is at least $n_i^s = n_i - n_i^l$ (see Figure 3.4 for the full LP).

Such LP has $n' + m\kappa + \pi$ constraints and at most $n' \kappa^m$ variables (recall n' is the number of small jobs). A basic feasible solution of this LP is guaranteed to have at most $m\kappa + \pi$ small jobs that actually receive fractional assignments and the remaining $n_i - m\kappa - \pi$ small jobs will have unique integral assignment to δ -intervals. We can simply ignore the fractional small jobs for now, because we can append them all at the end of B_1 with a cost of at most $(m\kappa + \pi + m - 1)\gamma$ (which is at most ϵB_1 , based on the discussion in Section 3.1.2). Call this schedule the first-step schedule, then it finishes at least n_i jobs before time B_i , for $i = 1, \dots, \pi$, as wanted. But it may not be feasible. In order

$$\begin{aligned}
& \sum_{t_1, t_2, \dots, t_{\lambda_j}} x_{j, (t_1, t_2, \dots, t_{\lambda_j})} = 1, j = 1, \dots, n', \\
& \sum_{\{j | M_1 \in P_j\}} p_j x_{j, (\dots, k, \dots)} \leq \alpha_1^k, k = 1, \dots, \kappa, \\
& \sum_{\{j | M_2 \in P_j\}} p_j x_{j, (\dots, k, \dots)} \leq \alpha_2^k, k = 1, \dots, \kappa, \\
& \vdots \\
& \sum_{\{j | M_m \in P_j\}} p_j x_{j, (\dots, k, \dots)} \leq \alpha_m^k, k = 1, \dots, \kappa, \\
& \sum_{\{j | t_{\lambda_j} \cdot \delta \leq B_i\}} x_{j, (\dots, t_{\lambda_j})} \geq n_i^s, i = 1, \dots, \pi, \\
& x \geq 0.
\end{aligned}$$

Figure 3.4: the modified LP. π new constraints are added to ensure that at least $n_i^s = n_i - n_i^l$ small jobs are finished before time B_i .

to turn it into a feasible schedule, we need to inject delays to machines.

Lemma 18 *After delaying the operations on machine M_i by $2(i-1)(\delta + \gamma)$ units, for $i = 2, \dots, m$. The resulting schedule becomes feasible. And the delays only stretch the schedule by a factor of $(1 + \epsilon)$.*

Proof. The proof is analogous to the proof of Lemma 7 and Theorem 11. ■

Theorem 17 follows immediately from the above discussion.

3.3.2 The PTAS

The main theorem that we prove in this subsection is:

Theorem 19 *If there is a polynomial time α -approximation algorithm for the segmented GPS problem, then there is a polynomial time $(1 + \epsilon)\alpha$ -approximation algorithm for the GPS min-sum minimization problem.*

Then Theorem 15 follows immediately from Theorems 17 and 19.

Proof of Theorem 19 is built upon ideas of [22] for minimum latency traveling repairman problem on Euclidean metrics. First we show that:

Lemma 20 *With a $(1 + \epsilon)$ -factor loss, we may assume that the makespan of the optimal schedule is polynomially bounded in m, n .*

Proof. Let p_{max}, p_{min} be the largest and smallest processing times, respectively. One can assume that $p_{min} \geq \epsilon p_{max}/(mn)$, otherwise all jobs with processing times smaller than $\epsilon p_{max}/(mn)$ can be removed, then they can be added to any schedule of the rest of jobs right before a job of size p_{max} and this will increase the total completion time of the schedule by at most a $(1 + \epsilon)$ factor. With this assumption, we can scale processing times so that $p_{min} = 1$ and $p_{max} \leq mn/\epsilon$. ■

Now, we are ready to present the reduction. Consider the time points $t_1, t_2, \dots, t_\Gamma$, where $t_1 = 1$ and $t_i/t_{i-1} = (1 + \epsilon)^\pi$, for some constant π that only depends on ϵ . We can assume $\Gamma = O(\log mn)$ by the previous lemma. The part of schedule between t_i and t_{i+1} is called the i th *subschedule*. We call a schedule is *well-structured* if each *subschedule* processes a subset of jobs completely. That is, if a job j starts processing at or after time t_i , it has to be finished on all its span before t_{i+1} . We first show that we can reduce the solution space to only the *well-structured* schedules by losing an ϵ -factor. This allows us to deal with each *subschedule* independently. Moreover, the time frame between t_i and t_{i+1} can be further partitioned into π sub-intervals such that the ratio of the end time and start time of each sub-interval is $(1 + \epsilon)$. Therefore, each *subschedule* can be viewed as an instance of the *segmented GPS* problem. However, we cannot afford to guess the subset of jobs to be processed on every *subschedule*, but we show that, for large enough π , in the i th *subschedule* we can simply re-do all the jobs that have been processed in the previous *subschedules*. As a result, we don't need to know the set of jobs to be processed on each *subschedule*, instead, we use Dynamic Programming to enumerate the number of jobs to be processed, which can be done in polynomial time. The first step is the following lemma.

Lemma 21 *There is a $(1 + \epsilon)$ -approximate well-structured schedule OPT' .*

Proof. We prove this by showing that we can modify an optimal schedule to satisfy the desired property and the total completion time of the modified schedule only increases by an ϵ factor. Assume $0 < \epsilon \leq 1$ and $\pi = O(1/\epsilon^2)$ is a constant depending on ϵ only. Let h be a random offset chosen u.a.r. in $\{0, 1, \dots, \pi - 1\}$. Let

$$A_i = (1 + \epsilon)^{(i-1)\pi+h}, \text{ for } i \geq 0.$$

Consider an optimal schedule OPT (use opt to denote its value), and let OPT_i be the partial schedule restricted to the jobs that finish no later than A_i , for $i = 1, \dots, \Gamma = O(\log mn)$. The modified schedule OPT' with value opt' is constructed by concatenating partial schedules OPT_i sequentially such that OPT_i starts at time $t_i = cA_{i-1}$, for $i = 1, \dots, \Gamma$ and some constant c . Observe that a job j that appears in OPT_i will also appear in $OPT_{i'}$ for $i' > i$, the completion time of a job is defined by its first appearance.

We first prove that OPT' is feasible. For that, we need to show $t_i \geq t_{i-1} + A_{i-1}$ for all i , i.e. $c/(c-1) \leq A_{i-1}/A_{i-2} = (1 + \epsilon)^\pi$. For this inequality to hold, it's enough to take $\pi \geq \frac{\log c - \log(c-1)}{\log(1+\epsilon)} = O(\frac{1}{\epsilon})$.

Next, let C_j and C'_j be the completion time of job j in OPT and OPT' , respectively. In order to show $\mathbb{E}[opt'] \leq (1 + \epsilon)opt$, it's sufficient to show $\mathbb{E}[C'_j] \leq (1 + \epsilon)C_j$, for any job j .

For an arbitrary job j , let i' be the first index such that $A_{i'} \geq C_j$. Thus, job j is first processed by $OPT_{i'}$ in OPT' . Let $(1 + \epsilon)^{q-1} < C_j \leq (1 + \epsilon)^q$ for some integer $q \geq 0$. Then the expected value of $A_{i'}$ over all possible values that h can take is:

$$\mathbb{E}[A_{i'}] = \frac{1}{\pi} \sum_{h=0}^{\pi-1} (1 + \epsilon)^{q+h} = \frac{(1 + \epsilon)^{\pi+q-1} - (1 + \epsilon)^q}{\pi\epsilon} < \frac{(1 + \epsilon)^{\pi+q}}{\pi\epsilon}$$

Observe that the completion time of job j in OPT is C_j , and j is scheduled by $\text{OPT}_{i'}$ in OPT' , so the completion time j in OPT' is $C'_j = t_{i'} + C_j$, i.e. $t_{i'} = C'_j - C_j$. Also, recall that $t_{i'} = cA_{i'-1} = \frac{cA_{i'}}{(1+\epsilon)^\pi}$, therefore,

$$\begin{aligned} \mathbb{E}[C'_j] - C_j &= \mathbb{E}[t_{i'}] = \frac{c}{(1+\epsilon)^\pi} \mathbb{E}[A_{i'}] \\ &< \frac{c}{(1+\epsilon)^\pi} \cdot \frac{(1+\epsilon)^{\pi+q}}{\pi\epsilon} \\ &= \frac{c(1+\epsilon)^q}{\pi\epsilon} \\ &< \frac{c(1+\epsilon)}{\pi\epsilon} C_j. \end{aligned}$$

The error is within an ϵ factor if $\pi \geq \frac{c(1+\epsilon)}{\epsilon^2} = O\left(\frac{1}{\epsilon^2}\right)$, as desired. We can assume the inequalities still hold without expectation by trying all values of h . ■

So using Lemma 21 we can focus on *well-structured* solutions. The proof of the lemma shows that solution OPT' is ① *well-structured*; ② each *subschedule* processes all jobs that appear in previous *subschedules*. Therefore, let D_j be the completion time of j th job on the first *subschedule* that processes at least j jobs in OPT' , we know that $\sum_{j=1}^n D_j = \sum_{j \in J} C'_j \leq (1+\epsilon)\text{opt}$. Therefore, in order to find such an OPT' , we can search for a solution among all schedules satisfying ① and ② that minimizes $\sum_{j=1}^n D_j$. We show this can be done by a DP that runs in polynomial time if we have an algorithm for the following subproblem.

Definition 22 (The subproblem) *An instance of the subproblem is given by $i \in \{1, \dots, \Gamma\}$ and integers $n' \leq n'' \in \{0, 1, \dots, n\}$. A solution is a schedule that starts at time t_i and finishes before time t_{i+1} that processes exactly n'' jobs. The goal is to find a schedule that minimizes the total completion time of jobs $n' + 1, \dots, n''$. For any feasible instance (i, n', n'') , let $\text{SUB}_i(n', n'')$ denote its optimal value.*

We say an algorithm is (α, β) -approximation for the subproblem if for any feasible instance (i, n', n'') , it find a schedule that starts at time αt_i and finishes before time αt_{i+1} , and the total completion time of jobs $n' + 1, \dots, n''$ is at most $\alpha\beta \text{SUB}_i(n', n'')$.

Lemma 23 *If there is an (α, β) -approximation for the subproblem, then there is an $\alpha\beta(1 + \epsilon)$ -approximation for the GPS min-sum minimization problem.*

Proof. Suppose there is an (α, β) -approximation algorithm ALG. And let $\text{ALG}_i(n', n'')$ be the value returned by ALG on instance (i, n', n'') . For any sequence $0 \leq \hat{n}_1 \leq \dots \leq \hat{n}_\Gamma = n$, we have

$$\sum_{i=1}^{\Gamma} \text{ALG}_i(\hat{n}_{i-1}, \hat{n}_i) \leq \alpha\beta \sum_{i=1}^{\Gamma} \text{SUB}_i(\hat{n}_{i-1}, \hat{n}_i)$$

The solution is simply concatenating the schedules returned by ALG on $(i, \hat{n}_{i-1}, \hat{n}_i)$. The sequence of \hat{n}_i 's that minimizes the total completion time can be found by the following DP: let $\text{ALG}_i(n'') = \text{ALG}_i(0, n'')$. For all $n'' \leq n$ and for $i = 2, \dots, \Gamma$, let

$$\text{ALG}_i(n'') = \min_{n' \leq n''} \text{ALG}_{i-1}(n') + \text{ALG}_i(n', n'').$$

Therefore, the minimum is given by $\text{ALG}_\Gamma(n)$, and the DP has Γn^2 entries (Γ choices for i and $n', n'' \leq n$). Let n_1^*, \dots, n_Γ^* be the values returned by the DP that minimize $\sum_{i=1}^{\Gamma} \text{ALG}_i(\hat{n}_{i-1}, \hat{n}_i)$, and let n_i be the number of jobs finished in OPT_i . Thus, the total completion time of the solution that we find is at most:

$$\begin{aligned} \sum_{i=1}^{\Gamma} \text{ALG}_i(n_{i-1}^*, n_i^*) &\leq \sum_{i=1}^{\Gamma} \text{ALG}_i(n_{i-1}, n_i) \\ &\leq \alpha\beta \sum_{i=1}^{\Gamma} \text{SUB}_i(n_{i-1}, n_i) \\ &= \alpha\beta \text{OPT}' \\ &\leq \alpha\beta(1 + \epsilon) \text{OPT} \end{aligned}$$

■

Now, we are ready to complete the proof for Theorem 19 by showing that:

Lemma 24 *If there is an α -approximation for the segmented GPS problem, then there is an $(\alpha, 1 + \epsilon)$ -approximation for the subproblem.*

Proof. Given an instance (i, n', n'') of the subproblem. We define the time points $t_i^{(h)}$ for $h = 1, 2, \dots, \pi$, such that

$$t_i^{(h)} = (1 + \epsilon)^h t_i.$$

Naturally, the bounds in the *segmented* GPS problem are defined as $B_h = t_i^{(h)} - t_i$, and the numbers of jobs that have to be finished before every bound, n_h , can take all possible values as long as $n_1 \leq \dots \leq n_\pi = n''$. This create $O(n^\pi)$ instances of *segmented* GPS. We use the α -approximation algorithm to solve all of them and simply return the one with smallest total completion time over jobs $n' + 1, \dots, n''$.

Let SUB be an optimal solution of the subproblem instance and let $\text{SUB}_i(n', n'')$ be its value. Consider the number of jobs that are finished before B_h , $h = 1, 2, \dots, \pi$ in SUB. This is among the enumerated instances of *segmented* GPS. Let ALG be the α -approximation on the enumerated instance that is consistent with SUB, where consistent means the number of jobs finished between B_i and B_{i+1} in ALG is the same as in SUB, and denote the value of ALG by $\text{ALG}_i(n', n'')$. Let ALG start at time αt_i , because the makespan of ALG is at most $\alpha(t_{i+1} - t_i)$, so ALG finishes before time αt_{i+1} and it processes exactly n'' jobs. Therefore, it only remains to show $\text{ALG}_i(n', n'') \leq \alpha(1 + \epsilon)\text{SUB}_i(n', n'')$.

Let C_j^{Alg} and C_j^{Sub} be the completion time of the j th job in ALG and SUB, respectively. Suppose $B_{i-1} \leq C_j^{\text{Sub}} \leq B_i$, then we have $\alpha B_{i-1} \leq C_j^{\text{Alg}} \leq \alpha B_i = \alpha(1 + \epsilon)B_{i-1}$ (since $B_i/B_{i-1} = (1 + \epsilon)$). Therefore, $C_j^{\text{Alg}} \leq \alpha(1 + \epsilon)C_j^{\text{Sub}}$, for $j \in \{1, \dots, n''\}$. This implies $\text{ALG}_i(n', n'') \leq \alpha(1 + \epsilon)\text{SUB}_i(n', n'')$, as desired.

■

Theorem 19 follows from Lemmas 21, 23, and 24. Combining with the $(1 + \epsilon)$ -approximation for the *segmented* GPS problem from Section 3.3.1, we obtain a PTAS for the GPS *min-sum* minimization problem for sub-logarithmic m .

Runtime: The number of subproblems that we need to consider is $O(\Gamma n^2) = O(n^2 \log mn)$, and for each subproblem, we enumerate $O(n^\pi) = O(n^{1/\epsilon^2})$ instances of *segmented* GPS. Therefore, the total runtime is $\tilde{O}(n^2 \log mn)$ multiplied by the runtime of the algorithm that approximates the *segmented* GPS problem.

3.4 An $O\left(\frac{\log m}{\log \log m}\right)$ -approximation for min-sum GPS

In this section, we prove

Theorem 25 *For GPS with min-sum objective, there is an $O\left(\frac{\log m}{\log \log m}\right)$ -approximation.*

As discussed earlier, the framework of using a min-max solver as a blackbox to approximate a min-sum objective has been used in the past extensively. To apply that here we first define the following variant of the problem:

Definition 26 (Throughput Maximization Given Bound B) *Given an instance of the GPS problem and a bound B , what is the maximum number q of jobs that can be finished before this bound?*

An α -approximation for this problem is an algorithm that finishes q jobs within time αB .

For the ease of notation, we denote this problem as problem A. Then

Lemma 27 *If there is an α -approximation for problem A, then there is an $O(\alpha)$ -approximation for the min-sum objective.*

Proof. Given a black box that can approximate problem A within factor α , we can obtain an $O(\alpha)$ -approximation for the *min-sum* objective as follows. Let S_j be

the set of jobs that finish between time 2^j and 2^{j+1} in the optimal schedule (regarding min-sum), and let $n_j = |S_j|$. Therefore, by invoking the solver for problem A, for each j , we can find a maximum set of jobs Q_j with size q_j that can be scheduled within time $\alpha 2^{j+1}$.

Our solution to the min-sum objective is the following: for $j = 1, 2, \dots$, schedule the jobs in Q_j as suggested by the solver of problem A. Note that a job might be scheduled multiple times in different Q_j 's, the completion time of a job is the first time when it is completely scheduled. Consider the i th job that finishes in our schedule, say $i \in S_j$. Then the completion time of i th job in the optimal schedule is at least 2^j . Consider the set of jobs Q_j and note the $q_j \geq i$. Therefore the completion time of the i th job in our schedule is at most $\alpha \sum_{k=1}^{j+1} 2^k \leq \alpha 2^{j+2}$. That is, the average completion time of our schedule is at most 4α times the value of the optimal solution. This completes the proof. The constant 4 can be further optimized to 3.59, but it doesn't help the ratio asymptotically, hence details are omitted.

■

So it is enough to get an $O(\frac{\log m}{\log \log m})$ -approximation for problem A (for general m). The algorithm is similar to the one in Section 3.4, so we only provide a sketch here. First, we select $h = O(\frac{\log^{1/6} m}{\log \log m})$ terminal machines that partition the machines into h equal size segments. Group the jobs that cross the terminal machines together and do it recursively, we obtain $\Delta = O(\log_h m) = O(\frac{\log m}{\log \log m})$ classes of jobs. Let OPT be an optimal schedule that completes $q = q_1 + \dots + q_\Delta$ jobs before given bound B , where q is the maximum possible jobs that can be finished before time B and q_i is the number of jobs from class i . Consider the instance of problem A on a single class of jobs. Similarly we consider the segment between two terminal machines as the role of a single machine and define δ and γ accordingly. Then each instance can be viewed as a special case of the *segmented* GPS problem (discussed in Section 3.3.1) when $\pi = 1$, so the PTAS for the *segmented* GPS problem can be applied here. That is, given bound B and let q_i^* be the maximum number of jobs that can be finished

before B if we only consider jobs in class i , we have an algorithm that finishes q_i^* jobs before time $(1 + \epsilon)B$. Note that $q_i^* \geq q_i$. Therefore, if we apply the PTAS on every class of jobs to obtain the q_i^* many jobs from each class i and sequentially glue them together, then we get a schedule that finishes $q_1^* + \dots + q_\Delta^* \geq q_1 + \dots + q_\Delta = q$ jobs before time $(1 + \epsilon)\Delta B$, which is a $\Delta = O(\frac{\log m}{\log \log m})$ -approximation for problem A . Combining with the result from Lemma 27, we get an $O(\frac{\log m}{\log \log m})$ -approximation for *min-sum* objective (for general m), which completes the proof of Theorem 25.

Chapter 4

Conclusions

In this thesis, we studied the Generalized Path Scheduling problem, which is proved to be NP-hard for both makespan and total completion time objectives (min-sum) in Chapter 2. Then in Chapter 3, we proposed several improved approximation algorithms for it.

For the case when the number of machines m is sub-logarithmic in the number of jobs, we presented a PTAS for both objectives. The PTAS for the makespan objective was built based on the notion of outline scheme and Linear Programming. Intuitively, we label the jobs as *big* and *small* and show that the number of *big* jobs is not too large so that we can afford to enumerate the schedule of *big* jobs and do the *small* jobs by an LP. The PTAS for the total completion time objective is given by showing that with an $O(\epsilon)$ loss, we can reduce the min-sum version of the problem to a so-called segmented version, which can be solved by an algorithm similar to the PTAS for the makespan objective.

For general m , we proposed two $O(\frac{\log m}{\log \log m})$ -approximation algorithms for the GPS problem under both objectives, which improve the previous best result of [5] by a double logarithmic factor. The idea of the algorithm for makespan is that we show if one can get a ρ -approximation for the instances with h terminal machines, then there is an $O(\rho \log_h m)$ -approximation for the general instances. And we prove that we can set $h = O(\log^{1/6} m / \log \log m)$ and have $\rho = O(1)$. Moreover, we showed that how

one can turn an α -approximation algorithm for makespan to an $O(\alpha)$ -approximation algorithm for min-sum objective, which leads to an approximation algorithm for min-sum with the same asymptotic approximation ratio.

Future directions: The problem of getting an $O(1)$ -approximation algorithm for the GPS is still open for both objectives. The *furthest-to-go* algorithm seems plausible as it gives the optimal makespan for the non-nested instances [13] and we do not know any example showing that the congestion/dilation lower bound is violated by more than a small constant factor by the *furthest-to-go* algorithm. It is also worth pointing out that, for the makespan objective, if one can show every fixed priority gives $O(1)$ -approximation for instances in which all jobs have the same end machine, then *furthest-to-go* gives $O(1)$ -approximation for GPS. This is because for any machine M_i in a GPS instance, if J_i is the set of jobs that use M_i , then none of the jobs in J_i will be delayed by any job in $J - J_i$ before machine M_i . The priority among jobs in J_i is defined by their destinations. Therefore, if every fixed priority gives $O(1)$ -approximation for instances with same end machine, then the *furthest-to-go* algorithm gives $O(1)$ -approximation for the completion time of every machine, hence for the makespan of the schedule.

Another research direction is extending the current results to more general settings. For example, what if the machines are unrelated? Suppose we don't bound m , if we still use the current method of partitioning the jobs into groups and solve them separately, then each instance of h terminal machines becomes an instance of the classic flow shop problem, for which we do not know any $O(1)$ -approximation algorithm. Therefore, the extension still requires new ideas. Generalizing the current results to a more complicated network, say a tree, is also an obvious direction. However, if we look at a machine M_i , jobs that use M_i might come from multiple different machines, so extra care is needed to ensure the feasibility when designing algorithms.

Bibliography

- [1] A. Antoniadis, N. Barcelo, D. Cole, K. Fox, B. Moseley, M. Nugent, and K. Pruhs, “Packet forwarding algorithms in a line network,” *Latin American Symposium on Theoretical Informatics. Springer Berlin Heidelberg*, 2014.
- [2] K. Chaudhuri, B. Godfrey, S. Rao, and K. Talwar, “Paths, trees, and minimum latency tours,” in *44th Symposium on Foundations of Computer Science (FOCS 2003), 11-14 October 2003, Cambridge, MA, USA, Proceedings*, IEEE Computer Society, 2003, pp. 36–45. DOI: 10.1109/SFCS.2003.1238179. [Online]. Available: <https://doi.org/10.1109/SFCS.2003.1238179>.
- [3] U. Feige and C. Scheideler, “Improved bounds for acyclic job shop scheduling,” pp. 624–633, 1998.
- [4] Z. Friggstad, A. Golestanian, K. Khodamoradi, C. Martin, M. Rahgoshay, M. Rezapour, M. R. Salavatipour, and Y. Zhang, “Scheduling problems over a network of machines,” *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*, 2017.
- [5] Z. Friggstad, A. Golestanian, K. Khodamoradi, C. S. Martin, M. Rahgoshay, M. Rezapour, M. R. Salavatipour, and Y. Zhang, “Scheduling problems over a network of machines,” *J. Sched.*, vol. 22, no. 2, pp. 239–253, 2019. DOI: 10.1007/s10951-018-0591-z. [Online]. Available: <https://doi.org/10.1007/s10951-018-0591-z>.
- [6] I. Gamzu and D. Segev, “A polynomial-time approximation scheme for the airplane refueling problem,” *J. Sched.*, vol. 22, no. 1, pp. 119–135, 2019. DOI: 10.1007/s10951-018-0569-x. [Online]. Available: <https://doi.org/10.1007/s10951-018-0569-x>.
- [7] M. R. Garey, D. S. Johnson, and R. Sethi, “The complexity of flowshop and jobshop scheduling,” *Mathematics of Operations Research*, vol. 1, no. 2, pp. 117–129, 1976. [Online]. Available: <https://EconPapers.repec.org/RePEc:inm:ormoor:v:1:y:1976:i:2:p:117-129>.
- [8] M. R. Garey and D. S. Johnson, *Computers and intractability: a guide to the theory of NP-completeness*, ser. Mathematical sciences. New York, NY: Freeman, 1979. [Online]. Available: <https://cds.cern.ch/record/210237>.
- [9] L. A. Goldberg, M. Paterson, A. Srinivasan, and E. Sweedyk, “Better approximation guarantees for job-shop scheduling,” pp. 599–608, 1997.

- [10] R. Graham, E. Lawler, J. Lenstra, and A. Rinnooy Kan, “Optimization and approximation in deterministic sequencing and scheduling : A survey,” English, *Annals of Discrete Mathematics*, vol. 5, pp. 287–326, 1979, ISSN: 0167-5060. DOI: 10.1016/S0167-5060(08)70356-X.
- [11] L. A. Hall, “Approximability of flow shop scheduling,” *Math. Program.*, vol. 82, pp. 175–190, 1998. DOI: 10.1007/BF01585870. [Online]. Available: <https://doi.org/10.1007/BF01585870>.
- [12] D. G. Harris and A. Srinivasan, “Constraint satisfaction, packet routing, and the lovasz local lemma,” in *Symposium on Theory of Computing Conference, STOC’13, Palo Alto, CA, USA, June 1-4, 2013*, D. Boneh, T. Roughgarden, and J. Feigenbaum, Eds., ACM, 2013, pp. 685–694. DOI: 10.1145/2488608.2488696. [Online]. Available: <https://doi.org/10.1145/2488608.2488696>.
- [13] R. Koch, B. Peis, M. Skutella, and A. Wiese, “Real-time message routing and scheduling,” *Approximation, , and Combinatorial Optimization. Algorithms and Techniques*, 2009.
- [14] D. Kowalski, Z. Nutov, and M. Segal, ““scheduling of vehicles in transportation networks,” *International Workshop on Communication Technologies for Vehicles*, 2012.
- [15] D. R. Kowalski, E. Nussbaum, M. Segal, and V. Milyeykovski, “Scheduling problems in transportation networks of line topology,” *Optimization Letters*, vol. 8, pp. 777–799, 2014.
- [16] F. T. Leighton, B. M. Maggs, and S. B. Rao, “Packet routing and job-shop scheduling in $o(\text{congestion} + \text{dilation})$ steps,” *Combinatorica*, 1994.
- [17] T. Leighton, B. Maggs, and A. Richa, “Fast algorithms for finding $o(\text{congestion} + \text{dilation})$ packet routing schedules,” English (US), *Combinatorica*, vol. 19, no. 3, pp. 375–401, 1999, ISSN: 0209-9683. DOI: 10.1007/s004930050061.
- [18] W. Li, M. Queyranne, M. Sviridenko, and J. Yuan, “Approximation algorithms for shop scheduling problems with minsum objective: A correction,” *Journal of Scheduling*, vol. 9, pp. 569–570, 2006.
- [19] S. Sevastianov and G. Woeginger, “Makespan minimization in open shops : A polynomial time approximation scheme,” English, *Mathematical Programming*, vol. 82, no. 1-2, pp. 191–198, 1998, ISSN: 0025-5610. DOI: 10.1007/BF01585871.
- [20] N. Shakhlevich, H. Hoogeveen, and M. Pinedo, “Minimizing total weighted completion time in a proportionate flow shop,” *Journal of Scheduling*, 1998.
- [21] D. B. Shmoys, C. Stein, and J. Wein, “Improved approximation algorithms for shop scheduling problems,” 1994.

- [22] R. Sitters, “Polynomial time approximation schemes for the traveling repairman and other minimum latency problems,” in *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, C. Chekuri, Ed., SIAM, 2014, pp. 604–616. DOI: 10.1137/1.9781611973402.46. [Online]. Available: <https://doi.org/10.1137/1.9781611973402.46>.
- [23] R. Sitters and L. Yang, “A $(2 + \epsilon)$ -approximation for precedence constrained single machine scheduling with release dates and total weighted completion time objective,” *CoRR*, vol. abs/1706.07604, 2017. arXiv: 1706.07604. [Online]. Available: <http://arxiv.org/abs/1706.07604>.
- [24] P. M. Vaidya, “Speeding-up linear programming using fast matrix multiplication (extended abstract),” in *30th Annual Symposium on Foundations of Computer Science, Research Triangle Park, North Carolina, USA, 30 October - 1 November 1989*, IEEE Computer Society, 1989, pp. 332–337. DOI: 10.1109/SFCS.1989.63499. [Online]. Available: <https://doi.org/10.1109/SFCS.1989.63499>.
- [25] V. V. Vazirani, *Approximation Algorithms*. Berlin, Heidelberg: Springer-Verlag, 2001, ISBN: 3540653678.
- [26] W. Yu, H. Hoogeveen, and J. K. Lenstra, “Minimizing makespan in a two-machine flow shop with delays and unit-time operations is np-hard,” *Journal of Scheduling*, vol. 7, no. 5, pp. 333–348, Sep. 2004, Copyright - Kluwer Academic Publishers 2004; Last updated - 2014-08-30.