

Speed-Robust Scheduling*

Sand, Bricks, and Rocks

Franziska Eberle[†] Ruben Hoeksma[‡] Nicole Megow[§] Lukas Nölke[§]
Kevin Schewior[¶] Bertrand Simon^{||}

July 18, 2022

Abstract

The speed-robust scheduling problem is a two-stage problem where, given m machines, jobs must be grouped into at most m bags while the processing speeds of the machines are unknown. After the speeds are revealed, the grouped jobs must be assigned to the machines without being separated. To evaluate the performance of algorithms, we determine upper bounds on the worst-case ratio of the algorithm’s makespan and the optimal makespan given full information. We refer to this ratio as the robustness factor. We give an algorithm with a robustness factor $2 - \frac{1}{m}$ for the most general setting and improve this to 1.8 for equal-size jobs. For the special case of infinitesimal jobs, we give an algorithm with an optimal robustness factor equal to $\frac{e}{e-1} \approx 1.58$. The particular machine environment in which all machines have either speed 0 or 1 was studied before by Stein and Zhong (ACM Trans. Alg. 2020). For this setting, we provide an algorithm for scheduling infinitesimal jobs with an optimal robustness factor of $\frac{1+\sqrt{2}}{2} \approx 1.207$. It lays the foundation for an algorithm matching the lower bound of $\frac{4}{3}$ for equal-size jobs.

1 Introduction

Scheduling problems with incomplete knowledge of the input data have been studied extensively. There are different ways to model such uncertainty, the major frameworks being *online optimization*, where parts of the input are revealed incrementally, *stochastic optimization*, where parts of the input are modeled as random variables, and *robust optimization*, where uncertainty in the data is bounded. Most scheduling research in this context assumes uncertainty about the job characteristics. Examples include online scheduling, where the job set is a priori unknown [1, 16], mixed-criticality scheduling, where the processing time comes from a given set [4], stochastic scheduling, where the processing times are modeled as random variables with known distributions [14, 15], robust scheduling, where the unknown processing times are within a given interval [13], two/multi-stage stochastic and robust scheduling with recourse, where the set of jobs that has to be scheduled stems

*Research was partially supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under project 146371743 within TRR 89 Invasive Computing and under contract ME 3825/1.

[†]London School of Economics and Political Sciences, Department of Mathematics, United Kingdom. Email: franziska.eberle@posteo.de

[‡]University of Twente, Department of Applied Mathematics, The Netherlands. r.p.hoeksma@utwente.nl

[§]University of Bremen, Department of Mathematics and Computer Science, Germany. {nmegow,noelke}@uni-bremen.de

[¶]University of Southern Denmark, Department of Mathematics and Computer Science, Odense, Denmark. kevs@sdu.dk

^{||}IN2P3 Computing Center, CNRS, Villeurbanne, France. bertrand.simon@cc.in2p3.fr

from a known superset and is revealed in stages [5, 17], and scheduling with explorable uncertainty, where the processing time of a job can potentially be reduced by testing the job at some extra cost [8].

A lot less research addresses uncertainty about the machine environment, particularly, where the processing speeds of machines change in an unforeseeable manner. The majority of such research focuses on the special case of scheduling with unknown non-availability periods, that is, machines break down temporarily [2, 7] or permanently [18]. Arbitrarily changing machine speeds have been considered for scheduling on a single machine [10].

Fluctuations in the processing speeds of machines are pervasive in real-world environments. For example, machines can be shared computing resources in data centers, where a sudden increase of workload may cause a general slowdown or, for some users, the machine may become unavailable due to priority issues. As another example, machines that are production units may change their processing speed due to aging or, unexpectedly, break down completely. In any case, (unforeseen) changes in the processing speed may have a drastic impact on the quality of a given schedule.

In this paper, we are concerned with the question of how to design a partial schedule by committing to groups of jobs, to be scheduled on the same machine, before knowing the actual machine speeds. This question is motivated, for example, by MapReduce computations done in large data centers. A MapReduce function typically groups workload before knowing the actual number or precise characteristics of the available computing resources [6].

We consider a two-stage robust scheduling problem in which we aim for a schedule of minimum makespan on multiple machines of unknown speeds. Given a set of n jobs and m machines, we ask for a partition of the jobs into m groups, we say *bags*, that have to be scheduled on the machines after their speeds are revealed without being split up. That is, in the second stage, when the machine speeds are known, a feasible schedule assigns all jobs in the same bag to the same machine. The goal is to minimize the second-stage makespan.

More formally, we define the *speed-robust scheduling* problem as follows. We are given n jobs with processing times $p_j \geq 0$, for $j \in \{1, \dots, n\}$, and the number of machines, $m \in \mathbb{N}$. Machines run in parallel but their speed is a priori unknown. In the first stage, the task is to group jobs into at most m bags. In the second stage, the machine speeds $s_i \geq 0$, for $i \in \{1, \dots, m\}$, are revealed. The time needed to execute job j on machine i is p_j/s_i , if $s_i > 0$. If a machine has speed $s_i = 0$, then it cannot process any job; we say the machine *fails*. Given the machine speeds, the second-stage task is to assign bags to the machines such that the makespan C_{\max} is minimized, where the makespan is the maximum sum of execution times of jobs assigned to the same machine.

Given a set of bags and machine speeds, the second-stage problem emerges as classical makespan minimization on related parallel machines. It is well-known that this problem can be solved arbitrarily close to optimality by polynomial-time approximation schemes [3, 11, 12]. As we are interested in the information-theoretic tractability, we allow superpolynomial running times for our algorithms – ignoring any computational concern – and assume that the second-stage problem is solved optimally. Thus, an *algorithm* for speed-robust scheduling defines a job-to-bag allocation, i.e., it gives a solution to the first-stage problem. We may use non-optimal bag-to-machine assignments to simplify the analysis.

We evaluate the performance of algorithms by a worst-case analysis comparing the makespan of the algorithm with the optimal makespan achievable when all machine speeds are known in advance. We say that an algorithm is ρ -*robust* if, for any instance, its makespan is within a factor $\rho \geq 1$ of the optimal solution. The *robustness factor* of the algorithm is defined as the infimum over all such ρ .

The special case of speed-robust scheduling where all machine speeds are either 0 or 1 has been studied previously by Stein and Zhong [18]. They introduced the problem with identical machines and an unknown number of machines that fail (speed 0) in the second stage. They present a

simple lower bound of $4/3$ on the robustness factor with equal jobs and design a general $5/3$ -robust algorithm. For infinitesimal jobs, they give an improved 1.2333 -robust algorithm complemented by a lower bound for each number of machines which tends to $(1 + \sqrt{2})/2 \approx 1.207$ for large m . Stein and Zhong also consider the objective of minimizing the maximum difference between the most loaded and the least loaded machine, motivated by questions on fair allocation.

Our Contribution

We introduce the speed-robust scheduling problem and present robust algorithms. The algorithmic difficulty of this problem is to construct bags in the first stage that are robust under any choice of machine speeds in the second stage. The straightforward approach of using any makespan-optimal solution on m identical machines is not sufficient. Lemma 3.1 shows that such an algorithm might have an arbitrarily large robustness factor. Using *Longest Processing Time* first (LPT) to create bags does the trick and is $(2 - \frac{1}{m})$ -robust for arbitrary job sizes (Theorem 3.3). While this was known for speeds in $\{0, 1\}$ [18], our result for general speeds is much less obvious.

Note that LPT aims at “balancing” the bag sizes which cannot lead to a better robustness factor than $2 - \frac{1}{m}$ as we show in Lemma 3.2. Hence, to improve upon this factor, we need to carefully construct bags with imbalanced bag sizes. There are two major challenges with this approach: (i) finding the ideal imbalance in the bag sizes independent from the actual job processing times that would be robust for all adversarial speed settings simultaneously and (ii) to adapt bag sizes to accommodate discrete jobs.

A major contribution of this paper is an optimal solution to the first challenge by considering infinitesimal jobs (Theorems 2.1 and 2.3). One can think of this as filling bags with *sand* to the desired level. In this case, the speed-robust scheduling problem boils down to identifying the best bag sizes as placing the jobs into bags becomes trivial. We give, for any number of machines, optimally imbalanced bag sizes and prove a best possible robustness factor of

$$\bar{\rho}(m) = \frac{m^m}{m^m - (m-1)^m} \leq \frac{e}{e-1} \approx 1.58.$$

For infinitely many infinitesimal jobs in the particular machine environment in which all machines have either speed 0 or 1, we obtain an algorithm (Theorem 2.4) with robustness factor

$$\bar{\rho}_{01}(m) = \max_{t \leq \frac{m}{2}, t \in \mathbb{N}} \frac{1}{\frac{t}{m-t} + \frac{m-2t}{m}} \leq \frac{1 + \sqrt{2}}{2} \approx 1.207 = \bar{\rho}_{01}.$$

This improves the previous upper bound of 1.233 by Stein and Zhong [18] and matches exactly their lower bound for each m . Furthermore, we show that the lower bound in [18] holds even for randomized algorithms (Theorem 2.1), so our algorithm is optimal for both, deterministic and randomized scheduling.

The above tight results for infinitesimal jobs are crucial for our further results for discrete jobs. Following the figurative notion of sand for infinitesimal jobs, we think of equal-size jobs as *bricks* and arbitrary jobs as *rocks*. Building on those ideal bag sizes, our approaches differ substantially from the methods in [18]. When all jobs have equal processing time, we obtain a 1.8 -robust solution through a careful analysis of the trade-off between using slightly imbalanced bags and a scaled version of the ideal bag sizes computed for the infinitesimal setting (Theorem 4.3).

When machines have only speeds in $\{0, 1\}$ and jobs have arbitrary equal sizes, i.e., unit size, we give an optimal $\frac{4}{3}$ -robust algorithm (Theorem 4.6). This is an interesting class of instances as the best known lower bound of $\frac{4}{3}$ for discrete jobs uses only unit-size jobs [18]. To achieve this result,

	General speeds		Speeds from $\{0, 1\}$	
	Lower bound	Upper bound	Lower bound	Upper bound
Discrete jobs (Rocks)	$\bar{\rho}(m)$ (Theorem 2.1)	$2 - \frac{1}{m}$ (Theorem 3.3)	$\frac{4}{3}$ [18]	$\frac{5}{3}$ [18]
Equal-size jobs (Bricks)	$\bar{\rho}(m)$ (Theorem 2.1)	1.8 (Theorem 4.3)	$\frac{4}{3}$ ([18], Theorem 4.6)	
Infinitesimal jobs (Sand)	$\bar{\rho}(m) \leq \frac{e}{e-1} \approx 1.58$ (Theorems 2.1 and 2.3)		$\bar{\rho}_{01}(m) \leq \frac{1+\sqrt{2}}{2} \approx 1.207$ ([18], Theorem 2.4)	

Table 1: Summary of results on speed-robust scheduling.

we, again, crucially exploit the ideal bag sizes computed for infinitesimal jobs by using a scaled variant of these sizes. Some cases, depending on m and the optimal makespan on m machines, have to be handled individually. Here, we use a direct way of constructing bags with at most four different bag sizes, and some cases can be solved by an integer linear program. We summarize our results in Table 1.

Inspired by traditional one-stage scheduling problems where jobs have *machine-dependent execution times* (unrelated machine scheduling), one might ask for such a generalization of our problem. However, it is easy to rule out any robustness factor for such a setting: Consider four machines and five jobs, where each job may be executed on a unique pair of machines. Any algorithm must build at least one bag with at least two jobs. For this bag, there is at most one machine to which it can be assigned with finite makespan. If this machine fails, the algorithm cannot complete the jobs whereas an optimal solution can split this bag on multiple machines to get a finite makespan.

2 Speed-Robust Scheduling with Infinitesimal Jobs

In this section, we assume that there are infinitely many jobs of infinitesimal processing time, we say *infinitesimal jobs*. We give optimal algorithms for speed-robust scheduling for both, the general case (Section 2.1) and the special case with speeds in $\{0, 1\}$ (Section 2.2).

2.1 General Speeds

We present an algorithm for speed-robust scheduling with infinitesimal jobs that achieves a best-possible robustness factor of $\bar{\rho}(m)$ for all $m \geq 1$, where

$$\bar{\rho}(m) = \frac{m^m}{m^m - (m-1)^m} \leq \frac{e}{e-1} \approx 1.58.$$

We first show that, even when we restrict the adversary to a particular set of speed configurations, no deterministic algorithm can achieve a robustness factor better than $\bar{\rho}(m)$. Note that since we can scale all speeds equally by an arbitrary factor without influencing the robustness factor, we can assume that the sum of the speeds is equal to 1. Similarly, we can assume that the total processing time of the jobs is equal to 1, such that the optimal makespan of the adversary is equal to 1 and the worst-case makespan of an algorithm is equal to its robustness factor.

Intuitively, the idea behind the set of m speed configurations is that the adversary can set $m-1$ machines to equal low speeds and one machine to a high speed. The low speeds are set such that

one particular bag size just fits on that machine when aiming for the given robustness factor. This immediately implies that all larger bags have to be put on the fast machine together. This way, the speed configuration can *target* a certain bag size. We provide specific bag sizes that achieve a robustness of $\bar{\rho}(m)$ and show that for the speeds targeting these bag sizes, other bag sizes would result in even larger robustness factors.

We define $U = m^m$, $L = m^m - (m-1)^m$, and $t_k = (m-1)^{m-k}m^{k-1}$ for $k \in \{1, \dots, m\}$. Intuitively, these values are chosen such that the bag sizes t_i/L are optimal and t_i/U corresponds to the low speed of the i -th speed configuration. It is easy to verify that $\bar{\rho}(m) = U/L$ and for all k we have

$$\sum_{i < k} t_i = (m-1)t_k - U + L. \quad (1)$$

In particular, this implies that $\sum_{i \leq m} t_i = mt_m - U + L = L$ and, hence, that the sum of the bag sizes is 1. Let $a_1 \leq \dots \leq a_m$ denote the bag sizes chosen by an algorithm and $s_1 \leq \dots \leq s_m$ the speeds chosen by the adversary.

Theorem 2.1. *For any $m \geq 1$, no deterministic algorithm for speed-robust scheduling with infinitesimal jobs can have a robustness factor less than $\bar{\rho}(m)$.*

Proof. We restrict the adversary to the following m speed configurations indexed by $k \in \{1, \dots, m\}$:

$$\mathcal{S}_k := \{s_1 = t_k/U, s_2 = t_k/U, \dots, s_{m-1} = t_k/U, s_m = 1 - (m-1)t_k/U\}.$$

Note that for all $k \in \{1, \dots, m\}$, we have $mt_k \leq U$ and, thus, $s_m \geq s_{m-1}$.

We show that for any bag sizes a_1, \dots, a_m , the adversary can force the algorithm to have a makespan of at least U/L with some \mathcal{S}_k . Since the optimal makespan is fixed to be equal to 1 by assumption, this implies a robustness factor of at least U/L .

Let k^* be the smallest index such that $a_k \geq t_k/L$. Such an index exists because the sum of the t_i 's is equal to L (Equation (1)) and the sum of the a_i 's is equal to 1. Now, consider the speed configuration \mathcal{S}_{k^*} . If one of the bags a_i for $i \geq k^*$ is not scheduled on the m -th machine, the makespan is at least $a_i/s_1 \geq a_{k^*}U/t_{k^*} \geq U/L$. Otherwise, all a_i for $i \geq k^*$ are scheduled on machine m . Then, using Equation (1), the load on that machine is at least

$$\sum_{i \geq k^*} a_i = 1 - \sum_{i < k^*} a_i \geq 1 - \frac{1}{L} \sum_{i < k^*} t_i = \frac{1}{L} (L - (m-1)t_{k^*} + U - L) = \frac{U}{L} s_m.$$

Thus, either a machine $i < m$ with a bag $i' \geq k^*$ or machine $i = m$ has a load of at least $s_i \cdot U/L$ and determines the makespan. \square

For given bag sizes, we call a speed configuration that maximizes the minimum makespan a *worst-case speed configuration*. Before we provide the strategy that obtains a matching robustness factor, we state a property of such best strategies for the adversary.

Lemma 2.2. *Given bag sizes and a worst-case speed configuration, for each machine i , there exists an optimal assignment of the bags to the machines such that only machine i determines the makespan.*

Proof. Consider a given set of bag sizes and a speed configuration $\{s_1, \dots, s_m\}$ that maximizes the minimum makespan for those bag sizes. Let C_{\max}^* be the minimum makespan of the best assignment of the bags given these speeds. This implies that, for any other speed configuration, there exists an assignment which has a makespan at most C_{\max}^* .

We prove the lemma by contradiction. If there exists a machine i that does not satisfy the lemma, we increase its speed by an additive factor of ε and we decrease the speed of all other machines by $\varepsilon/(m-1)$. Pick ε such that all (non-optimal) assignments that cause the load of machine i to be strictly greater than $C_{\max}^* \cdot s_i$ still satisfy that their respective load on machine i is strictly greater than $C_{\max}^* \cdot (s_i + \varepsilon)$. Denote the new speeds by s'_i for $1 \leq i' \leq m$. Now, consider any assignment. If the load of machine i is larger than $C_{\max}^* \cdot s_i$, it is also larger than $C_{\max}^* \cdot s'_i$ by the construction of s'_i and the assignment still has a makespan strictly larger than C_{\max}^* . Conversely, if the load of machine i is at most $C_{\max}^* \cdot s_i$, there must be a machine $i' \neq i$ with load at least $C_{\max}^* \cdot s_{i'}$. Otherwise, either i is the only machine that admits the makespan and already satisfies the lemma, or the assignment has a makespan smaller than C_{\max}^* , a contradiction. Consider any assignment where at least one other machine i' has load at least $C_{\max}^* \cdot s_{i'}$. Since we decreased the speed of all machines except machine i , the load of machine i' is strictly larger than $C_{\max}^* \cdot s'_{i'}$ leading to a makespan strictly greater than C_{\max}^* . This contradicts that the speed configuration maximizes the minimum makespan since every assignment with the new speeds has a makespan strictly larger than C_{\max}^* . \square

Note that, by Lemma 2.2, for a worst-case speed configuration, many bag-to-machine assignments obtain the optimal makespan. Lemma 2.2 also implies that, for such a speed configuration, all speeds are non-zero. Indeed, if a machine has a speed equal to zero, then it cannot determine the makespan in an optimal assignment (a better speed configuration would slow down other machines to increase its speed).

Let SAND denote the algorithm that creates m bags of the following sizes

$$a_1 = t_1/L, \quad a_2 = t_2/L, \quad \dots, \quad a_m = t_m/L.$$

Note that this is a valid algorithm since the sum of these bag sizes is equal to 1. Moreover, these bag sizes are exactly such that if we take the speed configurations from the proof of Theorem 2.1, placing bag j on a slow machine in configuration j results in a makespan that is equal to $\bar{\rho}(m)$.

We proceed to show that SAND has a robustness factor of $\bar{\rho}(m)$.

Theorem 2.3. *For any $m \geq 1$, SAND is $\bar{\rho}(m)$ -robust for speed-robust scheduling with infinitesimal jobs.*

Proof. Let a_1, \dots, a_m be the bag sizes as specified by SAND and let s_1, \dots, s_m be a speed configuration that maximizes the minimum makespan given these bag sizes. Further, consider an optimal assignment of bags to machines and let C_{\max}^* denote its makespan. We use one particular (optimal) assignment to obtain an upper bound on C_{\max}^* . By Lemma 2.2, there exists an optimal assignment where only machine 1 determines the makespan, i.e., machine 1 has load $C_{\max}^* \cdot s_1$ and any other machine i has load strictly less than $C_{\max}^* \cdot s_i$. Consider such an assignment. If there are two bags assigned to machine 1, then there is an empty machine with speed at least s_1 . Therefore, we can put one of the two bags on that machine and decrease the makespan. This contradicts C_{\max}^* being the optimal makespan, so there is exactly one bag assigned to machine 1. Let k be the index of the unique bag placed on machine 1, i.e., $C_{\max}^* = a_k/s_1$, and let ℓ be the number of machines of speed s_1 .

If $a_k > a_\ell$, machine $i \in \{1, \dots, \ell\}$ with speed s_1 can be assigned bag i with a load that is strictly less than $C_{\max}^* \cdot s_1$. Thus, given the current assignment, we can remove bag a_k from machine 1 and place the ℓ smallest bags on the ℓ slowest machines, one per machine, e.g., bag a_i on machine i for $i \in \{1, \dots, \ell\}$. This empties at least one machine of speed strictly larger than s_1 . Then, we can place bag a_k on this, now empty, machine, which yields a makespan that is strictly smaller

than C_{\max}^* . This contradicts the assumption that C_{\max}^* is the optimal makespan and, thus, $a_k \leq a_\ell$, which implies $k \leq \ell$.

Let P_i denote the total processing time of bags that are assigned to machine i and let C be the total *remaining capacity* of the assignment, that is, $C := \sum_{i=1}^m (s_i C_{\max}^* - P_i)$. We construct an upper bound on C , which allows us to bound C_{\max}^* .

Machines in the set $\{2, \dots, \ell\}$ cannot be assigned a bag of size larger than a_k since their load would be greater than $C_{\max}^* \cdot s_1$, causing a makespan greater than C_{\max}^* . Therefore, we assume without loss of generality that all bags $a_j < a_k$ are assigned to a machine with speed s_1 . The total remaining capacity on the first k machines is therefore equal to $(k-1)a_k - \sum_{i < k} a_i$.

Consider a machine $i > k$. If its remaining capacity is greater than a_k , then we can decrease the makespan of the assignment by moving bag k to machine i . Therefore, the remaining capacity on machine i is at most a_k .

Combining the above and using (1), we obtain:

$$C \leq (m-1)a_k - \sum_{i < k} a_i = \frac{1}{L} \left((m-1)t_k - \sum_{i < k} t_i \right) = \frac{1}{L} (U - L).$$

The total processing time is $\sum_{i=1}^m a_i = 1$, and the maximum total processing time the machines could process with makespan C_{\max}^* is $\sum_{i=1}^m s_i C_{\max}^* = C_{\max}^*$. Since the latter is equal to the total processing time plus the remaining capacity, we have $C_{\max}^* = 1 + C \leq U/L$, which proves the lemma. \square

While the robustness factor $\bar{\rho}(m)$ is best possible for every m for any deterministic algorithm, this is not true when we allow algorithms that make randomized decisions and compare to an oblivious adversary. For $m = 2$, uniformly randomizing between bag sizes $a_1 = a_2 = 1/2$ and $a_1 = 1/4, a_2 = 3/4$ yields a robustness factor that is slightly better than $\bar{\rho}(2) = 4/3$. Interestingly, with speeds in $\{0, 1\}$, the optimal robustness factor is equal for deterministic and randomized algorithms as we show in Theorem 2.8.

2.2 Speeds in $\{0, 1\}$

This section is devoted to showing that the best-possible robustness factor that can be achieved with speeds in $\{0, 1\}$ is precisely

$$\bar{\rho}_{01}(m) = \max_{t \in \mathbb{N}, t \leq \frac{m}{2}} \frac{1}{\frac{t}{m-t} + \frac{m-2t}{m}} \leq \frac{1 + \sqrt{2}}{2} = \bar{\rho}_{01} \approx 1.207$$

for both, deterministic and randomized algorithms. The main contribution is the upper bound, that is, the following theorem.

Theorem 2.4. *For all $m \geq 1$, there is a deterministic $\bar{\rho}_{01}(m)$ -robust algorithm for speed-robust scheduling with speeds in $\{0, 1\}$ for infinitesimal jobs.*

The lower bound for deterministic algorithms and some useful insights were already presented in [18]. We recall some of these insights here because they are used in the proof. To do so, we introduce some necessary notation used in the remainder of this paper. The number of failing machines (i.e., machines with speed equal to 0) is referred to as $t \geq 0$, and we assume w.l.o.g. that these are machines $1, \dots, t$. Furthermore, we assume for this subsection again w.l.o.g. that the total volume of infinitesimal jobs is m , and we will define bags $1, \dots, m$ with respective sizes $a_1 \leq \dots \leq a_m$ summing to at least m (the potential excess being unused).

Lemma 2.5 (Statement (3) in [18]). *For all $m \geq 1$ and $t \leq m/2$, there exists a makespan-minimizing allocation of bags to machines for speed-robust scheduling with speeds in $\{0, 1\}$ and infinitely many infinitesimal jobs that assigns the smallest $2t$ bags to machines $t + 1, \dots, 2t$.*

Since Lemma 2.5 only works for $t \leq m/2$, one may worry that, for larger t , there is a more difficult structure to understand. The following insight shows that this worry is unjustified. Indeed, if $m' < m/2$ is the number of machines that do *not* fail, one can simply take the solution for $2m'$ machines and assign the bags from any two machines to one machine. The optimal makespan is doubled and that of the algorithm is at most doubled, conserving the robustness.

Lemma 2.6 (Proof of Theorem 2.2 in [18]). *Let $\rho > 1$. For all $m \geq 1$, if an algorithm is ρ -robust for speed-robust scheduling with speeds in $\{0, 1\}$ and infinitely many infinitesimal jobs for $t \leq m/2$, it is ρ -robust for $t \leq m - 1$.*

We will thus focus on computing bag sizes such that the makespan of a best allocation according to Lemma 2.5 is within a $\bar{\rho}_{01}(m)$ factor of the optimal makespan when $t \leq m/2$. The approach in [18] to obtain the (as we show tight) lower bound $\bar{\rho}_{01}(m)$ is as follows. Given some $t \leq m/2$ and a set of bags allocated according to Lemma 2.5,

- (i) the load on machines $t + 1, \dots, 2t$ is at most $\bar{\rho}_{01}(m)$ times the optimal makespan $m/(m - t)$, and
- (ii) the load on machines $2t + 1, \dots, m$ is at most $\bar{\rho}_{01}(m)$ because those machines only hold a *single* bag after a simple “folding” strategy for assigning bags to machines, which we define below.

In particular, since $t = 0$ is possible, (ii) implies that all bag sizes are at most $\bar{\rho}_{01}(m)$. The fact that the total processing volume of m has to be accommodated and maximizing over t results in the lower bound given in Theorem 2.4.

To define the bag sizes leading to a matching upper bound, we further restrict our choices when $t \leq m/2$ machines fail. Of course, since we match the lower bound, the restriction is no limitation but rather a simplification. When $t \leq m/2$ machines fail, we additionally assume that the machines $t + 1, \dots, 2t$ receive exactly two bags each: Assuming $t \leq m/2$, the *simple folding* of these bags onto machines assigns bags $i \geq t + 1$ to machine i , and bags $i = 1, \dots, t$ (recall machine i fails) to machine $2t - i + 1$. Hence, bags $1, \dots, t$ are “folded” onto machines $2t, \dots, t + 1$ (sic), visualized in Figure 1.

For given m , let t^* be an optimal adversarial choice for t in Theorem 2.4. Assuming there are bag sizes a_1, \dots, a_m that match the bound $\bar{\rho}_{01}(m)$ through simple folding, by (i) and (ii), we precisely know the makespan on all machines after folding when $t = t^*$. That fixes $a_i + a_{2t+1-i} = \bar{\rho}_{01}(m) \cdot m/(m - t)$ for all $i = 1, \dots, t$ and $a_{2t+1}, \dots, a_m = \bar{\rho}_{01}(m)$, see Figure 1. In contrast to [18], we show that defining a_i for $i = 1, \dots, t$ to be essentially a linear function of i , and thereby fixing all bag sizes, suffices to match $\bar{\rho}_{01}(m)$. The word “essentially” can be dropped when replacing $\bar{\rho}_{01}(m)$ by $\bar{\rho}_{01}$.

A clean way of thinking about the bag sizes is through *profile functions* which reflect the distribution of load over bags in the limit case $m \rightarrow \infty$. Specifically, we identify the set $\{1, \dots, m\}$ with the interval $[0, 1]$ and define a continuous non-decreasing profile function $\bar{f} : [0, 1] \rightarrow \mathbb{R}_+$ integrating to 1. A simple way of getting back from the profile function to actual bag sizes of total size approximately m is by equidistantly “sampling” \bar{f} , that is, by defining $a_i := \bar{f}(\frac{i-1/2}{m})$ for all i .

Our profile function \bar{f} implements the above observations and ideas in the continuous setting. Indeed, our choice

$$\bar{f}(x) = \min \left\{ \frac{1}{2} + \bar{\rho}_{01} \cdot x, \bar{\rho}_{01} \right\} = \min \left\{ \frac{1}{2} + \frac{(1 + \sqrt{2}) \cdot x}{2}, \frac{1 + \sqrt{2}}{2} \right\}$$




Figure 1: Folding optimally sized bags when t^* machines fail.




Figure 2: The profile function \bar{f} and the equidistant “sampling” of SAND_{01} .

is linear up to $\beta = 2 - \sqrt{2}$, which turns out to be equal to $\lim_{m \rightarrow \infty} 2t^*/m$, and then constantly equal to $(1 + \sqrt{2})/2$ since

$$\frac{1}{2} + \frac{(1 + \sqrt{2}) \cdot x}{2} \leq \frac{1 + \sqrt{2}}{2} \Leftrightarrow x \leq \frac{\sqrt{2}}{1 + \sqrt{2}} = 2 - \sqrt{2}.$$

We give some intuition for why this function yields the desired bound using the continuous counterpart of folding. When $t \leq t^*$ machines fail, i.e., a continuum of machines with measure $x \leq \beta/2$, we fold the corresponding part of \bar{f} onto the interval $[x, 2x]$, yielding a rectangle of width x and height $\bar{f}(0) + \bar{f}(2x) = 2\bar{f}(x)$. We have to prove that the height does not exceed the optimal makespan $1/(1-x)$ by more than a factor of $\bar{\rho}_{01}$. Equivalently, we maximize $2\bar{f}(x)(1-x)$ (even over $x \in \mathbb{R}$) and observe the maximum of $\bar{\rho}_{01} = (1 + \sqrt{2})/2$ at $x = \beta/2$. When $x \in (\beta/2, 1/2]$, note that by folding, we *still* obtain a rectangle of height $2\bar{f}(x)$ (but width $\beta - x$), dominating the load on the other machines. Hence, the makespan is at most $\bar{\rho}_{01}/(1-x)$ for every $x \in [0, 1/2]$.

Directly “sampling” \bar{f} , we obtain a bound of $\bar{\rho}_{01}$, stated below. Later, we make use of the corresponding simpler algorithm. Let SAND_{01} denote the algorithm that creates m bags of size $a_i := \bar{f}(\frac{i-1/2}{m})$, for $i \in \{1, \dots, m\}$. See Figure 2 for a visualization.

Theorem 2.7. SAND_{01} is $\bar{\rho}_{01}$ -robust for speed-robust scheduling with speeds in $\{0, 1\}$ and infinitely many infinitesimal jobs for all $m \geq 1$.

Proof. Our proof naturally splits into two parts. In the first part, we show that the bag sizes are feasible, i.e., their total size is at least m . In the second part, we show that the bag sizes achieve the claimed robustness factor.

To show that the bag sizes are feasible, we prove that

$$a_i \geq m \cdot \int_{\frac{i-1}{m}}^{\frac{i}{m}} \bar{f}(x) dx \quad (2)$$

for all $i \in \{1, \dots, m\}$. Then the first part follows by summing over (2) for all i and indeed, as required for a profile function, \bar{f} integrates to 1:

$$\begin{aligned} \int_0^1 \bar{f}(x) dx &= \int_0^\beta \left(\frac{1}{2} + \bar{\rho}_{01} x \right) dx + \int_\beta^1 \bar{\rho}_{01} dx \\ &= \frac{\beta}{2} + \frac{\bar{\rho}_{01} \beta^2}{2} + (1 - \beta) \cdot \bar{\rho}_{01} = 1. \end{aligned}$$

For i with $\beta \notin (\frac{i-1}{m}, \frac{i}{m})$, we have that \bar{f} is linear on $[\frac{i-1}{m}, \frac{i}{m}]$ and therefore

$$\int_{\frac{i-1}{m}}^{\frac{i}{m}} \bar{f}(x) dx = \left(\frac{i}{m} - \frac{i-1}{m} \right) \cdot \bar{f} \left(i - \frac{1}{2} \right) = \frac{f(i-1/2)}{m} = \frac{a_i}{m}.$$

For the single i with $\beta \in (\frac{i-1}{m}, \frac{i}{m})$ — and there is at least one such i because β is irrational — we even get a stronger bound due to the fact that \bar{f} is strictly concave on that interval. Formally, we distinguish two cases, in which we use either of the two linear functions from the definition of \bar{f} as upper bound on \bar{f} .

If $\frac{i-1/2}{m} \leq \beta$, it follows that

$$\int_{\frac{i-1}{m}}^{\frac{i}{m}} \bar{f}(x) dx \leq \int_{\frac{i-1}{m}}^{\frac{i}{m}} \left(\frac{1}{2} + \bar{\rho}_{01} x \right) dx = \left(\frac{i}{m} - \frac{i-1}{m} \right) \cdot \bar{f} \left(i - \frac{1}{2} \right) = \frac{a_i}{m}.$$

For $\frac{i-1/2}{m} > \beta$, we have that

$$\int_{\frac{i-1}{m}}^{\frac{i}{m}} \bar{f}(x) dx \leq \int_{\frac{i-1}{m}}^{\frac{i}{m}} \bar{\rho}_{01} dx = \left(\frac{i}{m} - \frac{i-1}{m} \right) \cdot \bar{f} \left(i - \frac{1}{2} \right) = \frac{f(i-1/2)}{m} = \frac{a_i}{m}.$$

That finishes the proof of (2) for all i and verifies that our bag sizes are feasible.

It remains to show that our bag sizes achieve the claimed robustness factor of $\bar{\rho}_{01} \approx 1.207$. Essentially, the argument is a formal version of the intuitive argument we gave in the continuous setting, restricted to x (the measure of the continuum of failing machines) being $\frac{i-1/2}{m}$ for some $i \in \{1, \dots, m\}$. By Lemma 2.6, it suffices to consider the case that $t \leq m/2$ machines fail. By our self-imposed restriction, we only consider bag-to-machine assignments obtained through simple folding. Note that we only need to bound the load on machines that have two bags assigned to them after folding. Recall that these machines are machines $t+1, \dots, 2t$; for all $i \in \{1, \dots, t\}$, bags i and $2t+1-i$ are assigned to machine $2t+1-i$. Also recall that

$$a_i \leq \frac{1}{2} + \bar{\rho}_{01} \cdot \frac{i - \frac{1}{2}}{m}$$

for all $i \in \{1, \dots, m\}$. Hence, the load created by bags $i \in \{1, \dots, t\}$ and $2t+1-i$ on machine $2t+1-i$ is

$$a_i + a_{2t+1-i} \leq \frac{1}{2} + \bar{\rho}_{01} \cdot \frac{i - \frac{1}{2}}{m} + \frac{1}{2} + \bar{\rho}_{01} \cdot \frac{2t+1-i - \frac{1}{2}}{m} = 1 + 2\bar{\rho}_{01} \cdot \frac{t}{m}.$$

We would like to show that this load is at most a $\bar{\rho}_{01}$ factor away from the load of the optimum, that is,

$$1 + 2\bar{\rho}_{01} \cdot \frac{t}{m} \leq \bar{\rho}_{01} \cdot \frac{m}{m-t}.$$

Letting $x := t/m$ yields the inequality

$$(1 + 2\bar{\rho}_{01}x)(1-x) \leq \bar{\rho}_{01}.$$

The left hand side of this inequality takes its maximum $\bar{\rho}_{01}$ (even among all $x \in \mathbb{R}$) at $x = \beta = 2 - \sqrt{2}$, which shows that the inequality is indeed true and therefore completes the proof. \square

To show Theorem 2.4, however, we have to match $\bar{\rho}_{01}(m)$ for every m . To do so, we need to design bag sizes more carefully. Indeed, $\bar{\rho}_{01}(m)$ is smaller than $\bar{\rho}_{01}$ for all values of m , and therefore bags of sizes $\bar{\rho}_{01}$ are not allowed anymore. For every $x \in [0, 1]$, the bag size of the $\lceil x \cdot m \rceil$ -th smallest bag still approaches $\bar{f}(x)$ as $m \rightarrow \infty$, rather than being defined *through* \bar{f} . Specifically, in the appendix we give a family of bag sizes parameterized by some $\delta > 0$ that allow a simple computation of the robustness factor. The remainder of the proof is then concerned with showing algebraically that, for each m , δ can be chosen so as to fulfill the constraints imposed by feasibility and robustness.

We close this section with showing that no better robustness factor can be achieved even by a randomized algorithm.

Theorem 2.8. *For all $m \geq 1$ and $\varepsilon > 0$, there is no randomized algorithm that achieves a robustness factor of $\bar{\rho}_{01}(m) - \varepsilon$ against an oblivious adversary.*

Proof. The outline of this proof is based on the same result from [18] for deterministic algorithms. Consider any randomized algorithm, the size of each bag follows some probability distribution which can be correlated. The problem can be described as follows: the adversary first selects the number t of machine failures, knowing the distribution of the bag sizes; then, the actual bag sizes are revealed; finally, the algorithm schedules these bags on $m - t$ machines. We assume by contradiction that, for every t , the expectation of the resulting makespan is smaller than $\bar{\rho}_{01}(m)/(1 - \frac{t}{m})$.

We consider an adversary with two possible strategies: make zero machines fail, or make $t \leq m/2$ machines fail, the value of t being fixed later. For large m , t/m will approach $1 - \sqrt{2}/2$ and $\bar{\rho}_{01}(m)$ will approach $\bar{\rho}_{01}$.

The expected size of each bag must be smaller than $\bar{\rho}_{01}(m)$; otherwise the expected makespan on m machines would be too large. For every realization of bag sizes, there exists an optimal bag-to-machine allocation on $m - t$ machines that uses all machines, so has at least $m - 2t$ machines containing a single bag. Reorder the machines so that machines $t+1$ to $m-t$ have a single bag. The expected load of each of the first t machines is smaller than $\bar{\rho}_{01}(m)/(1 - \frac{t}{m})$ as the optimal makespan on $m - t$ machines is $1/(1 - \frac{t}{m})$. The expected load of each of the other machines is smaller than $\bar{\rho}_{01}(m)$ as they contain a single bag. By linearity of expectation, and due to the expected total load being equal to m , we obtain the following contradiction:

$$\begin{aligned} m &< \min_{t \leq \frac{m}{2}, t \in \mathbb{N}} \frac{t \cdot \bar{\rho}_{01}(m)}{1 - \frac{t}{m}} + (m - 2t)\bar{\rho}_{01}(m) \\ \bar{\rho}_{01}(m) &> \max_{t \leq \frac{m}{2}, t \in \mathbb{N}} \frac{1}{\frac{t}{m-t} + \frac{m-2t}{m}} = \bar{\rho}_{01}(m). \end{aligned}$$

This completes the proof. \square

3 Speed-Robust Scheduling with Discrete Jobs

In this section, we consider the most general version of speed-robust scheduling, i.e., discrete jobs scheduled on machines with arbitrary unknown speeds. While in Sections 2 and 4 we crucially use in our algorithm design the assumption that jobs are infinitesimally small (sand) or of the same size (bricks), respectively, here, their sizes can vary arbitrarily (rocks). By a scaling argument, we may assume w.l.o.g. that the machine speeds satisfy $\sum_{i=1}^m s_i = \sum_{j=1}^n p_j$. We first note in the following lemma that obtaining a robust algorithm is not trivial in this case, as even algorithms minimizing the largest bag size may not have a constant robustness factor. This contrasts with the case where machine speeds are restricted to $\{0, 1\}$, in which such algorithms are $(2 - \frac{2}{m})$ -robust. To see this, once the number m' of speed-1 machines is revealed, simply combine the two smallest bags repetitively if $m' < m$. The makespan is then at most twice the average load on $m' + 1$ machines, so $\frac{2m'}{m'+1}$ times the average load on m' machines. This is largest for $m' = m - 1$ which gives the desired robustness.

Lemma 3.1. *There exists an algorithm for speed-robust scheduling minimizing the size of the largest bag which does not have a constant robustness factor.*

Proof. Consider any integer $k \geq 1$, a number of machines $m = k^2 + 1$, k^2 unit-size jobs and one job of processing time k . The maximum bag size is equal to k , so an algorithm building $k + 1$ bags of size k respects the conditions of the lemma. Consider the speed configuration where k^2 machines have speed 1 and one machine has speed k . It is possible to schedule all jobs within a makespan 1 on these machines. However, the algorithm must either place a bag on a machine of speed 1 or all bags on the machine of speed k , hence leading to a makespan of k , and proving the result. Note that by adding k^2 unit-size jobs, we can build a similar example where the algorithm does not leave bags empty, which is always beneficial. \square

A feature, that is exploited in the lower bound, of the algorithms considered in Lemma 3.1 is that bags sizes are too unbalanced. A way to prevent this would be to maximize the size of a minimum bag as well. But this criterion becomes useless if we consider the same example as above with $m = k^2 + 2$. Then, the minimum bag size is 0 as there are more machines than jobs, and the same lower bound holds.

Hence, in order to obtain a robust algorithm in the general case, we focus on algorithms that aim at balanced bag sizes, for which the best lower bound is described in the following lemma. An algorithm is called *balanced* if, for an instance of unit-size jobs, the bag sizes created by the algorithm differ by at most one unit. In particular, a balanced algorithm creates m bags of size k when confronted with mk unit-size jobs and m bags. For balanced algorithms, we give a lower bound in Lemma 3.2 and a matching upper bound in Theorem 3.3.

Lemma 3.2. *No balanced algorithm for speed-robust scheduling can obtain a better robustness factor than $2 - \frac{1}{m}$ for any $m \geq 1$.*

Proof. Consider any $m \geq 1$ and km unit-size jobs, with $k = 2m - 1$. Assume the adversary puts m jobs on the first machine of speed m and $2m$ jobs on each of the remaining machines of speed $2m$ each. An algorithm that uses evenly balanced bags builds m bags of size k . It must either place a bag of size k on the machine of speed m or $2k$ jobs on a machine of speed $2m$. In any case, the robustness factor is at least $2 - \frac{1}{m}$. \square

We now show that this lower bound is attained by a simple algorithm, commonly named as *Longest Processing Time First* (LPT) which considers jobs in non-increasing order of processing

times and assigns each job to the bag that currently has the smallest size, i.e., the minimum allocated processing time.

Theorem 3.3. *LPT is $(2 - \frac{1}{m})$ -robust for speed-robust scheduling for all $m \geq 1$.*

Proof. While we may assume that the bags are allocated optimally to the machines once the speeds are given, we use a different allocation for the analysis. This can only worsen the robustness factor.

Consider the m bags and let b denote the size of a largest bag, B , that consists of at least two jobs. Place all bags of size strictly larger than b , each containing only a single job, on the same machine as OPT places the corresponding jobs. We define for each machine i with given speed s_i a capacity bound of $(2 - \frac{1}{m}) \cdot s_i$. Then, we consider the remaining bags in non-increasing order of bag sizes and iteratively assign them to the – at the time of assignment – least loaded machine with sufficient remaining capacity.

With the assumption $\sum_{i=1}^m s_i = \sum_{j=1}^n p_j$ and the capacities $(2 - \frac{1}{m}) \cdot s_i$, it is sufficient to show that LPT can successfully place all bags.

The bags of size larger than b fit by definition as they contain a single job. Suppose for the sake of contradiction that there is a bag which cannot be assigned. Consider the first such bag and let T be its size. Let $k < m$ be the number of bags that have been assigned already. Further, denote by w the size of a smallest bag. Since we used LPT in creating the bags, we can show that $w \geq \frac{1}{2}b$. To see that, consider bag B and notice that the smallest job in it has size at most $\frac{1}{2}b$. When this job was assigned to its bag, B was a bag with the smallest size, and this size was at least $\frac{1}{2}b$ since we allocate jobs in LPT-order. Hence, the size of a smallest bag is $w \geq \frac{1}{2}b \geq \frac{1}{2}T$, where the second inequality is true as all bags larger than b can be placed.

We use this inequality to give a lower bound on the total remaining capacity on the m machines when the second-stage algorithm fails to place the $(k+1)$ -st bag. The $(m-k)$ bags that were not placed have a combined volume of at least $V_\ell = (m-k-1)w + T \geq (m-k+1)\frac{T}{2}$. The bags that were placed have a combined volume of at least $V_p = kT$. The remaining capacity is then at least $C = (2 - \frac{1}{m})V_\ell + (1 - \frac{1}{m})V_p$, and we have

$$C \geq \frac{2m-1}{m} \frac{m-k+1}{2} T + \frac{m-1}{m} kT \geq mT + T - \frac{m+k+1}{2m} T \geq mT.$$

Thus, there is a machine with remaining capacity T which contradicts the assumption that the bag of size T does not fit. \square

4 Speed-Robust Scheduling with Equal-Size Jobs

In this section, we consider instances where all jobs are of equal size (bricks) as this case seems to capture the complexity of the general problem. This intuition stems from the fact that all known lower bounds already hold for this type of instances (see [18] and Lemma 4.4). By a scaling argument, we may assume that all jobs have unit processing time. Therefore, we consider unit-size jobs for the remainder of the section.

Before focusing on a specific speed setting, we show that in both settings we can use any algorithm for infinitesimal jobs with a proper scaling to obtain a corresponding algorithm for unit-size jobs. Its robustness factor improves with increasing average load per bag $n/m =: \lambda$. Assume $\lambda > 1$, as otherwise the problem is trivial. We define the algorithm SANDFORBRICKS that builds on the optimal algorithm for infinitesimal jobs, SAND*, which is SAND for general speeds (Section 2.1) or SAND₀₁ for speeds in $\{0, 1\}$ (Section 2.2). Let a_1, \dots, a_m be the bag sizes constructed by SAND* scaled such that a total processing volume of n can be assigned, that is, $\sum_{i=1}^m a_i = n$. For unit-size jobs, we define bag sizes as $a'_i = (1 + \frac{1}{\lambda}) \cdot a_i$ and assign the jobs greedily to the bags.

Lemma 4.1. *For n jobs with unit processing times and m machines, SANDFORBRICKS for speed-robust scheduling is $(1 + \frac{1}{\lambda}) \cdot \rho(m)$ -robust, where $\rho(m)$ is the robustness factor for SAND* for m machines.*

Proof. To prove the lemma, it is sufficient to show that all n unit-size jobs can be assigned to the constructed bags of sizes a'_1, \dots, a'_m . Suppose for the sake of contradiction that there is a job j that does not fit into any bag without exceeding the bag size. The remaining volume in the bags is at least the total capacity minus the processing volume of all jobs except j , that is,

$$\sum_{i=1}^m a'_i - (n - 1) = \left(1 + \frac{1}{\lambda}\right) \cdot n - n + 1 > \frac{1}{\lambda} \cdot n = m.$$

Hence, there must exist some bag that has a remaining capacity of at least 1 and can fit job j . \square

4.1 General Speeds

For unit-size jobs, we show how to beat the robustness factor $2 - \frac{1}{m}$ (Theorem 3.3) for $m > 4$ and obtain a 1.8-robust algorithm. For $m = 2$ and $m = 3$, we give algorithms with best-possible robustness factors $\frac{4}{3}$ and $\frac{3}{2}$, respectively.

Theorem 3.3 shows that LPT has a robustness factor of $2 - \frac{1}{m}$, even for unit-size jobs. BUILDODD has a robustness factor increasing with the ratio between the number of jobs and the number of machines. BUILDODD builds bags of three possible sizes: for $q \in \mathbb{N}$ such that $\lambda \in [2q - 1, 2q + 1]$, bags of sizes $2q - 1$ and $2q + 1$ are built, with possibly one additional bag of size $2q$ (recall that $\lambda = n/m$ is the average load per bag).

Lemma 4.2. *For n unit-size jobs, m machines and $q \in \mathbb{N}$ with $\lambda \in [2q - 1, 2q + 1]$, BUILDODD is $(2 - \frac{1}{q+1})$ -robust for speed-robust scheduling.*

Proof. The proof is along the lines of the proof of Theorem 3.3. Recall that BUILDODD builds m bags of sizes belonging to $\{2q - 1, 2q, 2q + 1\}$, with at most one bag of size $2q$. This is possible by building first m bags of size $2q - 1$ then putting 2 additional jobs per bag until zero or one job remains. Let m_s be the number of small bags (size $2q - 1$), m_m be the number of medium bags (size $2q$) and m_b be the number of big bags (size $2q + 1$). We have $m_s + m_m + m_b = m$ and $n = (2q + 1)m_b + 2qm_m + (2q - 1)m_s$.

We assume $q \geq 1$ since, if $q = 0$, only bags of size 1 are built and the problem is trivial.

Note that a small bag can be executed at speed q and a large or medium bag can be executed at speed $q + 1$ while respecting the prescribed makespan of $2 - \frac{1}{q+1} = \frac{2q+1}{q+1}$. We define the *weight* of a small bag as q and the weight of a large or medium bag as $q + 1$.

In the second stage, when the speed $s_i \in \mathbb{N}$ is given for each machine i , associate a *capacity* s_i with each machine. Assign the bags in LPT order to the machines, each bag to the least loaded machine such that the total *weight* of bags assigned to a machine does not exceed the capacity. The total capacity of all bins is equal to n . If all bags can be assigned to the machines, then the total *size* of the bags assigned to a machine of speed s_i is at most $(2 - \frac{1}{q+1})s_i$, which gives the result.

Suppose for the sake of contradiction that there is a bag that cannot be assigned to a machine. Let T be the *weight* of this bag. It suffices to show that the total remaining capacity on all machines is at least $m(T - 1) + 1$. Indeed, weights and capacities are integers, so if the average remaining capacity per machine is strictly larger than $T - 1$, one machine has a remaining capacity at least T and the bag fits.

Assume first that the bag is small, i.e., $T = q$. The total weight placed so far is at most $(q + 1)(m_b + m_m) + q(m_s - 1)$, so the remaining capacity is:

$$\begin{aligned}
C &\geq n - (q + 1)m_b - (q + 1)m_m - qm_s + q \\
&= (2q + 1)m_b + 2qm_m + (2q - 1)m_s - (q + 1)m_b - (q + 1)m_m - qm_s + q \\
&= qm_b + (q - 1)m_m + (q - 1)m_s + q \\
&\geq m(T - 1) + q.
\end{aligned}$$

Assume now that the bag is big or medium, so $T = q + 1$. The total weight placed so far is at most $(q + 1)m_b$ and, thus, the remaining capacity is at least:

$$\begin{aligned}
C &\geq n - (q + 1)m_b + q \\
&= (2q + 1)m_b + 2qm_m + (2q - 1)m_s - (q + 1)m_b + q \\
&= qm_b + qm_m + qm_s + (q - 1)m_s + qm_m + q \\
&\geq m(T - 1) + q.
\end{aligned}$$

We conclude that all bags can be assigned to the machines without exceeding the capacity. Hence, this algorithm is $(2 - \frac{1}{q+1})$ -robust. \square

Notice that the robustness guarantees in Lemmas 4.1 and 4.2 are functions that are decreasing in λ and increasing in λ , respectively. By carefully choosing between BUILDODD and SANDFORBRICKS, depending on the input, we obtain an improved algorithm for bricks. For $\lambda < 8$, we execute BUILDODD, which yields a robustness factor of at most 1.8 by Lemma 4.2, as $q \leq 4$ for $\lambda < 8$. Otherwise, when $\lambda \geq 8$, we run SANDFORBRICKS and obtain a guarantee of $\frac{9}{8} \cdot \frac{e}{e-1} \approx 1.78$ by Lemma 4.1.

Theorem 4.3. *There is an algorithm for speed-robust scheduling with unit-size jobs that has a robustness factor of at most 1.8 for any $m \geq 1$.*

We give a general lower bound on the best achievable robustness factor. Note that the lower bound of $\bar{\rho}(m)$ from Theorem 2.1 remains valid in this setting and is larger than 1.5 for $m \geq 6$.

Lemma 4.4. *For every $m \geq 3$, no algorithm for speed-robust scheduling can have a robustness factor smaller than 1.5, even restricted to unit-size jobs.*

Proof. Consider an instance with $2m$ unit-size jobs. If an algorithm places 3 jobs in a bag, the adversary selects identical speeds which leads to a makespan $3/2$ times larger than the optimal. Otherwise, the adversary chooses a speed 1 for $m - 1$ machines and a speed $m + 1$ for the remaining machine, thus being able to complete the instance within a makespan 1. The algorithm then has to put all the bags on the fastest machine to obtain a robustness factor smaller than 2. The factor is equal to $2m/(m + 1)$ which is at least 1.5 for $m \geq 3$. \square

For special cases with few machines, we give best possible algorithms which match the previously mentioned lower bounds. We also show for $m = 6$ a lower bound larger than $\bar{\rho}(6) > 1.5$; proofs can be found in Appendix B. Similar lower bounds have been found by a computer search for many larger values of m , for which the difference to $\bar{\rho}(m)$ tends towards zero when m grows.

Lemma 4.5. *An optimal algorithm for speed-robust scheduling for unit-size jobs has robustness factor $4/3$ on $m = 2$ machines and $3/2$ on $m = 3$ machines, and larger than $\bar{\rho}(6) > 1.5$ for $m = 6$.*

4.2 Speeds in $\{0, 1\}$

When considering speeds in $\{0, 1\}$, bricks are of particular interest as the currently best known lower bound for rocks (arbitrary jobs) is $\frac{4}{3}$ and uses only bricks (unit-size jobs) [18]. We present a matching upper bound.

Theorem 4.6. *There exists a $\frac{4}{3}$ -robust algorithm for speed-robust scheduling with $\{0, 1\}$ -speeds and unit-size jobs.*

In the proof, we handle different cases depending on m and OPT_m by carefully tailored methods. (Recall, OPT_m denotes the optimal makespan on m machines.) Again, we use λ to denote the average load per bag, i.e., $\lambda := \frac{n}{m}$.

For $\text{OPT}_m \geq 11$, we use the algorithm SANDFORBRICKS based on SAND₀₁; see Corollary 4.8. For $\text{OPT}_m \in \{9, 10\}$, we refine this approach and show that for $m \geq 40$ it is still possible to use bag sizes based on SAND₀₁; see Lemma 4.9. For $\text{OPT}_m \in \{3, \dots, 8\}$ and $m \geq 50$, we explicitly give a packing in Lemma 4.11 while for $\text{OPT}_m \in \{1, 2\}$ and $m \geq 50$ packing bags according to the optimal schedule on m machines is sufficient; see Lemma 4.12. The remaining cases, $\text{OPT}_m \leq 10$ and $m \leq 50$, can be verified by enumerating all possible instances and using an integer linear program to verify that there is a solution of bag sizes that is $\frac{4}{3}$ -robust; see Lemma 4.13.

While proving these results, inductively applying the following lemma allows us to restrict ourselves to instances where $\text{OPT}_m < \text{OPT}_{m-1}$. Hence, we can express $n = m\text{OPT}_m - \ell$, where $0 \leq \ell < \min\{m, \text{OPT}_m\}$.

Lemma 4.7. *Fix a job set I and some ρ . If $\text{OPT}_m = \text{OPT}_{m-1}$, then any solution for I that is ρ -robust for $m - 1$ bags on $m - 1$ machines is ρ -robust in the instance on m machines as well.*

Proof. Let I be the set of jobs. Compute a solution with $m - 1$ bags that is ρ -robust for the instance with jobs I and $m - 1$ bags. If $m' < m$ machines actually work, i.e., $m - m'$ machines fail, return the schedule computed by the ρ -robust algorithm with $m - 1$ bags and m' machines. If m machines work, assign each of the $m - 1$ bags to its private machine. By assumption, the load of the largest bag is at most $\rho\text{OPT}_{m-1} = \rho\text{OPT}_m$ which gives the result. \square

Corollary 4.8. *If $\text{OPT}_m \geq 11$, SANDFORBRICKS based on SAND₀₁ is $\frac{4}{3}$ -robust.*

Proof. This follows directly from Lemma 4.1 and the fact that, for $\text{OPT}_m \geq 11$, we have $\lambda = \frac{n}{m} \geq 10$. Thus, the robustness ratio of SANDFORBRICKS is at most $(1 + \frac{1}{\lambda})\bar{\rho}_{01} \leq 1.1 \cdot \bar{\rho}_{01} < 1.33$ which is less than $4/3$. \square

Lemma 4.9. *For $\text{OPT}_m \in \{9, 10\}$ and $m \geq 50$, there is a $\frac{4}{3}$ -robust algorithm.*

Proof. Consider bags created by SAND₀₁ and scale them by a factor of $\frac{4}{3\bar{\rho}_{01}}$. We obtain bag sizes $a'_i = \frac{n}{m} \cdot \min\{\frac{4}{3}, \frac{2}{3\bar{\rho}_{01}} + \frac{4}{3} \cdot \frac{i-\frac{1}{2}}{m}\}$, for $i \in \{1, \dots, m\}$. Round down a'_i to the nearest integer and denote the rounded bag size by a_i .

The total volume of bags before rounding is $\sum_{i=1}^m a'_i = \frac{4}{3\bar{\rho}_{01}}n > n$ and is, thus, larger than the total processing volume of all jobs. We will show that after rounding down the bag sizes to the nearest integer, the remaining volume is still at least n . Therefore, we can guarantee that all unit-size jobs can be assigned to the bags, so the robustness factor is not larger than the robustness factor $\bar{\rho}_{01}$ of SAND₀₁ times the scaling factor $4/(3\bar{\rho}_{01})$, which proves the lemma.

To argue that the volume that remains after rounding is at least n , we show that the loss of volume due to rounding is bounded by the term $(\frac{4}{3\bar{\rho}_{01}} - 1)n$. We do this by carefully analyzing the loss incurred on three different portions of bags that correspond to different parts of the profile function

$\bar{f}(x) = \frac{n}{m} \cdot \min \left\{ \frac{4}{3}, \frac{2}{3\beta_{01}} + \frac{4}{3} \cdot x \right\}$, which, too, is obtained by scaling the profile function of SAND₀₁. Denote by ω the order of $4n$ in the additive group \mathbb{Z}_{3m^2} , that is, $\omega = \min\{i \in \mathbb{N} \setminus \{0\} \mid i \cdot 4n \equiv 0 \pmod{3m^2}\}$. We claim:

- (i) On the *plateau*, that is, for all bags with $a'_i = \frac{n}{m} \cdot \frac{4}{3}$, we loose a volume of at most $L_{\max} := 1 - \frac{1}{3m}$ per bag.
- (ii) On the remaining portion, i.e., the *slope*, we loose for any ω consecutive bags an average volume of at most $\frac{1}{2} + \frac{4n-1}{3m^2}$ per bag.
- (iii) On *leftover bags* of the slope, i.e., bags that remain after partitioning bags on the slope into ω -sized groups, we loose a total volume of $\frac{3m^2}{32n} + \frac{n}{6m^2} + \frac{1}{4}$ additional to the average loss of $\frac{1}{2} + \frac{4n-1}{3m^2}$ per bag from (ii).

See Figure 3 for an illustration of the different parts.




Figure 3: Bags a'_i obtained from SAND₀₁ by scaling (striped) and bags a_i rounded down to the nearest integer for $m = 30$ and $n = 270$. The plateau is indicated in dark grey, the sloped part in mid grey, and the leftover bags in light grey. The continuous sloped line indicates the function \bar{f} with the appropriate scaling. Note that, here, $\omega = 5$, and the volume of the bags a_i is $288 \geq n$.

The first case, Claim (i), follows directly from the fact that bag sizes are $\frac{4n}{3m}$.

For Claim (ii), note that since the slope of the profile function is $\Delta := \frac{4n}{3m^2}$, the loss due to rounding is periodic. Specifically, it has a period of ω . For ω consecutive bags, and starting with a worst-case loss of L_{\max} for the first bag, the average loss due to rounding equals

$$\frac{1}{\omega} \sum_{i=0}^{\omega-1} \left(L_{\max} - \frac{i}{\omega} \right) = 1 - \frac{1}{3m^2} - \frac{1}{2} \frac{\omega - 1}{\omega} = 1 - \frac{1}{2} - \frac{1}{3m^2} + \frac{1}{2\omega}.$$

Denote this term by L_ω . Since $\omega \geq \frac{3m^2}{4n}$, the average loss is at most $\frac{1}{2} + \frac{4n-1}{3m^2}$.

Finally, for Claim (iii), we construct adversarial leftover bags that maximize the loss due to rounding. The first bag has a loss of L_{\max} , the second a loss of $L_{\max} - \Delta$, and so on until the last one has a loss which is just above L_ω . Adding further leftover bags would lead to averaging as in Claim (ii) and ultimately a smaller rounding loss. For such adversarial leftover bags, we define $n_\ell := \frac{L_{\max} - L_\omega}{\Delta}$, so $\lfloor n_\ell \rfloor$ is the number of leftover bags. Then, the overall additional loss, when

compared to L_ω , is given by

$$\begin{aligned} \sum_{i=0}^{\lfloor n_\ell \rfloor - 1} ((L_{\max} - L_\omega) - i \cdot \Delta) &\leq \lfloor n_\ell \rfloor (L_{\max} - L_\omega) - \frac{\lfloor n_\ell \rfloor (n_\ell - 2)}{2} \cdot \Delta \\ &\leq \frac{n_\ell (L_{\max} - L_\omega)}{2} + n_\ell \cdot \Delta \leq \frac{(\frac{1}{2} - \frac{1}{2\omega})^2}{2\Delta} + \frac{1}{2}. \end{aligned}$$

Again, we use $\omega \geq \frac{3m^2}{4n}$ to obtain an upper bound for the last term. Then

$$\sum_{i=0}^{\lfloor n_\ell \rfloor - 1} ((L_{\max} - L_\omega) - i \cdot \Delta) \leq \frac{(\frac{1}{2} - \frac{2n}{3m^2})^2}{\frac{8n}{3m^2}} + \frac{1}{2} = \frac{3m^2}{32n} + \frac{n}{6m^2} + \frac{1}{4},$$

which concludes the proof of Claim (iii).

Let $\lceil x \rceil$ denote the value of x rounded to the closest integer. We can now bound the overall loss due to rounding by

$$[(1 - \beta)m] \cdot L_{\max} + [\beta m] \cdot L_\omega + \left(\frac{3m^2}{32n} + \frac{n}{6m^2} + \frac{1}{4} \right).$$

Using $(1 - \beta)m$ and βm instead of rounding to the nearest integer decreases this term by at most $\frac{1}{2}$ since m is an integer and $L_{\max}, L_\omega \leq 1$. Thus the total loss due to rounding is less or equal to $(\frac{4}{3\bar{\rho}_{01}} - 1)n$, if

$$\left(\frac{4}{3\bar{\rho}_{01}} - 1 \right) n - (1 - \beta)m - \beta m \cdot \left(\frac{1}{2} + \frac{4n - 1}{3m^2} \right) - \frac{3m^2}{32n} - \frac{n}{6m^2} - \frac{1}{4} \geq 0,$$

Using $m\text{OPT}_m - \text{OPT}_m \leq n \leq m\text{OPT}_m$, algebraic computations show that for $\text{OPT}_m = 9$ and $\text{OPT}_m = 10$ this is the case when $m \geq 40$ and $m \geq 30$, respectively (see [9]). \square

When $\text{OPT}_m \leq 8$, we only require constantly many different bag sizes, which we describe explicitly. To simplify the analysis, we assume that, once the number m' of non-failing machines is revealed, bags are assigned to the machines in LPT order. Since this assignment cannot be better than the optimal bag-to-machine assignment, the ratio between the makespan attained by LPT and the optimal solution is not smaller than the robustness factor.

The following observation on bags assigned to machines by LPT is crucial. For a given set of bags, let $\text{LPT}_{m'}$ be the makespan attained by assigning the bags in LPT order to the currently least loaded machine when there are m' machines with speed $s_i = 1$ (and $m - m'$ machines with speed $s_i = 0$).

Lemma 4.10. *Let a be the size of a bag determining the makespan of LPT on m' machines. If $a \leq \frac{\text{OPT}_{m'}}{3}$, then $\text{LPT}_{m'} \leq \frac{4}{3}\text{OPT}_{m'}$.*

Proof. Let b be a bag of size a that determines the makespan and is placed on machine i . As LPT assigns bags in decreasing size to the currently least loaded machine, the load on machine i right before assigning b was at most $\text{OPT}_{m'}$. Hence, it follows that $\text{LPT}_{m'} \leq \text{OPT}_{m'} + a \leq \frac{4}{3}\text{OPT}_{m'}$. \square

For $\text{OPT}_m \in \{3, \dots, 8\}$, we pack four different types of bags depending on OPT_m . For $l \in \{0, 1, 2, 3\}$, we denote by a_l the size of the l -th bag type and by x_l its multiplicity. The idea is to have $x_0 + x_1 = x_3 \approx \frac{2}{5}m$ while $x_2 \approx \frac{1}{5}m$.

More precisely, let $a_1 = \lceil \frac{2}{3} \text{OPT}_m \rceil$, $a_2 = \text{OPT}_m$, and $a_3 = \lfloor \frac{4}{3} \text{OPT}_m \rfloor$ be the three *standard* bag sizes. Since $a_1 + a_3 = 2a_2$, packing as many smallest as largest bags, i.e., $x_1 = x_3$, ensures that $\sum_{l=1}^3 a_l x_l = m \text{OPT}_m$. Recall that $n = m \text{OPT}_m - \ell$ with $0 \leq \ell < \text{OPT}_m$. Hence, we decrease x_1 by ℓ and pack $x_0 = \ell$ many bags of size $a_0 = a_1 - 1 = \lceil \frac{2}{3} \text{OPT}_m \rceil - 1$ in order to pack exactly n jobs in our m bags. As we aim for a tight robustness guarantee, we have to be careful about the exact number of bags in this section. The following table defines $x_0 + x_1$, x_2 , and x_3 depending on $m \pmod{5}$.

$m \pmod{5}$	0	1	2	3	4	a_l
$x_0 + x_1$	$\frac{2m}{5}$	$2 \lfloor \frac{m}{5} \rfloor$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$\lceil \frac{2}{3} \text{OPT}_m \rceil$
x_2	$\frac{m}{5}$	$\lfloor \frac{m}{5} \rfloor + 1$	$\lfloor \frac{m}{5} \rfloor$	$\lfloor \frac{m}{5} \rfloor + 1$	$\lfloor \frac{m}{5} \rfloor + 2$	OPT_m
x_3	$\frac{2m}{5}$	$2 \lfloor \frac{m}{5} \rfloor$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$\lfloor \frac{4}{3} \text{OPT}_m \rfloor$

The analysis of assigning these bags via LPT is a tedious case distinction based on which types of bags are assigned to the same machine and does not provide new insights. Therefore, we defer it to Appendix C.

Lemma 4.11. *If $\text{OPT}_m \in \{3, \dots, 8\}$ and $m \geq 50$, there is a $\frac{4}{3}$ -robust algorithm.*

Lemma 4.12. *If $\text{OPT}_m \leq 2$ and $m \geq 37$, there exists a $\frac{4}{3}$ -robust algorithm.*

Proof. If $\text{OPT}_m = 1$, each job gets a unique bag. Hence, $\text{LPT}_{m'} = \text{OPT}_{m'}$.

For $\text{OPT}_m = 2$, if no machine fails, the packing achieves a makespan of OPT_m . By Lemma 4.7, we may assume that $n \in \{2m, 2m - 1\}$. Therefore, if $\frac{m}{2} \leq m' \leq m - 1$, we have $\text{OPT}_{m'} \geq 3$ while $\text{LPT}_{m'} \leq 4$ since at most 2 bags are assigned to the same machine. For the remaining cases, with $m' \leq \frac{m}{2} - 1$, we have $\text{OPT}_{m'} \geq 5$. When $\text{OPT}_{m'} = 5$, LPT assigns at most 3 bags to each machine which guarantees that $\text{LPT}_{m'} \leq 6 \leq \frac{4}{3} \text{OPT}_{m'}$. If $\text{OPT}_{m'} \geq 6$, Lemma 4.10 implies that $\text{LPT}_{m'} \leq \frac{4}{3} \text{OPT}_{m'}$. \square

Lemma 4.13. *For $\text{OPT}_m \leq 10$ and $m \leq 50$, there exists a $\frac{4}{3}$ -robust algorithm.*

Proof. We enumerate all instances with $\text{OPT}_m \neq \text{OPT}_{m-1}$ for which the properties of the statement hold. These instances consist of $n = m \text{OPT}_m - \ell$ unit-size jobs, with $\text{OPT}_m \in \{1, 2, \dots, 10\}$ and $\ell \in \{0, 1, \dots, \min(\text{OPT}_m - 1, m - 1)\}$. We solve each such instance by computing a feasible solution for an integer linear program (ILP) which we now describe. Denote by p_{\max} the largest possible size of a bag. To be able to achieve a robustness of $\frac{4}{3}$, we must have $p_{\max} = \lfloor \frac{4}{3} \text{OPT}_m \rfloor$. The decision variables used are y_p and $x_{p,i,m'}$ to indicate how many bags of size p are created and how many of these bags of size p are assigned to machine i when m' machines are working, respectively. The ILP is as follows.

$$\begin{aligned}
\sum_{p=1}^{p_{\max}} y_p &= m \\
\sum_{p=1}^{p_{\max}} p \cdot y_p &= n \\
\sum_{i=1}^{m'} x_{p,i,m'} &= y_p && \text{for all } m' \in [m], p \in [p_{\max}] \\
\sum_{p=1}^{p_{\max}} p \cdot x_{p,i,m'} &\leq \frac{4}{3} \cdot \text{OPT}_{m'} && \text{for all } m' \in [m], i \in [m'] \\
x_{p,i,m'} &\in \mathbb{Z}_{\geq 0} && \text{for all } p \in [p_{\max}], m' \in [m], i \in [m'] \\
y_p &\in \mathbb{Z}_{\geq 0} && \text{for all } p \in [p_{\max}]
\end{aligned}$$

The first equation ensures that exactly m bags are created, and the second that, in total, they consist of exactly n jobs. The third equation enforces that, for every value of m' , all bags are assigned to a machine. Finally, the fourth equation checks that, for every value of m' and for every remaining machine, the makespan of the optimal solution is not exceeded by more than a factor of $\frac{4}{3}$, yielding the robustness guarantee. The last two constraints dictate integrality and non-negativity of the decision variables.

A file containing the solutions produced by this ILP and a simple program verifying them are available in [9]. \square

5 Concluding Remarks

In this work, we have established matching lower and upper bounds for the speed-robust scheduling problem with infinitesimal jobs and design optimal algorithms for either infinitesimal jobs or equal-size jobs when speeds are restricted to $\{0, 1\}$. We believe that the insights from our optimal algorithms will be useful to improve the more general upper bounds.

We have also shown that randomization does not help when the speeds belong to $\{0, 1\}$ and jobs are infinitesimal. However, the other known lower bounds do not hold in a randomized setting, so designing better randomized algorithms remains an interesting challenge.

The following observation about adversarial strategies might be useful for further research. We give two somewhat orthogonal examples proving the lower bound of $\frac{4}{3}$ for speed-robust scheduling with unit processing time jobs and speeds from $\{0, 1\}$. In both examples, there are only two relevant adversarial strategies: either one machine fails or none. This may seem sub-optimal, but the lower bound of $\frac{4}{3}$ is tight for unit-size jobs (Theorem 4.6). Further, we show in the proof of Theorem 2.4 (Appendix A) that, for infinitesimal jobs, an adversary only requires two strategies to force all algorithms to have a robustness factor at least $\bar{\rho}_{01}(m)$, which is optimal.

Example 1 (from [18]). Consider $2m$ jobs and $m > 2$ machines. If an algorithm places 2 jobs per bag, let one machine fail. This leads to a makespan of 4 while the optimal makespan is 3 which gives a lower bound of $\frac{4}{3}$. Otherwise, one bag has at least three jobs, and, if no machine fails, the algorithm's makespan is 3 while the optimal makespan is 2, yielding a lower bound of $\frac{3}{2}$.

Example 2. Our new dual example has $3m$ jobs for $m > 3$ machines. If an algorithm places 3 jobs per bag, let one machine fail. This leads to a makespan of 6 while the optimal makespan is 4, implying a lower bound of $\frac{3}{2}$. Otherwise, one bag has at least 4 jobs, and if no machine fails, the algorithm's makespan is 4 whereas the optimal makespan is 3, which again gives a lower bound of $\frac{4}{3}$.

We conclude with a remark on a natural generalization of our model in which we allow to pack more bags than the number of machines, say $M \geq m$. Our upper bounds obviously still hold by using the algorithms with m bags as presented. One would expect that the approximation ratio improves when M/m increases, and it would be interesting to quantify the achievable robustness factor in terms of both M and m . Indeed, for the general problem with rocks and arbitrary speeds, it is not difficult to see that Theorem 3.3 can be generalized by parameterization and adjusting the inequalities in the proof (see Appendix D). Packing $M \geq m$ bags by LPT yields a robustness factor of $1 + \frac{m-1}{M}$, which interpolates nicely between $2 - \frac{1}{m}$ (for $M = m$) and 1 (for $M \rightarrow \infty$).

Concerning our lower bounds, we expect Theorems 2.1 and 2.8 to become invalid for $M > m$ as these results are tight for $M = m$ and an algorithm should be able, with infinitesimal jobs, to exploit any additional bag to reduce the competitive ratio. However, the lower bound of $4/3$ for unit-time jobs and speeds in $\{0, 1\}$ (both examples above) has some slack and therefore holds when M is moderately larger than m . Specifically, assuming for simplicity that 3 divides m , one can easily check that for an instance with $M = \frac{4}{3}m - 1$ bags and $2m$ unit jobs, no algorithm can be

better than $4/3$ -competitive (the worst cases being 0 and $\frac{m}{3}$ machine failures). Quantifying exactly how all results in this paper evolve when $M > m$ remains an open problem.

References

- [1] S. Albers and M. Hellwig. Online makespan minimization with parallel schedules. *Algorithmica*, 78(2):492–520, 2017. doi: 10.1007/s00453-016-0172-5.
- [2] S. Albers and G. Schmidt. Scheduling with unexpected machine breakdowns. *Discret. Appl. Math.*, 110(2-3):85–99, 2001. doi: 10.1016/s0166-218x(00)00266-3.
- [3] N. Alon, Y. Azar, G. J. Woeginger, and T. Yadid. Approximation schemes for scheduling on parallel machines. *Journal of Scheduling*, 1(1):55–66, 1998. doi: 10.1002/(sici)1099-1425(199806)1:1<55::aid-jos2>3.0.co;2-j.
- [4] S. K. Baruah, V. Bonifaci, G. D’Angelo, H. Li, A. Marchetti-Spaccamela, N. Megow, and L. Stougie. Scheduling real-time mixed-criticality jobs. *IEEE Trans. Computers*, 61(8):1140–1152, 2012. doi: 10.1109/tc.2011.142.
- [5] L. Chen, N. Megow, R. Rischke, L. Stougie, and J. Verschae. Optimal algorithms for scheduling under time-of-use tariffs. *Annals of Operations Research*, 2021. doi: 10.1007/s10479-021-04059-3.
- [6] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. *Commun. ACM*, 51(1):107–113, Jan. 2008. doi: 10.1145/1327452.1327492.
- [7] F. Diedrich, K. Jansen, U. M. Schwarz, and D. Trystram. A survey on approximation algorithms for scheduling with machine unavailability. In *Algorithmics of Large and Complex Networks*, volume 5515 of *Lecture Notes in Computer Science*, pages 50–64. Springer, 2009. doi: 10.1007/978-3-642-02094-03.
- [8] C. Dürr, T. Erlebach, N. Megow, and J. Meißner. An adversarial model for scheduling with testing. *Algorithmica*, 82(12):3630–3675, 2020. doi: 10.1007/s00453-020-00742-2.
- [9] F. Eberle, R. Hoeksma, N. Megow, L. Nölke, K. Schewior, and B. Simon. Speed robust scheduling supplemental material, Dec 2021.
- [10] L. Epstein, A. Levin, A. Marchetti-Spaccamela, N. Megow, J. Mestre, M. Skutella, and L. Stougie. Universal sequencing on an unreliable machine. *SIAM J. Comput.*, 41(3):565–586, 2012. doi: 10.1137/110844210.
- [11] D. S. Hochbaum and D. B. Shmoys. Using dual approximation algorithms for scheduling problems theoretical and practical results. *J. ACM*, 34(1):144–162, 1987. doi: 10.1145/7531.7535.
- [12] K. Jansen. An EPTAS for scheduling jobs on uniform processors: Using an MILP relaxation with a constant number of integral variables. *SIAM J. Discrete Math.*, 24(2):457–485, 2010. doi: 10.1137/090749451.
- [13] P. Kouvelis and G. Yu. *Robust Discrete Optimization and Its Applications*. Springer, 1997. doi: 10.1007/978-1-4757-2620-6.

- [14] J. Niño-Mora. Stochastic scheduling. In *Encyclopedia of Optimization*, pages 3818–3824. Springer, 2009. doi: 10.1007/978-0-387-74759-0665.
- [15] M. L. Pinedo. *Scheduling: Theory, Algorithms, and Systems*. Springer Publishing Company, Incorporated, 4th edition, 2012.
- [16] K. Pruhs, J. Sgall, and E. Torng. Online scheduling. In *Handbook of Scheduling*. Chapman and Hall/CRC, 2004. doi: 10.1007/978-3-319-99849-7.
- [17] D. B. Shmoys and M. Sozio. Approximation algorithms for 2-stage stochastic scheduling problems. In *IPCO*, volume 4513 of *Lecture Notes in Computer Science*, pages 145–157. Springer, 2007. doi: 10.1007/978-3-540-72792-712.
- [18] C. Stein and M. Zhong. Scheduling when you do not know the number of machines. *ACM Trans. Algorithms*, 16(1):9:1–9:20, 2020. doi: 10.1145/3340320.

Appendices

A Proofs for Section 2.2 – Infinitesimal Jobs and Speeds in $\{0, 1\}$

In the following, we establish the upper bound of $\bar{\rho}_{01}(m)$ when speeds are in $\{0, 1\}$, which dominates the more coarse bound of $\bar{\rho}_{01}$ from Theorem 2.7 but also requires substantially more work.

Theorem 2.4. *For all $m \geq 1$, there is a deterministic $\bar{\rho}_{01}(m)$ -robust algorithm for speed-robust scheduling with speeds in $\{0, 1\}$ for infinitesimal jobs.*

Proof. We give bag sizes that guarantee a robustness factor of $\bar{\rho}_{01}(m)$ for every m . While the load distribution in the limit approaches that given by \bar{f} , we do not work with \bar{f} explicitly anymore. We fix $m \geq 3$ in the following; the other cases are trivial. Furthermore, let again

$$t^* \in \arg \max_{t \leq \frac{m}{2}, t \in \mathbb{N}} \frac{1}{\frac{t}{m-t} + \frac{m-2t}{m}}.$$

To show the theorem, we distinguish two cases: $t^* \leq 1$ and $t^* > 1$.

We start with $t^* \leq 1$ and show that this implies $m \in \{3, 4, 5\}$. First, note that the above expression defining t^* equals 1 for $t = 0$ and $t = m/2$ and is larger for values of t in between. As we assumed $m \geq 3$, we have that $m/2 \geq 1$, and thus $t^* \geq 1$. Hence, for $m \in \{3, 4\}$, we have $t^* = 1$.

For $m \geq 5$, in order to have $t^* = 1$, it is necessary that $\frac{1}{\frac{t}{m-t} + \frac{m-2t}{m}}$ is larger at $t = 1$ than at $t = 2$. Noting that

$$\frac{1}{\frac{t}{m-t} + \frac{m-2t}{m}} = \frac{m^2 - tm}{m^2 - 2tm + 2t^2}$$

and reformulating, we obtain the necessary condition

$$\frac{m^2 - 2m}{m^2 - 4m + 8} - \frac{m^2 - m}{m^2 - 2m + 2} < 0.$$

The left-hand side is equal to

$$\frac{m(m-2)}{(m-2)^2 + 4} - \frac{m^2 - m}{m(m-2) + 2} = \frac{m(m^2 - 6m + 4)}{((m-2)^2 + 4) \cdot (m(m-2) + 2)}.$$

As $m \geq 5$, this term is negative if and only if $m(m^2 - 6m + 4) < 0$. Since the roots of this polynomial are 0 , $3 - \sqrt{5}$, and $3 + \sqrt{5}$, the expression is strictly negative if $m \leq 5$ and strictly positive for $m \geq 6$, which implies $t^* \geq 2$ in the latter case.

It remains to consider $m = 5$. By the above calculations we know that $t^* = 1$. Therefore, $t^* \leq 1$ implies $m \in \{3, 4, 5\}$. For these three cases, the optimal bag sizes are as follows:

- For $m = 3$, the bag sizes are $\{0.9, 0.9, 1.2\}$ and $\bar{\rho}_{01}(m) = 1.2$.
- For $m = 4$, the bag sizes are $\{0.8, 0.8, 1.2, 1.2\}$ and $\bar{\rho}_{01}(m) = 1.2$.
- For $m = 5$, the bag sizes are $\{\frac{25}{34}, \frac{25}{34}, \frac{40}{34}, \frac{40}{34}, \frac{40}{34}\}$ and $\bar{\rho}_{01}(m) = 40/34$.

It can be easily verified that, if at most $m/2$ machines fail, the obtained makespan is at most $\bar{\rho}_{01}(m)$.

Consider now the case that $t^* \geq 2$. By the discussion above, this implies $m \geq 6$. We also claim that there exist bag sizes that achieve a robustness factor of $\bar{\rho}_{01}(m)$. It turns out that, for many different values of m , these sizes are not unique. We impose additional constraints on the bag sizes




Figure 4: The structure of the optimal solution of Theorem 2.4 illustrated for $m = 32$.

to get bag sizes that are easier to analyze, giving some intuition along the way for why the imposed restrictions do not remove all bag sizes achieving $\bar{\rho}_{01}(m)$. When we are left with a single degree of freedom, we impose lower and upper bounds on the corresponding variable so that fulfilling these bounds implies the robustness guarantee $\bar{\rho}_{01}(m)$. We then show that the largest lower bound does not exceed the smallest upper bound, implying that there is a feasible choice for said variable.

As observed in Section 2.2, if we want to guarantee a robustness factor of $\bar{\rho}_{01}(m)$, we need to fix $a_{2t^*+1}, \dots, a_m = \bar{\rho}_{01}(m)$; so, this restriction does not remove any set of bag sizes that achieves $\bar{\rho}_{01}(m)$. Since the optimal factor could be achieved by a set of bag sizes that achieves this factor with the simple folding, we also assume such bag sizes here. As before, this implies $a_i + a_{2t^*+1-i} = \bar{\rho}_{01}(m) \cdot m / (m - t^*)$ for all $i = 1, \dots, t^*$.

Note that, when considering a bag $i \leq t$ for any $t < m/2$ and keeping track of the bag which i is matched up with while increasing t , odd bags are always folded on top of even bags and vice versa. That motivates considering bags in pairs and thinking of pairs that get matched up rather than particular bags. Specifically, we choose $a_{2i-1} = a_{2i}$ for $i = 1, \dots, t^*$.

To imitate the increasing linear part of \bar{f} , we further impose $a_{2i+1} = a_{2i-1} + \delta$ for some $\delta \geq 0$ and all $i = 1, \dots, t^* - 1$. Since \bar{f} has two different slopes, similar to our previous set of bag sizes, the difference between $\bar{\rho}_{01}(m)$ and the largest bag that is smaller than $\bar{\rho}_{01}(m)$ may differ from δ ; we call this value $\delta' := a_{2t^*+1} - a_{2t^*}$. In fact, it can be shown that $\delta' \neq \delta$ is even necessary given our previous assumption on a_i for some values of m . Figure 4 visualizes this structure for $m = 32$. Note that Figure 1 visualizes folding on the structure in the $m = 20$ case.

Note that there is a single degree of freedom left if we want to define bags of total volume precisely m : If we choose δ , then this fixes δ' , and vice versa. The goal is to show that the set that we can choose δ from, so as to guarantee $\bar{\rho}_{01}(m)$, is nonempty. In what follows, we derive lower and upper bounds on δ . Fulfilling these bounds implies a robustness factor of $\bar{\rho}_{01}(m)$.

To make our computations simpler, we assume that $\delta \geq \delta'$. This yields our first lower bound L_1 on δ . We express all bounds in terms of m , t^* , and $\bar{\rho}_{01}(m)$. To compute δ' from these values, observe that the total size of bags a_1, \dots, a_{2t^*} is $m - \bar{\rho}_{01}(m) \cdot (m - 2t^*)$ using that the total volume is m and that bags a_{2t^*+1}, \dots, a_m have size $\bar{\rho}_{01}(m)$. Dividing by $2t^*$ yields the average bag size of

bags a_1, \dots, a_{2t^*} , denoted by

$$\bar{a} = \frac{m - \bar{\rho}_{01}(m) \cdot (m - 2t^*)}{2t^*}.$$

From \bar{a} it takes $(t^* - 1)/2$ steps of size δ to get to a_{2t^*} . Arguing in terms of volume, which allows arguing in terms of half steps, we get: If $\delta' \leq \delta$, then we have $\bar{\rho}_{01}(m) = a_{2t^*} + \delta' \leq \bar{a} + \delta(t^* - 1)/2 + \delta = \bar{a} + \delta(t^* + 1)/2$. This is equivalent to

$$\begin{aligned} \delta &\geq \frac{2}{t^* + 1} \cdot (\bar{\rho}_{01}(m) - \bar{a}) = \frac{2}{t^* + 1} \cdot \left(\bar{\rho}_{01}(m) - \frac{m - \bar{\rho}_{01}(m) \cdot (m - 2t^*)}{2t^*} \right) \\ &= \frac{m(\bar{\rho}_{01}(m) - 1)}{t^*(t^* + 1)} =: L_1. \end{aligned} \quad (3)$$

Now, we give two bounds on δ that ensure that the bag sizes are feasible in that we have $a_i \in [0, \bar{\rho}_{01}(m)]$ for all i . Using that $a_{2t^*} = \bar{a} + (t^* - 1)/2 \cdot \delta$ and imposing that $a_{2t^*} \leq \bar{\rho}_{01}(m)$, i.e., that $\delta' = a_{2t^*+1} - a_{2t^*} \geq 0$, yields $\bar{\rho}_{01}(m) \geq \bar{a} + (t^* - 1)\delta/2$. Hence,

$$\begin{aligned} \delta &\leq \frac{2}{t^* - 1} \cdot (\bar{\rho}_{01}(m) - \bar{a}) = \frac{2}{t^* - 1} \cdot \left(\bar{\rho}_{01}(m) - \frac{m - \bar{\rho}_{01}(m) \cdot (m - 2t^*)}{2t^*} \right) \\ &= \frac{m(\bar{\rho}_{01}(m) - 1)}{t^*(t^* - 1)} =: U_1. \end{aligned} \quad (4)$$

We also need $a_1 \geq 0$, which we impose by letting the sum of the increments not exceed $\bar{\rho}_{01}(m)$. This yields an upper bound on δ of

$$U_2 := \frac{\bar{\rho}_{01}(m)}{t^* - 1}. \quad (5)$$

For the upper and lower bounds ensuring the robustness, recall that $a_i + a_{2t^*+1-i} = \bar{\rho}_{01}(m) \cdot \frac{m}{m-t^*}$ for all $i = 1, \dots, t^*$. That is, the robustness ratio is attained *exactly* when $t = t^*$. This implies that the ratio of the increase in the algorithm's cost and the increase the optimum's cost when the number t of failing machines is increased from t^* to, say, $t^* + k$, should be at most $\bar{\rho}_{01}(m)$. To compute the change in the algorithm's cost, note that, using (3), *any* bag size a_i can be bounded from above by $a_1 + \lfloor i/2 \rfloor \cdot \delta$. That implies that the algorithm's cost increases by at most $k\delta$. On the other hand the optimum's cost changes from $\frac{m}{m-t^*}$ to $\frac{m}{m-t^*-k}$. So the aforementioned ratio is

$$\frac{k \cdot \delta}{\frac{m}{m-t^*-k} - \frac{m}{m-t^*}} = \frac{\delta \cdot (m - t^*)(m - t^* - k)}{m},$$

yielding an upper bound on δ of

$$U_3^k := \frac{m \cdot \bar{\rho}_{01}(m)}{(m - t^*)(m - t^* - k)}.$$

This upper bound is minimized for $k = 1$, so we only need to remember $U_3 := U_3^1$. Similarly, we consider the case when t is decreased from t^* to $t^* - k$. Then the algorithm's cost decreases by *precisely* $k\delta$ and that of the optimum by $\frac{m}{m-t^*} - \frac{m}{m-t^*+k}$. Imposing that the ratio of these quantities is *at least* $\bar{\rho}_{01}(m)$ yields

$$\delta \geq \frac{m \cdot \bar{\rho}_{01}(m)}{(m - t^*)(m - t^* + k)} =: L_2^k.$$

Again, this bound is *maximized* for $k = 1$, so we only need to remember $L_2 := L_2^1$.

With these lower and upper bounds on δ , we can complete the proof. It boils down to showing that the interval $[\max\{L_1, L_2\}, \min\{U_1, U_2, U_3\}]$ is nonempty because then we can choose δ from that

interval and thereby define bag sizes with a robustness of $\bar{\rho}_{01}$. To achieve this, we first rewrite $\bar{\rho}_{01}(m)$ as

$$\bar{\rho}_{01}(m) = \frac{1}{\frac{t^*}{m-t^*} + \frac{m-2t^*}{m}} = \frac{m(m-t^*)}{m^2 - 2mt^* + 2(t^*)^2}. \quad (6)$$

We use (6) to expand the lower and upper bounds as follows:

$$\begin{aligned} L_1 &= \frac{m(\bar{\rho}_{01}(m) - 1)}{t^*(t^* + 1)} = \frac{m(m - 2t^*)}{(t^* + 1)(m^2 - 2mt^* + 2(t^*)^2)} \\ L_2 &= \frac{m\bar{\rho}_{01}(m)}{(m - t^*)(m - t^* + 1)} = \frac{m^2}{(m - t^* + 1)(m^2 - 2mt^* + 2(t^*)^2)}, \\ U_1 &= \frac{m(\bar{\rho}_{01}(m) - 1)}{t^*(t^* - 1)} = \frac{m(m - 2t^*)}{(t^* - 1)(m^2 - 2mt^* + 2(t^*)^2)}, \\ U_2 &= \frac{\bar{\rho}_{01}(m)}{t^* - 1} = \frac{m(m - t^*)}{(t^* - 1)(m^2 - 2mt^* + 2(t^*)^2)}, \\ U_3 &= \frac{m\bar{\rho}_{01}(m)}{(m - t^*)(m - t^* - 1)} = \frac{m^2}{(m - t^* - 1)(m^2 - 2mt^* + 2(t^*)^2)}. \end{aligned}$$

First notice that U_1 and U_2 differ only by mt^* in the numerator, so U_1 is not greater than U_2 and therefore we can ignore U_2 . Moreover, when comparing L_2 with U_3 and L_1 with U_1 , we observe that the numerators do not differ while the denominators are smaller in the upper bounds. Therefore, it is immediate that $L_2 \leq U_3$ and $L_1 \leq U_1$.

For the remaining two comparisons, $L_2 \leq U_1$ and $L_1 \leq U_3$, we use that $t = t^*$ maximizes the expression

$$g(t) := \frac{m(m-t)}{m^2 - 2mt + 2t^2} = \frac{m(m-t)}{(m-t)^2 + t^2}.$$

Hence, $g(t^*) - g(t^* - 1) \geq 0$. Therefore, we have that

$$\begin{aligned} 0 &\leq \frac{m(m-t^*)}{(m-t^*)^2 + (t^*)^2} - \frac{m(m-(t^*-1))}{(m-(t^*-1))^2 + (t^*-1)^2} \\ &= m \frac{(m-t^*)^2 + (m-t^*)(-2t^*+2) - (t^*)^2}{((m-t^*)^2 + (t^*)^2)((m-t^*+1)^2 + (t^*-1)^2)}. \end{aligned} \quad (7)$$

Since m and the denominator of (7) are both strictly positive, we obtain that

$$m^2 - 4mt^* + 2m + 2(t^*)^2 - 2t^* \geq 0. \quad (8)$$

Similarly, $g(t^*) - g(t^* + 1) \geq 0$ and we have that

$$\begin{aligned} 0 &\leq \frac{m(m-t^*)}{(m-t^*)^2 + (t^*)^2} - \frac{m(m-(t^*+1))}{(m-(t^*+1))^2 + (t^*+1)^2} \\ &= m \frac{-(m-t^*)^2 + (m-t^*)(2t^*+2) + (t^*)^2}{((m-t^*)^2 + (t^*)^2)((m-t^*-1)^2 + (t^*+1)^2)}. \end{aligned} \quad (9)$$

Again, m and the denominator of (9) are both strictly positive, so we obtain that

$$-m^2 + 4mt^* + 2m - 2(t^*)^2 - 2t^* \geq 0. \quad (10)$$

Now, we show that $L_2 \leq U_1$. This is equivalent to showing that

$$\frac{m}{m-t^*+1} \leq \frac{m-2t^*}{t^*-1} \quad \Leftrightarrow \quad 0 \leq m^2 - 4mt^* + 2m + 2(t^*)^2 - 2t^*,$$

which is true by (8). Finally, we show that $L_1 \leq U_3$. This is equivalent to showing that

$$\frac{m - 2t^*}{t^* + 1} \leq \frac{m}{m - t^* - 1} \quad \Leftrightarrow \quad m^2 - 4mt^* - 2m + 2(t^*)^2 + 2t^* \leq 0,$$

which is true by (10). Thus, the interval $[\max\{L_1, L_2\}, \min\{U_1, U_2, U_3\}]$ is non-empty and there are valid choices for δ . The validity of the upper bound follows. \square

B Proofs for Section 4.1 – Equal-Size Jobs and General Speeds

In the following we show the results of Lemma 4.5 for few machines.

Lemma 4.5.a. *The optimal algorithm for speed-robust scheduling for unit-size jobs has robustness factor $4/3$ on $m = 2$ machines.*

Proof. The lower bound is implied by Theorem 2.1. Let n be the number of jobs of the instance. Consider an algorithm that builds two bags containing at most a_1 and a_2 jobs as follows.

$$a_1 := \left\lfloor \frac{4}{3} \cdot \lceil n/4 + 1 \rceil \right\rfloor \quad ; \quad a_2 := \left\lfloor \frac{4}{3} \lceil n/2 \rceil \right\rfloor \geq \left\lfloor \frac{2}{3}n \right\rfloor.$$

We now show that (i) for every adversary, the algorithm can schedule these bags within a makespan of $4/3$ and (ii) that $a_1 + a_2 \geq n$, so the bags contain all jobs.

The adversary places at least $M_2 := \lceil n/2 \rceil$ jobs on one machine, say machine 2. So the algorithm is always able to place at least a_2 jobs on machine 2. If the adversary places at least $\lceil 3n/4 \rceil$ jobs on machine 2, then the algorithm can place both bags on this machine. Otherwise, the adversary can place at most $\lceil 3n/4 \rceil - 1$ jobs on machine 2. This implies that there are at least $M_1 := n - \lceil 3n/4 \rceil + 1 = \lfloor n/4 \rfloor + 1$ jobs on machine 1. This means that the algorithm can place at least $\lfloor \frac{4}{3}M_1 \rfloor$ jobs on machine 1; this is exactly a_1 . So the makespan achieved by the algorithm is at most $4/3 \cdot C_{\max}^*$, where C_{\max}^* is the optimal makespan.

Hence, the algorithm is $4/3$ -robust if $a_1 + a_2 \geq n$. Consider the four natural integers $s < 4$, $t < 3$, k and q such that $n = 4k + s = 3q + t$. Note that

$$a_1 = \left\lfloor \frac{4}{3} \cdot (k + 1) \right\rfloor = \left\lfloor \frac{4k + 4}{3} \right\rfloor = \left\lfloor \frac{3q + t + (4 - s)}{3} \right\rfloor = q + \left\lfloor \frac{4 + t - s}{3} \right\rfloor.$$

We consider several cases which together complete the proof:

- $4 + t - s \geq 3$, i.e., $t \geq s - 1$: we have $a_1 \geq q + 1 \geq \lceil n/3 \rceil$ so $a_1 + a_2 \geq n$.
- $t = 0$: we have $a_1 \geq q = \lceil n/3 \rceil$ so $a_1 + a_2 \geq n$.
- $0 < t < s - 1 \leq 2$, which means $t = 1$ and $s = 3$: we have $a_1 = q$ and $n = 4k + 3$ is odd, so $M_2 = \frac{1}{2}(n + 1) = \frac{1}{2}(3q + 2)$ and

$$a_2 = \left\lfloor \frac{4}{3}M_2 \right\rfloor = \left\lfloor 2q + \frac{4}{3} \right\rfloor = 2q + 1 = n - a_1.$$

\square

Lemma 4.5.b. *The optimal algorithm for speed-robust scheduling for unit-size jobs has robustness factor $3/2$ on $m = 3$ machines.*

Proof. The lower bound is implied by Lemma 4.4. Let $n > 3$ be the number of jobs. Consider an algorithm that builds three bags as follows:

$$a_1 := \left\lfloor \frac{3}{2} \cdot \frac{1}{2} \left\lfloor \frac{n}{3} + 1 \right\rfloor \right\rfloor \quad ; \quad a_3 := \left\lfloor \frac{3}{2} \left\lfloor \frac{n}{3} \right\rfloor \right\rfloor \geq \lfloor n/2 \rfloor \geq \frac{n-1}{2} \quad ; \quad a_2 := n - a_1 - a_3 .$$

Order the machines in increasing order of adversary load first. The adversary places at least $\lceil \frac{n}{3} \rceil$ jobs on the most loaded machine, machine 3, so the algorithm can always put at least a_3 jobs, i.e., the third bag on machine 3.

If the adversary places at least $\lceil \frac{2}{3}n \rceil$ jobs on machine 3, then the algorithm can put all bags on it and the claim holds. Assume now the adversary places at most $\lceil \frac{2}{3}n \rceil - 1$ jobs on machine 3, so machines 1 and 2 receive at least $\lfloor \frac{n}{3} + 1 \rfloor$ jobs combined. This means in particular that the algorithm can always put bag a_1 on machine 2. We now consider several cases that could prevent the algorithm from reaching a robustness factor of $3/2$. They all implicitly assume that the algorithm cannot simultaneously put the bags a_1 and a_2 on machines 1 and 2 and cannot put the bags a_2 and a_3 jobs on machine 3, as the contrary allows to fit all bags. We therefore show a contradiction in each case.

- (i) The algorithm cannot put a_2 on machine 2. This means than $M_2 < \frac{2}{3}a_2$ and therefore changed from itemize (with dashes) to enumeration

$$M_1 \leq M_2 \leq \frac{2}{3}a_2 - \frac{1}{3} \quad ; \quad M_3 \leq \frac{2}{3}(a_2 + a_3) - \frac{1}{3} .$$

This implies

$$n = M_1 + M_2 + M_3 \leq 2a_2 + \frac{2}{3}a_3 - 1 .$$

Using that $a_2 = n - a_1 - a_3$ and rearranging yields

$$n \leq 2n - 2a_1 - 2a_3 + \frac{2}{3}a_3 - 1 \quad \Leftrightarrow \quad 2a_1 + \frac{4}{3}a_3 - n + 1 \leq 0 .$$

However, letting $n = 3k + t$ with $t \in \{0, 1, 2\}$ and $k > 0$, we have

$$a_1 = \left\lfloor \frac{3k+3}{4} \right\rfloor \geq \frac{3k}{4} \quad ; \quad a_3 \geq \frac{n-1}{2} .$$

This leads to the following contradiction:

$$2a_1 + \frac{4}{3}a_3 - n + 1 \geq \frac{3k}{2} + \frac{2}{3}n - \frac{2}{3} - n + 1 = \frac{3k}{2} - \frac{n}{3} + \frac{1}{3} \geq k + \frac{k}{2} - k - \frac{1}{3} > 0 .$$

- (ii) The algorithm cannot put a_3 on machine 2, a_2 can be put there. Consequently, the algorithm cannot put a_1 on machine 1 as this allows to place simultaneously a_1 and a_2 on machines 1 and 2. So $M_1 < \frac{2}{3}a_1$, which means $M_1 \leq \frac{2}{3}a_1 - \frac{1}{3}$. Similarly, the algorithm cannot put simultaneously a_1 and a_2 on machine 2 nor a_2 and a_3 on machine 3. Therefore, we have

$$M_1 \leq \frac{2}{3}a_1 - \frac{1}{3} \quad ; \quad M_2 \leq \frac{2}{3} \min(a_1 + a_2, a_3) - \frac{1}{3} \quad ; \quad M_3 \leq \frac{2}{3}(a_2 + a_3) - \frac{1}{3} .$$

Noting that $\min(a_1 + a_2, a_3) \leq n/2$ as $a_1 + a_2 + a_3 = n$, we get the following contradiction

$$\begin{aligned} n = M_1 + M_2 + M_3 &\leq \frac{2}{3}(a_1 + a_2 + a_3 + \min(a_1 + a_2, a_3)) - 1 \\ &\leq \frac{2}{3}(n + \frac{1}{2}n) - 1 = n - 1 . \end{aligned}$$

- (iii) The algorithm cannot put $a_1 + a_2$ on machine 3. By definition, we have $a_3 \geq \frac{n-1}{2}$, so $a_1 + a_2 \leq \frac{n+1}{2}$ and $M_3 < \frac{2}{3}(a_1 + a_2) = \frac{n+1}{3}$. Therefore, all machine loads equal $n/3$ so one bag per machine fits, which is a contradiction. \square

Lemma 4.5.c. *For $m = 6$, the optimal algorithm for speed-robust scheduling for unit-size jobs has a robustness factor larger than $\bar{\rho}(m)$.*

Proof. Consider $n = 756$ unit-size jobs and $m = 6$ machines. Consider any algorithm building 6 bags of sizes $a_1 \leq a_2 \leq a_3 \leq a_4 \leq a_5 \leq a_6$ out of these jobs. Consider an adversary setting where five machines are set to a speed $s_1 \in \mathbb{N}$ and one machine to a speed $s_6 \in \mathbb{N}$ with $5s_1 + s_6 \geq n$ such that an optimal schedule of n jobs on these machines has a makespan at most 1. For the algorithm and for each i , either a bag of size at least a_i is scheduled on a machine of speed s_1 or all bags of size at least a_i are scheduled on the machine of speed s_6 . Hence, for each $i \in \{1, \dots, 6\}$, the number $\phi_i := \min(a_i/s_1, \sum_{j=i}^6 a_j/s_6)$ is a lower bound on the algorithm's makespan, in other words, on its robustness factor.

Consider Table 2. Observe that the Conditions (2) and (3) are equivalent as n , the number of jobs, is fixed. If none of the Conditions (1) are satisfied, then the total size of the bags is at most $76 + 91 + 109 + 132 + 158 + 189 = 755 < n$, which is a contradiction. Hence, let i be the first row such that Condition (1) is satisfied. The fact that Condition (1) is not satisfied for $i' < i$ then implies that Condition (3) is met by definition. Hence, the bag sizes of the algorithm satisfy the conditions of at least one row. Observe that this implies that for this particular i , ϕ_i is then a lower bound on the makespan of the algorithm's assignment. Overall, $\phi = \min_i \phi_i$ is then a lower bound on the robustness factor of the algorithm.

We have that $\phi = \frac{589}{391} \approx 1.506 > \bar{\rho}(6) \approx 1.503$. Hence, the robustness factor of *any* algorithm for unit-size jobs and m machines is indeed strictly larger than the robustness factor $\bar{\rho}(m)$ for infinitesimal jobs on the same number of machines; see Theorem 2.3. \square

Conditions on bag sizes			s_1	s_6	ϕ
(1)	(2)	(3)			
$a_1 \geq 77$			51	501	$\min(\frac{77}{51}, \frac{756}{501})$
$a_2 \geq 92$	$\sum_{j \geq 2} a_j \geq 680$	$a_1 \leq 76$	61	451	$\min(\frac{92}{61}, \frac{680}{451})$
$a_3 \geq 110$	$\sum_{j \geq 3} a_j \geq 589$	$a_1 + a_2 \leq 167$	73	391	$\min(\frac{110}{73}, \frac{589}{391})$
$a_4 \geq 133$	$\sum_{j \geq 4} a_j \geq 480$	$\sum_{j < 4} a_j \leq 276$	88	318	$\min(\frac{133}{88}, \frac{480}{318})$
$a_5 \geq 159$	$a_5 + a_6 \geq 348$	$\sum_{j < 5} a_j \leq 408$	105	231	$\min(\frac{159}{105}, \frac{348}{231})$
$a_6 \geq 190$			126	226	$\frac{190}{126}$

Table 2: Speed instance in function of the bag sizes in Lemma 4.5.c.

C Proofs for Section 4.2 – Equal-Size Jobs and Speeds in $\{0, 1\}$

In this section, we give the missing proof of Lemma 4.11. For the remainder of this section, we use $\bar{\lambda} := \text{OPT}_m$ to denote the optimal makespan on m machines.

Lemma 4.11. *If $\text{OPT}_m \in \{3, \dots, 8\}$ and $m \geq 50$, there is a $\frac{4}{3}$ -robust algorithm.*

The proof of Lemma 4.11 consists of two major cases depending on the number of bags that LPT assigns to the same machine. The first part of the proof is to consider m' such that LPT assigns at most 2 bags to any machine (Lemma C.1), and the second part consists of m' such that there is at least one machine to which LPT assigns at least 3 bags (Lemma C.2).

Recall that for $3 \leq \bar{\lambda} \leq 8$, we pack four different types a_0, \dots, a_3 of bags depending on $\bar{\lambda}$, where a_l denotes the size of the l -th bag type and x_l its multiplicity. We give their values in the following two tables.

$\bar{\lambda}$	3	4	5	6	7	8	$m \pmod{5}$	0	1	2	3	4	a_l
a_0	1	2	3	3	4	5	$x_0 + x_1$	$\frac{2m}{5}$	$2 \lfloor \frac{m}{5} \rfloor$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$\lceil \frac{2}{3} \bar{\lambda} \rceil$
a_1	2	3	4	4	5	6	x_2	$\frac{m}{5}$	$\lfloor \frac{m}{5} \rfloor + 1$	$\lfloor \frac{m}{5} \rfloor$	$\lfloor \frac{m}{5} \rfloor + 1$	$\lfloor \frac{m}{5} \rfloor + 2$	$\bar{\lambda}$
a_2	3	4	5	6	7	8	x_3	$\frac{2m}{5}$	$2 \lfloor \frac{m}{5} \rfloor$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$2 \lfloor \frac{m}{5} \rfloor + 1$	$\lfloor \frac{4}{3} \bar{\lambda} \rfloor$
a_3	4	5	6	8	9	10							

(B)

Lemma C.1. *Let $LPT_{m'}$ denote the makespan of LPT that assigns bags as described by (B) to m' machines. If $m' \geq \frac{m}{2}$, then $LPT_{m'} \leq \frac{4}{3} \text{OPT}_{m'}$.*

Proof. To prove this lemma, we consider all cases of how LPT assigns two bags to the same machine. Let $LPT_{m'}$ denote the resulting makespan. We start by bounding the number of failing machines t depending on the value of $LPT_{m'}$. If $LPT_{m'} = a_{l'} + a_l$ with $l' \leq l$, then

$$t \geq \sum_{l''=0}^{l'-1} x_{l''} + \left\lceil \sum_{l''=l'}^{l-1} x_{l''} / 2 \right\rceil + 1. \quad (11)$$

$LPT_{m'} \in \{2a_0, a_0 + a_1, 2a_1\}$: Since $\bar{\lambda} < \text{OPT}_{m-1}$ by our assumption based on Lemma 4.7, we have $2a_1 \leq \frac{4}{3} \text{OPT}_{m'}$ if $m' < m$. So if $LPT_{m'} \in \{2a_0, a_0 + a_1, 2a_1\}$, then $LPT_{m'} \leq \frac{4}{3} \text{OPT}_{m'}$.

$LPT_{m'} = a_0 + a_2$: Observe that $\lfloor (x_0 + x_1)/2 \rfloor \geq \lfloor \frac{m}{5} \rfloor$ by definition, see (B). With Equation (11), we obtain $m' = m - t \leq m - (\lfloor m/5 \rfloor + 1) \leq \frac{4}{5}m$. This implies that

$$\text{OPT}_{m'} \geq \frac{n}{m'} \geq \frac{m\bar{\lambda} - \ell}{4/5m} = \frac{5}{4} \left(\bar{\lambda} - \frac{\ell}{m} \right).$$

As $\ell < \bar{\lambda} \leq 8$ and $m \geq 37$ by assumption, we have that $\frac{5}{4} \frac{\ell}{m} \leq \frac{1}{4}$. Hence,

$$\frac{4}{3} \text{OPT}_{m'} \geq \frac{4}{3} \left(\frac{5}{4} \bar{\lambda} - \frac{1}{4} \right) \geq \frac{5}{3} \bar{\lambda} - \frac{1}{3}.$$

Using that $a_0 + a_2 = (\lfloor 2/3 \bar{\lambda} \rfloor - 1) + \bar{\lambda} \leq \frac{5}{3} \bar{\lambda} - 1$, we conclude that $LPT_{m'} \leq \frac{4}{3} \text{OPT}_{m'}$.

For the remaining cases, let l and l' be the indices of the bag types that are assigned to the same machine and let $y_{ll'} := \sum_{l''=0}^{l'-1} x_{l''} + \left\lceil \sum_{l''=l'}^{l-1} x_{l''} / 2 \right\rceil$. Then, $t \geq y_{ll'} + 1$ by Equation (11). Showing that $LPT_{m'} \leq \frac{4}{3} \text{OPT}_{m'}$ is equivalent to showing that

$$a_{l'} + a_l \leq \frac{4}{3} \left\lceil \frac{n}{m - (y_{ll'} + 1)} \right\rceil$$

for all possible combinations of $a_{l'}$ and a_l . With $a_{ll'} := \lceil \frac{3}{4}(a_{l'} + a_l) \rceil$, this inequality holds if

$$y_{ll'} \geq \left\lfloor \frac{(a_{ll'} - 1)m - n}{a_{ll'} - 1} \right\rfloor = m - \left\lfloor \frac{n}{a_{ll'} - 1} \right\rfloor = \left\lfloor m - \frac{\bar{\lambda}m}{a_{ll'} - 1} + \frac{\ell}{a_{ll'} - 1} \right\rfloor, \quad (12)$$

where we used the fact that $y_{ll'} \in \mathbb{Z}$.

LPT $_{m'}$ = $\mathbf{a}_1 + \mathbf{a}_2$: By Equation (12), it suffices to verify $x_0 + \lfloor \frac{x_1}{2} \rfloor \geq m - \left\lfloor \frac{n}{\lceil \frac{3}{4}(a_1 + a_2) \rceil - 1} \right\rfloor$. If $\bar{\lambda} = 3$, then the right hand side becomes $m - \lceil \frac{3m - \ell}{3} \rceil = 0$. Hence, the inequality is satisfied. For $\bar{\lambda} \geq 4$, the right hand side is at most $\lfloor \frac{m}{5} + \frac{\ell}{5} \rfloor \leq \lfloor \frac{m}{5} \rfloor + \lfloor \frac{4}{5} + \frac{\ell}{a-1} \rfloor$ with $0 \leq \ell < \bar{\lambda}$. Observe that $\lfloor \frac{x_1}{2} \rfloor + x_0 \geq \lfloor \frac{m}{5} \rfloor + \lfloor \frac{x_0}{2} \rfloor$. Using the definition of $x_0 + x_1$ and $x_0 = \ell$,

$$\left\lfloor \frac{m}{5} \right\rfloor + \left\lfloor \frac{4}{5} + \frac{\ell}{a-1} \right\rfloor \leq \begin{cases} \left\lfloor \frac{m}{5} \right\rfloor + 0 \leq \lfloor \frac{x_1}{2} \rfloor + x_0 & \text{if } x_0 = 0 \\ \left\lfloor \frac{m}{5} \right\rfloor + 0 \leq \lfloor \frac{x_1}{2} \rfloor + x_0 & \text{if } x_0 = 1 \text{ and } \bar{\lambda} \geq 5 \\ \left\lfloor \frac{m}{5} \right\rfloor + 1 \leq \lfloor \frac{x_1}{2} \rfloor + x_0 & \text{if } x_0 \geq 2, \end{cases}$$

which shows the validity of Inequality (12) in all cases except the combination of $\ell = x_0 = 1$ and $\bar{\lambda} = 4$. For this particular case, a careful case distinction based on $m \pmod{5}$ shows that Inequality (12) still holds.

LPT $_{m'}$ = $\mathbf{a}_2 + \mathbf{a}_2$: By Equation (12), it suffices to verify $x_0 + x_1 \geq m - \left\lfloor \frac{n}{\lceil \frac{3a_2}{2} \rceil - 1} \right\rfloor$. The right hand side can be transformed into $\lfloor \frac{m}{3} + \frac{2}{3} \frac{\ell}{m} \rfloor \leq \frac{m}{3} + \frac{2}{3} = \frac{5}{15}m + \frac{10}{15}$. We have $x_0 + x_1 \geq 2 \lfloor \frac{m}{5} \rfloor \geq \frac{6}{15}m - \frac{24}{15}$. Using that $m \geq 50$, we obtain $x_0 + x_1 \geq \frac{5}{15}m + \frac{21}{15}$, which concludes the proof of Inequality (12) for this case.

LPT $_{m'}$ = $\mathbf{a}_0 + \mathbf{a}_3$: We need to verify

$$\frac{x_0 + x_1 + x_2}{2} \geq \left\lfloor m - \frac{\bar{\lambda}}{\lceil \frac{3\bar{\lambda}}{2} \rceil - 1} m + \frac{\ell}{\lceil \frac{3\bar{\lambda}}{2} \rceil - 1} \right\rfloor.$$

Based on $m \pmod{7}$, the last term can be bounded from above by

$$\left\lfloor \frac{2}{7}m + \frac{2}{3} \right\rfloor = \frac{2}{7}m - 2 \frac{m \pmod{7}}{7} + \left\lfloor 2 \frac{m \pmod{7}}{7} + \frac{2}{3} \right\rfloor \leq \frac{2}{7}m + \frac{40}{70}.$$

Consider $\frac{x_0 + x_1 + x_2}{2}$, the left hand side of the inequality. We can express this as

$$\frac{x_0 + x_1 + x_2}{2} = \frac{3}{10}m + R \geq \frac{2}{7}m + \frac{43}{70},$$

where $R \in \{0, \frac{2}{10}, -\frac{1}{10}, \frac{1}{10}, \frac{3}{10}\}$ and we used that $m \geq 50$. This shows the validity of Inequality (12) for the current case.

LPT $_{m'}$ = $\mathbf{a}_1 + \mathbf{a}_3$: Since $a_1 + a_3 = 2\bar{\lambda}$, we need to verify

$$x_0 + \frac{x_1 + x_2}{2} \geq \left\lfloor m - \frac{\bar{\lambda}}{\lceil \frac{3\bar{\lambda}}{2} \rceil - 1} + \frac{\ell}{\lceil \frac{3\bar{\lambda}}{2} \rceil - 1} \right\rfloor.$$

Using $\bar{\lambda} \leq 8$ and that the second term on the right hand side is increasing in $\bar{\lambda}$ and depends on the parity of $\bar{\lambda}$, we can upper bound the right hand side by $\left\lfloor \frac{3}{10}m + \frac{2}{3} \frac{\ell}{\bar{\lambda}-1} \right\rfloor$. Note that the left hand side

is slightly larger than $\frac{3}{10}m$ but not sufficiently large for a crude upper bound. Hence, we rewrite the left hand side as

$$x_0 + \frac{x_1 + x_2}{2} = \frac{3}{10}m + \frac{x_0}{2} + R,$$

with $R \in \{0, \frac{2}{10}, -\frac{1}{10}, \frac{1}{10}, \frac{3}{10}\}$ depending on $m \pmod{5}$.

If $x_0 = \ell = 0$, then the right hand side of the inequality is $\lfloor \frac{3}{10}m \rfloor$. Except for the case $m \pmod{5} = 2$, the term $\frac{x_1+x_2}{2}$ clearly satisfies the inequality. If $m \pmod{5} = 2$, we have $m \pmod{10} \in \{2, 7\}$, which implies that $\lfloor \frac{3}{10}m \rfloor \leq \frac{3}{10}m - \frac{6}{10} \leq \frac{x_1+x_2}{2}$ by the case distinction above.

If $x_0 = \ell > 0$, then $\frac{x_0}{2} \geq \frac{1}{2}$ and, thus, the case distinction yields $x_0 + \frac{x_1+x_2}{2} \geq \frac{3}{10}m + \frac{4}{10}$. Using that the right hand side is upper bounded by $\lfloor \frac{3}{10}m + \frac{2}{3} \rfloor$, we use a similar case distinction based on $m \pmod{10}$ to derive

$$\left\lfloor \frac{3}{10}m + \frac{2}{3} \right\rfloor = \frac{3}{10}m - \frac{3}{10}(m \pmod{10}) + \left\lfloor \frac{3}{10}(m \pmod{10}) + \frac{2}{3} \right\rfloor \leq \frac{3}{10}m + \frac{3}{10},$$

which concludes the proof of Inequality (12).

By our choice of x_0, \dots, x_3 , the occurrence of $\text{LPT}_{m'} \in \{a_2 + a_3, a_3 + a_3\}$ implies that $m' < \frac{m}{2}$ which is not considered in this lemma. \square

Lemma C.2. *Let $\text{LPT}_{m'}$ denote the makespan of LPT that assigns bags as described by (B) to m' machines. If $m' < \frac{m}{2}$, then $\text{LPT}_{m'} \leq \frac{4}{3}\text{OPT}_{m'}$.*

Proof. Assume for the sake of contradiction that LPT fails to place all bags onto the machines such that $\text{LPT}_{m'} \leq \frac{4}{3}\text{OPT}_{m'}$. We consider the first bag $b \in [m]$ whose assignment to the currently least loaded machine causes the failure, i.e., the completion time of this machine exceeds $\frac{4}{3}\text{OPT}_{m'}$. For simplicity, let this be machine i and let C_i be the completion time of i before adding bag b . Let a be the size of bag b .

As $m' < \frac{m}{2}$, $m \geq 37$, and $\bar{\lambda} < \text{OPT}_{m-1}$, we have that $\text{OPT}_{m'} \geq 2\bar{\lambda} + 1$. Hence, if bag b is the first or second bag on machine i , then $C_i + a \leq 2 \lfloor \frac{4}{3}\bar{\lambda} \rfloor \leq \frac{4}{3}\text{OPT}_{m'}$; a contradiction.

Consider the case where bag b is the fourth bag (or larger) on machine i and restrict the instance to consist only of the jobs assigned by LPT so far plus the jobs in bag b . Let $\text{OPT}'_{m'}$ be the optimum of this restricted instance on m' machines. As LPT has not assigned all bags yet, i.e., there is still unscheduled volume, we have that $\text{OPT}'_{m'} \geq C_i + 1$. Since there are already at least three bags of size at least a on machine i , we have $\text{OPT}'_{m'} \geq C_i \geq 3a$. Moreover, as b is the first bag to violate $C_i + a \leq \frac{4}{3}\text{OPT}_{m'}$, bag b determines the makespan of LPT on the restricted instance; a contradiction by Lemma 4.10.

Hence, bag b is the third bag on machine i . Based on its size, we distinguish four cases.

$a = a_0$: If a bag of size a_0 is the first bag whose completion time violates $\frac{4}{3}\text{OPT}_{m'}$, then a bag of size a_0 determines $\text{LPT}_{m'}$. Note that $3a_0 = 3(\lceil \frac{2}{3}\bar{\lambda} \rceil - 1) \leq 3\frac{2}{3}\bar{\lambda} = 2\bar{\lambda} \leq \text{OPT}_{m'}$. Hence, by Lemma 4.10, we have $\text{LPT}_{m'} \leq \frac{4}{3}\text{OPT}_{m'}$; a contradiction.

$a = a_1$: If $C_i = 2a_3$, then $\frac{4}{3}\text{OPT}_{m'} \geq \frac{10}{3}\bar{\lambda}$ while $C_i + a_1 \leq \frac{10}{3}\bar{\lambda}$; a contradiction. If $C_i = a_2 + a_3$, then $\frac{4}{3}\text{OPT}_{m'} \geq \frac{28}{9}\bar{\lambda}$, while $C_i + a_1 = 3\bar{\lambda}$; a contradiction. If $C_i = 2\bar{\lambda} = 2a_2 = a_1 + a_3$, then $\frac{4}{3}\text{OPT}_{m'} \geq \frac{8}{3}\bar{\lambda} + \frac{4}{3}$ while $C_i + a_1 \leq \frac{8}{3}\bar{\lambda} + \frac{2}{3}$; a contradiction. If $C_i = a_1 + a_2$, then $\frac{4}{3}\text{OPT}_{m'} \geq a_2 + a_1 + (a_2 + a_1)/3 + \frac{4}{3}$ while $C_i + a_1 = a_2 + 2a_1$. For $3 \leq \bar{\lambda} \leq 8$, one can check that $(a_2 + a_1)/3 + \frac{4}{3} \geq a_1$; a contradiction. As $m' < \frac{m}{2}$, we have covered all possibilities for C_i . Hence, a bag of size a_1 cannot cause LPT to fail.

$\mathbf{a} = \mathbf{a}_2$: Let x be the number of bags of size a_2 that were successfully assigned by LPT before bag b . Denote by V a volume of $m' \cdot \frac{4}{3}\text{OPT}_{m'}$ minus the volume of already assigned bags which are x_3 bags of size a_3 and x bags of size a_2 . We have

$$V \geq \frac{4}{3}(a_0x_0 + a_1x_1 + (x_2 - x)a_2) + \frac{1}{3}(xa_2 + a_3x_3).$$

Using $x_0 + x_1 = x_3$, this implies

$$V \geq a_1x_3 - \frac{4}{3}x_0 + \frac{1}{3}x_3(a_1 + a_3) + \frac{1}{3}x_2a_2 + (x_2 - x)a_2 \geq \frac{4}{3}a_2x_3,$$

where the second inequality follows from $\frac{4}{3}x_0 < \frac{4}{3}\bar{\lambda} \leq \frac{1}{3}x_2\bar{\lambda} + (x_2 - x)\bar{\lambda}$. As b is the third bag on machine i by the above discussion, we have $m' < \frac{x_3 + x_2}{2} \leq x_3$. Thus, we conclude that $V \geq \frac{4}{3}a_2$. Hence, the total volume left on the m' machines is at least $m'a_2$. Hence, there has to be one machine i' where b still fits, i.e., $C_{i'} + a_2 \leq \frac{4}{3}\text{OPT}_{m'}$. As i is the least loaded machine when LPT assigns bag b and b violates $\frac{4}{3}\text{OPT}_{m'}$, we obtain a contradiction.

$\mathbf{a} = \mathbf{a}_3$: Let x be again the number of bags of size a_3 successfully assigned to machines by LPT before bag b . If a_3 is the size of a third bag on machine i , then $m' < \frac{x_3}{2}$. By definition of x_2 , this additionally implies that $m' \leq x_2$. Let V be the remaining volume after having assigned x bags of size a_3 . Then,

$$V = \frac{4}{3}(a_0x_0 + a_1x_1 + a_2x_2 + a_3(x_3 - x)) + \frac{1}{3}a_3x \geq \frac{8}{3}a_1m' + \frac{4}{3}a_2m' + \frac{2}{3}a_3m',$$

where we used $\frac{4}{3}x_0 \leq a_3 \leq (x_3 - x)a_3$ and $m' < \frac{x_3}{2}$ as well as $m' \leq x_2$. Thus, the remaining volume satisfies $V \geq a_3m'$. Hence, there is at least one machine with remaining volume at least a_3 contradicting LPT's choice of machine i .

As all possible cases for the size of the first bag that causes LPT to fail lead to a contradiction, this proves the statement. \square

Lemma 4.11. Combining the results of Lemmas C.1 and C.2 shows that packing the bags according to (\mathcal{B}) is $\frac{4}{3}$ -robust. \square

D Tradeoff Version of Theorem 3.3 for Packing More Bags

For the sake of completeness, we show an easy way how Theorem 3.3 can be extended to the model in which we may partition the set of jobs into $M \geq m$ bags.

Theorem D.1. For $M \geq m \geq 1$, let $\alpha := 1 + \frac{m-1}{M}$. LPT is α -robust for speed-robust scheduling.

Proof. The proof follows the same lines as for Theorem 3.3 with the adjustment that LPT packs M bags instead of m and, for assigning bags to machines, we use a parameterized capacity bound $\alpha \cdot s_i$ for machine i for some $\alpha \geq 1$.

We slightly adjust the last part of the proof, where we give a lower bound on the total remaining capacity on the m machines when the second-stage algorithm fails to place the $(k+1)$ -st bag. The $(M-k)$ bags that were not placed have a combined volume of at least $V_\ell = (M-k-1)w + T \geq (M-k+1)\frac{T}{2}$. The bags that were placed have a combined volume of at least $V_p = kT$. The remaining capacity is then at least $C = \alpha \cdot V_\ell + (\alpha - 1) \cdot V_p$, and we have

$$\begin{aligned} C &\geq \alpha \cdot (M-k+1)\frac{T}{2} + (\alpha-1) \cdot kT = \frac{\alpha}{2} \cdot TM + T \left(\frac{\alpha k}{2} - \frac{\alpha k}{2} \right) + \frac{\alpha}{2} \cdot T - kT \\ &= \frac{\alpha}{2} \cdot TM + T \left(\frac{\alpha k}{2} + \frac{\alpha}{2} - k \right) = \frac{\alpha}{2} \cdot TM + T \left(\frac{\alpha}{2} - k \cdot \left(1 - \frac{\alpha}{2} \right) \right). \end{aligned}$$

Using $\alpha < 2$ and $k \leq M - 1$, we obtain

$$C \geq \frac{\alpha}{2} \cdot TM + T \left(\frac{\alpha}{2} - (M - 1) \cdot \left(1 - \frac{\alpha}{2} \right) \right).$$

We define $x := M - m$ and replace M by $m + x$. Then

$$\begin{aligned} C &\geq Tm - \left(1 - \frac{\alpha}{2} \right) \cdot Tm + \frac{\alpha}{2} \cdot Tx + T \left(\frac{\alpha}{2} - (m + x - 1) \cdot \left(1 - \frac{\alpha}{2} \right) \right) \\ &\geq Tm + T \cdot \left(- \left(1 - \frac{\alpha}{2} \right) \cdot m + \frac{\alpha}{2} \cdot x + \frac{\alpha}{2} - (m + x - 1) \cdot \left(1 - \frac{\alpha}{2} \right) \right) \\ &= Tm + T \cdot (\alpha \cdot (m + x) + 1 - 2m - x) = Tm. \end{aligned}$$

Notice that the last equality follows directly from $\alpha = 1 + \frac{m-1}{M}$.

Thus, there is a machine with remaining capacity T , which contradicts the assumption that the bag of size T does not fit. \square