

When to Give Up on a Parallel Implementation

Nathan S. Sheffield, Alek Westover

Abstract

In the Serial Parallel Decision Problem (SPDP), introduced by Kuszmaul and Westover [SPAA'24], an algorithm receives a series of tasks online, and must choose for each between a serial implementation and a parallelizable (but less efficient) implementation. Kuszmaul and Westover describe three decision models: (1) *Instantly-committing* schedulers must decide on arrival, irrevocably, which implementation of the task to run. (2) *Eventually-committing* schedulers can delay their decision beyond a task's arrival time, but cannot revoke their decision once made. (3) *Never-committing* schedulers are always free to abandon their progress on the task and start over using a different implementation. Kuszmaul and Westover gave a simple instantly-committing scheduler whose total completion time is 3-competitive with the offline optimal schedule, and proved two lower bounds: no eventually-committing scheduler can have competitive ratio better than $\phi \approx 1.618$ in general, and no instantly-committing scheduler can have competitive ratio better than 2 in general. They conjectured that the three decision models should admit different competitive ratios, but left upper bounds below 3 in any model as an open problem.

In this paper, we show that the powers of instantly, eventually, and never committing schedulers are distinct, at least in the “massively parallel regime”. The massively parallel regime of the SPDP is the special case where the number of available processors is asymptotically larger than the number of tasks to process, meaning that the *work* associated with running a task in serial is negligible compared to its *runtime*. In this regime, we show (1) The optimal competitive ratio for instantly-committing schedulers is 2, (2) The optimal competitive ratio for eventually-committing schedulers lies in $[1.618, 1.678]$, (3) The optimal competitive ratio for never-committing schedulers lies in $[1.366, 1.500]$. We additionally show that our instantly-committing scheduler is also 2-competitive outside of the massively parallel regime, giving proof-of-concept that results in the massively parallel regime can be translated to hold with fewer processors.

1 Introduction

1.1 Background

Many computational tasks can be performed quickly in parallel over a large number of processors — but such parallel implementations may be less work-efficient than a serial implementation on a single processor, requiring substantially more total computation time across all machines. When several different tasks must be completed in as little total time as possible, this trade-off between work and time can necessitate running different tasks in different modes: small tasks can be done in serial to save work, while large tasks must be parallelized to prevent their serial runtimes from dominating the overall computation.

To formalize this problem, Kuszmaul and Westover introduced the *Serial Parallel Decision Problem (SPDP)* [KW24]. In their model, each task has exactly two possible implementations: a parallel implementation which can be worked on by multiple machines at once (where the rate of progress on the implementation is proportional to the number of processors assigned to it), and a serial implementation which can only be worked on by a single processor at a time. If all tasks are available at time 0, it is easy to efficiently determine the optimal strategy: all jobs with serial completion time smaller than some threshold can be run in serial, and the larger tasks must be run in parallel. The model becomes interesting when previously-unknown tasks are allowed to arrive at arbitrary times, and one wishes to minimize the competitive ratio between the total completion time of an *online* algorithm compared to the offline optimal completion time.

Kuszmaul and Westover define three distinct versions of this model, parameterized by the degree to which the online scheduler is able to reverse its decisions.

- a) An *instantly-committing* scheduler must choose an implementation for each task as soon as the task arrives, and is not allowed to revisit this choice.
- b) An *eventually-committing* scheduler may delay choosing an implementation, but must choose one irrevocably before assigning its work to a processor.
- c) A *never-committing* scheduler can, at any time, discard all as-yet completed work on an implementation and re-start the task with the other implementation.

The distinction between the eventually- and never-committing models is motivated by potential practical concerns: if a task involves mutating an input in memory, it may not be feasible to cancel an implementation once it begins running. Westover and Kuszmaul present an instantly-committing scheduler achieving competitive ratio 3, and show competitive ratio lower bounds of 2 and $\phi \approx 1.618$ in the instantly-committing and eventually-committing models, respectively. They conjecture that the ability to delay or cancel choices should allow for more competitive online algorithms, but leave open the problem of finding better competitive ratio upper bounds than 3.

1.2 This Work

In this work, we consider Kuszmaul and Westover’s SPDP when the number of available processors is much larger than the number of tasks, noting that all of their upper and lower bounds hold in this parameter regime. This is a particularly simple setting, since the work associated to a serial implementation is now negligible compared to its completion time — running a task in serial means accepting a lower bound on completion time, but requires essentially no work. We can think of this setting as an unrelated-machines scheduling problem with an unlimited number of identical “slow” machines, and a single unrelated “fast” machine, representing a massively parallel implementation of the task across many processors — note that this could also describe scenarios with a literal fast machine, such as a single piece of accelerated hardware.

Our main results are tight bounds on the competitive ratio of instantly-committing schedulers in this regime, and separations between the strength of all 3 models. Our results are summarized in [Table 1](#).

Model	Lower Bound	Upper Bound
Instantly-Committing Schedulers	2 [KW24]	2 *
Eventually-Committing Schedulers	1.618 [KW24]	1.678 *
Never-Committing Schedulers	1.366 *	1.5 *

Table 1: Main Results. * = this work.

In each case, the upper bound comes from a simple heuristic in which the algorithm compares its projected completion time to its current estimate of the optimal completion time. More precisely, at each time t , our scheduler computes an offline optimal strategy “ $\text{opt}(t)$ ” on the truncation of the task sequence to tasks that arrive before time t , and makes decisions based on the completion time of $\text{opt}(t)$.

Our main technical contribution is the analysis of the schedulers. Working with $\text{opt}(t)$ is challenging, because the schedules $\text{opt}(t)$ and $\text{opt}(t')$ can be quite different for $t \neq t'$. For instantly-committing schedulers we use an invariant-based approach to bound, at all times, the work taken by our scheduler in terms of the minimum work and completion time among all schedulers. For eventually- and never-committing schedulers this approach is no longer feasible: there is no well defined notion of the “work taken” by our scheduler, because the scheduler may not have committed to a decision yet. Instead, these analyses rely on choosing a couple of critical times to observe the state of our scheduler and $\text{opt}(t)$, and then establishing a dichotomy: either (1) “real fast tasks” (tasks that opt runs on the fast machine) arrive quickly, in which case our scheduler prioritizes real fast tasks on the fast machine and will run most other tasks on slow machines, or (2) real fast tasks arrive slowly, in which case our scheduler never falls too far behind opt , despite making suboptimal use of the fast machine.

In addition to these results, we show that with some effort our instantly-committing scheduler can be adapted to work for any number of processors, fully resolving the question of the optimal competitive

ratio of instantly-committing schedulers in the general SPDP, and giving a proof-of-concept that results for a large number of processors can be adapted to hold when the work associated with serial tasks is also a concern.

1.3 Related Work

There is a long line of work studying the phenomenon of work-inefficient parallel implementations in multi-processor scheduling. Typically, the models of limited parallelism considered involve one of three types of jobs:

- a) *Rigid* jobs, which come with a number p_i specifying a fixed number of processors the job must be run on at each timestep of its execution.
- b) *Moldable* jobs, where the scheduler may choose the (fixed) number of processors the job is run on, and the amount of work scales depending on this choice according to some *speedup curve*.
- c) *Malleable* jobs, which like moldable jobs have an associated speedup curve, but where the job may be assigned to different numbers of processors at different timesteps (as opposed to the scheduler choosing a fixed value at the start of the task’s runtime).

In each of these cases, there is interest in minimizing the total completion time (*makespan*) in both the offline setting — where problems tend to be NP-hard, but may have approximation algorithms [TWY92; MRT99; LT94; Tur+94a; Tur+94b] — and the online setting, where the goal is to minimize competitive ratio [Gra69; DM07; BS83; HP08; GK10; YHZ09; YCZ18]. Kuszmaul and Westover’s Serial Parallel Decision Problem is related to this line of work, but doesn’t quite fit into the usual framework — in their model, instead of dealing with an arbitrary speedup curve, there is a single binary decision between a completely serial and perfectly parallelizable implementation.

As noted, the massively-parallel regime of the SPDP considered in this paper can be naturally viewed as a scheduling problem with an unlimited number of identical “slow” machines, and a single unrelated “fast” machine. Standard scheduling problems in the unrelated machines model have also been well-studied, in terms of both offline approximation algorithms and hardness results [HS76; LST90; SS02; VHW14; Mol; PSM20; Gup+21; IL23; DLR23], and online algorithms [Awe+95; Asp+97; SS02; Car08; AGM11; Gup+17; Gup+20; Zha+22]. We note, however, that since we treat “slow” machines as an unbounded resource, and there is only a single fast machine available, most of the typical difficulties of multi-processor scheduling problems do not arise. In particular, unlike a typical load-balancing problem where NP-hardness follows from a standard set-partition reduction, the One-Fast-Many-Slow Decision Problem (without dependencies) is easily solvable offline, simply by putting all tasks which finish below a certain threshold on distinct slow machines.

1.4 Open Questions

We leave three main open questions as directions for future work.

Question 1.1. What are the optimal competitive ratios for eventually/never-committing schedulers?

In [Appendix A](#) we identify barriers, showing that improving on our eventually/never-committing schedulers will require substantially different algorithms — but we suspect that such improvements may be possible.

Question 1.2. Are randomized schedulers more powerful than deterministic schedulers?

In the main body of the paper we consider only deterministic schedulers; however, for many online problems randomized algorithms can do substantially better than deterministic ones. In [Appendix C](#) we give some lower bounds against randomized schedulers, but these bounds are weaker than those known for deterministic schedulers.

Question 1.3. Is there a general transformation between schedulers for the massively parallel regime (i.e. the One-Fast-Many-Slow Decision Problem) and the general SPDP?

The fundamental difficulty of the SPDP is deciding between implementations which take a lot of work, and implementations which take a lot of time. This tradeoff is absolute in the massively parallel regime, since the large number of processors means the amount of work associated with a serial implementation

is negligible, whereas in the general SPDP it is possible for all processors to be saturated with serial implementations to run. Intuitively, one might expect that having work associated to the serial implementations only makes the problem *easier*, since it makes the tradeoff less dramatic — indeed, Kuzmaul and Westover’s competitive ratio lower bounds become weaker when the number of processors is small. So, one might hope that algorithms in the massively parallel regime can be generically translated to limited-processor settings. Formalizing this connection is an interesting direction for future research.

2 Preliminaries

2.1 The One-Fast-Many-Slow Decision Problem

In this section we formally define the *One-Fast-Many-Slow Decision Problem*, where the goal is to distribute work between a single *fast machine* and an unlimited number of *slow machines*. An instance of the problem is a *Task Arrival Process (TAP)* $\mathcal{T} = (\tau_1, \dots, \tau_n)$, where each task τ_i consists of a tuple (σ_i, π_i, t_i) indicating runtime on a slow machine, runtime on the fast machine, and arrival time, respectively, such that $t_1 \leq \dots \leq t_n$. A valid schedule associates at most one task to each machine at each point in time¹ such that no work is done on any task before its arrival time, each task runs on at most 1 machine, and each task τ_i is either run for a total of σ_i time on some slow machine, or a total of π_i time on the fast machine. The *completion time* (also known as *makespan*) of the schedule is the time when the last task is finished.

We will be interested in online algorithms for this problem. An online scheduler learns about each task only at its arrival time, and at each time t must already have fixed the prefix of the schedule on times less than t . We define three distinct models for how these online decisions are made:

- a) For each task τ , an *instantly-committing* scheduler must fix at τ ’s arrival time the machine that τ will run on.
- b) An *eventually-committing* scheduler need not fix a machine for any task until that task begins running.
- c) A *never-committing* scheduler is an eventually-committing scheduler with the additional power to, at any time, “cancel” a task from the schedule, erasing all work previously done on the task and allowing it to be re-assigned to a new machine.

In each case, we are interested in minimizing the *competitive ratio* of an online scheduler, which is the supremum over all TAPs of the ratio of the online scheduler’s completion time to the completion time of an optimal scheduler on that TAP.

2.2 Connection to the SPDP

In Kuzmaul and Westover’s Serial Parallel Decision Problem, a scheduler must allocate work to p equally-powerful processors, where each task is specified by the work of the serial implementation, the work of the parallel implementation, and the runtime. The scheduler must choose whether to run each task in serial or parallel, and then must assign the resulting work to the p processors, where parallel work can run on multiple processors at once but serial work cannot.

We can define the *massively parallel* regime of this problem to be the limit as the number of processors becomes large compared to the number of tasks. Letting n be the number of tasks, if $n \leq \epsilon p$ then restricting the serial implementations to run on only the first n many processors, and the parallel implementations to run on only the last $p - n$ many processors, the completion time can increase by at most a $\frac{1}{1-\epsilon}$ factor. This corresponds directly to the One-Fast-Many-Slow Decision Problem: we think of each of these serial processors as a “slow machine”, noting that since we have as many as we have tasks there are effectively an unlimited number. We think of the parallel processors collectively as a “fast machine”, noting that we can assume without loss of generality that, at any point in time, all parallel processors are running the same parallel implementation.

¹In order for the notions like “amount of work performed on τ_i ” to be well-defined, we must additionally mandate that a schedule be measurable. Alternatively, one can assume that time is discretized into appropriately fine timesteps.

2.3 Notation

We now introduce our notation for describing and analyzing schedulers. For algorithm alg and TAP \mathcal{T} , we let $C_{\text{alg}}^{\mathcal{T}}$ be the completion time of alg on \mathcal{T} . Let \mathcal{T}^t be the truncation of TAP \mathcal{T} consisting of the tasks τ_i with $t_i \leq t$. When \mathcal{T} is clear from context we will write C_{alg}^t to denote $C_{\text{alg}}^{\mathcal{T}^t}$, and we will write C_{alg} to denote $C_{\text{alg}}^{\infty} = C_{\text{alg}}^{\mathcal{T}}$. We will also use \tilde{C}_{alg} to denote the completion time of the fast machine — that is, the final time when the fast machine has work.

It will be useful to be able to talk about the optimal completion time of a prefix of the TAP. Define the schedule $\text{opt}(t)$ to be a schedule for \mathcal{T}^t with minimal completion time. Note that $\text{opt}(t)$ is only defined as an offline strategy, but that an online algorithm can compute it at time t , thus obtaining a lower bound on C_{opt} , which will be useful to inform the algorithm's future decisions. For ease of notation, we'll often abbreviate $C_{\text{opt}(t)}^t$ as C^t .

There may be many sets of decisions which result in the optimal completion time; as opposed to letting $\text{opt}(t)$ be an arbitrary such scheduler, it will be useful to fix a canonical one, which we will do by letting $\text{opt}(t)$ run as many tasks in serial as possible.

Scheduler 2.1. The scheduler $\text{opt}(t)$, defined on \mathcal{T}^t , makes decisions as follows:

- If τ_i has $\sigma_i + t_i \leq C^t$, run τ_i on a slow machine when it arrives.
- Otherwise, run τ_i on the fast machine. Prioritize tasks with larger $\sigma_i + t_i$, and break ties by taking tasks with smaller i .

Finally, we let $[n] = \{1, \dots, n\}$, and for a set J of tasks we will write π_J to denote $\sum_{j \in J} \pi_j$.

3 A 2-Competitive Instantly-Committing Scheduler

In this section we present and analyze a 2-competitive instantly-committing scheduler. Kuszmaul and Westover showed that a competitive ratio of $(2 - \varepsilon)$ is impossible for instantly-committing schedulers, so our scheduler is optimal. The scheduler, which we call ins is defined in [Scheduler 3.1](#).

Scheduler 3.1. When task τ_i arrives:

- If $\sigma_i + t_i > 2C^{t_i}$ run τ_i on the fast machine, with the fast machine processing tasks in order of arrival.
- Otherwise run τ_i on a slow machine.

We analyze ins by showing inductively that \tilde{C}_{ins} (the completion time of the fast machine) is small compared to the work and completion time of *any* other schedule. For length n TAP \mathcal{T} , scheduler alg , and $i \in [n]$, we define the quantity K_{alg}^i to be the sum of π_j for all tasks $\tau_j \in \mathcal{T}^t$ that alg runs on the fast machine. The key to analyzing ins is the following lemma.

Lemma 3.2. *Fix a length n TAP. For all $i \in [n]$, and for all instantly-committing schedulers alg ,*

$$\tilde{C}_{\text{ins}}^{t_i} \leq C_{\text{alg}}^{t_i} + K_{\text{alg}}^{t_i}. \quad (1)$$

Proof. We prove the lemma by induction on i . For $i = 1$ the claim is trivial. Now, fix $i \in [n - 1]$, alg and assume the lemma for i and for all alg' ; we will prove the lemma for $i + 1$, alg .

If ins runs τ_{i+1} on a slow machine then $\tilde{C}_{\text{ins}}^{t_{i+1}} = \tilde{C}_{\text{ins}}^{t_i}$, and $C_{\text{alg}}^{t_i} + K_{\text{alg}}^i \leq C_{\text{alg}}^{t_{i+1}} + K_{\text{alg}}^{i+1}$. Thus, the invariant (1) is maintained. We always have $\tilde{C}_{\text{ins}}^{t_{i+1}} \leq \tilde{C}_{\text{ins}}^{t_i} + \pi_{i+1}$, so if alg runs τ_{i+1} on the fast machine then the invariant (1) is also maintained, since then $C_{\text{alg}}^{t_{i+1}} + K_{\text{alg}}^{t_{i+1}} \geq C_{\text{alg}}^{t_i} + K_{\text{alg}}^{t_i} + \pi_{i+1} \geq \tilde{C}_{\text{ins}}^{t_i} + \pi_{i+1}$ by the inductive hypothesis.

The final case to consider is when alg runs τ_{i+1} on a slow machine, while ins runs τ_{i+1} on the fast machine. From the definition [Scheduler 3.1](#) of ins , the fact that ins ran τ_{i+1} on a slow machine implies

$$2C^{t_{i+1}} < \sigma_{i+1} + t_{i+1}. \quad (2)$$

On the other hand, alg ran τ_{i+1} on the fast machine. Thus,

$$\sigma_{i+1} + t_{i+1} \leq C_{\text{alg}}^{t_{i+1}}. \quad (3)$$

Now, we use the invariant for $(i, \text{opt}_{t_{i+1}})$ to bound $\tilde{C}_{\text{ins}}^{t_{i+1}}$. We have:

$$\tilde{C}_{\text{ins}}^{t_{i+1}} \leq \tilde{C}_{\text{ins}}^{t_i} + \pi_{i+1} \leq K_{\text{opt}(t_{i+1})}^{t_i} + C_{\text{opt}(t_{i+1})}^{t_i} + \pi_{i+1}. \quad (4)$$

Because of (2) we know that $\text{opt}(t_{i+1})$ must run τ_{i+1} on the fast machine. So, we have

$$K_{\text{opt}(t_{i+1})}^{t_i} + C_{\text{opt}(t_{i+1})}^{t_i} + \pi_{i+1} = K_{\text{opt}(t_{i+1})}^{t_{i+1}} + C_{\text{opt}(t_{i+1})}^{t_{i+1}} \leq 2C^{t_{i+1}}. \quad (5)$$

Stringing together the above inequalities (4), (5), (2), and (3), we get

$$\tilde{C}_{\text{ins}}^{i+1} < C_{\text{alg}}^{t_{i+1}}.$$

Thus, the invariant (1) holds. \square

Using Lemma 3.2 it is easy to show that **ins** is 2-competitive.

Theorem 3.3. *ins is a 2-competitive instantly-committing scheduler.*

Proof. By Lemma 3.2 we have $\tilde{C}_{\text{ins}}^{t_n} \leq 2C_{\text{opt}}$. Thus, **ins** finishes using the fast machine before time $2C_{\text{opt}}$. Any task that **ins** runs on a slow machine must have $\sigma_i + t_i \leq 2C_{\text{opt}}$, so these tasks finish before $2C_{\text{opt}}$ as well. \square

4 A 1.678-Competitive Eventually-Committing Scheduler

In this section we present and analyze a ξ -competitive eventually-committing scheduler, where $\xi \approx 1.678$ is the real root of the polynomial $2x^3 - 3x^2 - 1$. Kuszmaul and Westover gave a lower bound of $\phi \approx 1.618$ on the competitive ratio of any eventually-committing scheduler and conjectured that this lower bound is tight. Our scheduler represents substantial progress towards resolving Kuszmaul and Westover's conjecture, improving on their previous best algorithm which had a competitive ratio of 3. Our scheduler, which we call **eve**, is defined in Scheduler 4.1.

Scheduler 4.1. At time t :

- If task τ_i , which has arrived but not yet been started, has $\sigma_i + t \leq \xi C^t$, then start τ_i on a slow machine.
- Maintain up to one **active** task at a time. The fast machine is always allocated to the active task.
- When there is no active task, but there are unstarted tasks present, choose as the new active task the unstarted task with the largest $\sigma_i + t_i$ value (breaking ties arbitrarily).

Theorem 4.2. *eve is a ξ -competitive eventually-committing scheduler.*

Proof. Fix TAP \mathcal{T} . Let \tilde{C}_{eve} denote the time when **eve** completes the last task run on the fast machine. If τ_i is run on a slow machine at any time t , then τ_i finishes before $\xi C^t \leq \xi C_{\text{opt}}$. Thus, it suffices to show that $\tilde{C}_{\text{eve}} \leq \xi C_{\text{opt}}$.

For any $x \in [0, C_{\text{opt}}]$, let $R(x)$ be the first time that an online algorithm becomes aware that the optimal schedule requires at least x completion time — that is, $R(x) = \inf \{t : C^t \geq x\}$. Let \mathcal{A} (“actual”) be the set of tasks that **opt** runs on the fast machine, and \mathcal{F} (“fake”) be the set of tasks that **opt** runs on a slow machine but **eve** runs on the fast machine. We can bound the sizes and arrival times of tasks in \mathcal{F} as follows.

Claim 4.2.1. *All tasks $\tau_i \in \mathcal{F}$ arrive before time $R(C_{\text{opt}}/\xi)$, and have $\pi_i < \sigma_i/\xi$.*

Proof. All tasks $\tau_i \in \mathcal{F}$ are run on the fast machine by **eve**, and on slow machines by **opt**. In particular this means

$$\xi C^{t_i} < \sigma_i + t_i \leq C_{\text{opt}} \leq \xi C^{R(C_{\text{opt}}/\xi)}.$$

Thus, $t_i < R(C_{\text{opt}}/\xi)$. To show $\pi_i < \sigma_i/\xi$, note that $\pi_i + t_i \leq C^{t_i} < \frac{\sigma_i + t_i}{\xi}$. \square

To analyze when tasks in \mathcal{F} get run it will be useful to partition \mathcal{F} into $\mathcal{F}_{\text{big}} = \{\tau_i \in \mathcal{F} : \sigma_i + t_i > C_{\text{opt}}/\xi\}$ and $\mathcal{F}_{\text{small}} = \{\tau_i \in \mathcal{F} : \sigma_i + t_i \leq C_{\text{opt}}/\xi\}$. Now we show that, without loss of generality, **eve** does not start any tasks in $\mathcal{F}_{\text{small}}$ too late.

Claim 4.2.2. *If eve starts a task $\tau \in \mathcal{F}_{\text{small}}$ at any time $t \geq R(C_{\text{opt}}/\xi)$, then $\tilde{C}_{\text{eve}} \leq \xi C_{\text{opt}}$.*

Proof. Note that no task $\tau_i \in \mathcal{F}_{\text{small}}$ can be started after time C_{opt} : since $C_{\text{opt}} + \sigma_i \leq C_{\text{opt}} + C_{\text{opt}}/\xi < \xi C_{\text{opt}}$, any task $\tau_i \in \mathcal{F}_{\text{small}}$ present but not already running at time C_{opt} would be run on a slow machine. Let t_* be the last time after $R(C_{\text{opt}}/\xi)$ when *eve* starts a task $\tau \in \mathcal{F}$. If t_* does not exist the claim is vacuously true. In light of our previous observation, $t_* < C_{\text{opt}}$. Let τ_i be the task that *eve* starts at time t_* . Because *eve* prioritizes making tasks with larger $\sigma_j + t_j$ values active, at time t_* there are no tasks $\tau \in \mathcal{F}_{\text{big}} \cup \mathcal{A}$ present. After time t_* , no more tasks from \mathcal{F} can arrive by [Claim 4.2.1](#), and at most $C_{\text{opt}} - t_*$ work in \mathcal{A} can arrive because *opt* must be able to complete this work. Thus,

$$\tilde{C}_{\text{eve}} \leq t_* + (C_{\text{opt}} - t_*) + \pi_i = C_{\text{opt}} + \pi_i. \quad (6)$$

Now, because $\tau_i \in \mathcal{F}_{\text{small}}$ we have $\pi_i \leq C_{\text{opt}}/\xi^2$; using this in (6) we find $\tilde{C}_{\text{eve}} \leq (1 + 1/\xi^2)C_{\text{opt}} \leq \xi C_{\text{opt}}$. \square

This means that we can assume that, after time $R(C_{\text{opt}}/\xi)$, the only tasks that *eve* runs on the fast machine are \mathcal{A} , \mathcal{F}_{big} , and whatever the active task was at time $R(C_{\text{opt}}/\xi)$. We call the active task at time $R(C_{\text{opt}}/\xi)$, if one exists, the **stuck** task, denoted τ_s . We split into cases depending on how large this stuck task is.

Case 1: There is no stuck task.

In this case, we in fact have $\tilde{C}_{\text{eve}} \leq C_{\text{opt}}$. Since there is no active task at time $R(C_{\text{opt}}/\xi)$, there are no tasks present but not started on slow machines. By [Claim 4.2.1](#), *eve* will run all tasks $\tau \notin \mathcal{A}$ arriving after time $R(C_{\text{opt}}/\xi)$ on slow machines. Thus, at all time steps $t \in [R(C_{\text{opt}}/\xi), C_{\text{opt}}]$, *eve* either has no active task on the fast machine, or has some $\tau \in \mathcal{A}$ as the active task on the fast machine, so *eve* completes \mathcal{A} at least as quickly as *opt*.

Case 2: There is a stuck task, with $\sigma_s + t_s > C_{\text{opt}}/\xi$.

Define $\mathcal{A}_{\text{early}} = \{\tau_i \in \mathcal{A} : t_i < R(C_{\text{opt}}/\xi)\}$ and $\mathcal{A}_{\text{late}} = \mathcal{A} \setminus \mathcal{A}_{\text{early}}$. Let $t < R(C_{\text{opt}}/\xi)$ be a time when all tasks in $\{\tau_s\} \cup \mathcal{F}_{\text{big}} \cup \mathcal{A}_{\text{early}}$ have already arrived; such a time must exist by [Claim 4.2.1](#). Observe that *opt*(t) runs all tasks in $\{\tau_s\} \cup \mathcal{F}_{\text{big}} \cup \mathcal{A}_{\text{early}}$ on the fast machine due to $C^t < C_{\text{opt}}/\xi$. This further implies that $\pi_{\mathcal{F}_{\text{big}} \cup \{\tau_s\} \cup \mathcal{A}_{\text{early}}} \leq C_{\text{opt}}/\xi$. Also $\pi_{\mathcal{A}_{\text{late}}} \leq C_{\text{opt}} - R(C_{\text{opt}}/\xi)$, simply because *opt* must complete the work on tasks $\mathcal{A}_{\text{late}}$ after these tasks arrive. By [Claim 4.2.2](#) we may assume without loss of generality that, after time $R(C_{\text{opt}}/\xi)$, *eve* is always running a task from $\{\tau_s\} \cup \mathcal{A} \cup \mathcal{F}_{\text{big}}$. Thus,

$$\begin{aligned} \tilde{C}_{\text{eve}} &\leq R(C_{\text{opt}}/\xi) + \pi_{\mathcal{A} \cup \mathcal{F}_{\text{big}} \cup \{\tau_s\}} \\ &= R(C_{\text{opt}}/\xi) + \pi_{\mathcal{A}_{\text{early}} \cup \mathcal{F}_{\text{big}} \cup \{\tau_s\}} + \pi_{\mathcal{A}_{\text{late}}} \\ &\leq R(C_{\text{opt}}/\xi) + C_{\text{opt}}/\xi + (C_{\text{opt}} - R(C_{\text{opt}}/\xi)) \\ &\leq \xi C_{\text{opt}} \end{aligned}$$

Case 3: There is a stuck task, with $\sigma_s + t_s \leq C_{\text{opt}}/\xi$.

This case will be the most difficult to handle of the three. It will be useful to focus now on the tasks of \mathcal{F}_{big} that arrive after the stuck task is started. Let t_* be the time when *eve* starts running τ_s , and let $\mathcal{F}_{\text{big}}' = \{\tau_i \in \mathcal{F}_{\text{big}} : t_i \geq t_*\}$ be the fake tasks arriving after t_* . We first observe that, if no such tasks arrive, *eve* performs very well.

Claim 4.2.3. *If $\mathcal{F}_{\text{big}}' = \emptyset$ then $\tilde{C}_{\text{eve}} \leq \xi C_{\text{opt}}$.*

Proof. At time t_* no tasks $\tau_i \in \mathcal{A} \cup \mathcal{F}_{\text{big}}$ can be present, since all such tasks have $\sigma_i + t_i > C_{\text{opt}}/\xi$, so *eve* would prioritize running them on the fast machine instead of the stuck task. By [Claim 4.2.2](#), we know that after time t_* *eve* will always be running tasks from $\{\tau_s\} \cup \mathcal{A} \cup \mathcal{F}_{\text{big}}'$. The total work on tasks from \mathcal{A} that arrives after time t_* is at most $C_{\text{opt}} - t_*$, so if $\mathcal{F}_{\text{big}}' = \emptyset$ we have

$$\tilde{C}_{\text{eve}} \leq t_* + \pi_s + C_{\text{opt}} - t_* \leq \sigma_s/\xi + C_{\text{opt}} \leq \xi C_{\text{opt}}.$$

\square

By [Claim 4.2.3](#) we may assume $\mathcal{F}_{\text{big}}' \neq \emptyset$. So, let $\sigma_{\min} = \min_{\tau_i \in \mathcal{F}_{\text{big}}'}(\sigma_i)$; we will be able to control how much work arrives in the TAP by the fact that *eve* never decides to run the task $\tau_i \in \mathcal{F}_{\text{big}}'$

with $\sigma_i = \sigma_{\min}$ on a slow machine. Split \mathcal{A} into $\mathcal{A}_{\text{early}} = \{\tau \in \mathcal{A} : t_* \leq t_i < R(\sigma_{\min})\}$ and $\mathcal{A}_{\text{late}} = \{\tau \in \mathcal{A} : t_i \geq \max(t_*, R(\sigma_{\min}))\}$ (note that we use a different threshold to define earliness here than we did in case 2). First we need the following analogue of [Claim 4.2.2](#).

Claim 4.2.4. *If eve starts a task $\tau \in \mathcal{F}_{\text{big}}'$ at any time $t \in [R(\mathbf{C}_{\text{opt}}/\xi), \mathbf{C}_{\text{opt}}]$, then $\tilde{\mathbf{C}} \leq \xi \mathbf{C}_{\text{opt}}$.*

Proof. Let t_* be the last time in $[R(\mathbf{C}_{\text{opt}}/\xi), \mathbf{C}_{\text{opt}}]$ when eve starts a task $\tau \in \mathcal{F}'$. If t_* does not exist the claim is vacuously true. Let τ_i be the task that eve starts at time t_* . Because eve prioritizes making tasks with larger $\sigma_j + t_j$ values active, at time t_* there are no tasks $\tau \in \mathcal{A}$ present. After time t_* at most $\mathbf{C}_{\text{opt}} - t_*$ work in \mathcal{A} can arrive because opt must be able to complete this work. Recalling that $\pi_{\mathcal{F}_{\text{big}}'} \leq \mathbf{C}_{\text{opt}}/\xi$, we have

$$\tilde{\mathbf{C}}_{\text{eve}} \leq t_* + \pi_{\mathcal{F}_{\text{big}}'} + (\mathbf{C}_{\text{opt}} - t_*) \leq \xi \mathbf{C}_{\text{opt}}.$$

□

Claim 4.2.5. $t_* + \pi_{\mathcal{F}_{\text{big}}'} + \pi_{\mathcal{A}_{\text{early}}} < (\sigma_{\min} + R(\sigma_{\min}))/\xi$.

Proof. Fix a time $t < R(\sigma_{\min})$ after all tasks in $\mathcal{F}_{\text{big}}' \cup \mathcal{A}_{\text{early}}$ have arrived, and fix a task $\tau_i \in \mathcal{F}_{\text{big}}'$ with $\sigma_i = \sigma_{\min}$. First, note that by [Claim 4.2.4](#) we can assume that eve has not started τ_i by time t . Thus, eve is free to start τ_i on a slow machine, but chooses not to. This implies

$$\xi \mathbf{C}^t < \sigma_i + t < \sigma_{\min} + R(\sigma_{\min}). \quad (7)$$

We also observe that opt(t) must run $\mathcal{F}_{\text{big}}' \cup \mathcal{A}_{\text{early}}$ on the fast machine, since running any of them on the slow machine would finish after time σ_{\min} . Thus,

$$t_* + \pi_{\mathcal{F}_{\text{big}}'} + \pi_{\mathcal{A}_{\text{early}}} \leq \mathbf{C}^t. \quad (8)$$

Combining (8) and (7) gives the desired statement. □

The other observation we make is that $R(\sigma_{\min})$ cannot happen too early.

Claim 4.2.6. $R(\sigma_{\min}) > (\xi - 1)\sigma_{\min}$.

Proof. Let $\tau_i \in \mathcal{F}_{\text{big}}'$ be a task with $\sigma_i = \sigma_{\min}$. Note that $t_i \leq R(\sigma_{\min})$ by [Claim 4.2.1](#). Then, by [Claim 4.2.4](#) we may assume without loss of generality that at time $R(\sigma_{\min})$ eve is not running τ_i , and does not choose to start τ_i on a slow machine. So,

$$R(\sigma_{\min}) + \sigma_{\min} > \xi \mathbf{C}^{R(\sigma_{\min})} \geq \xi \sigma_{\min}.$$

□

Now we conclude the theorem.

Claim 4.2.7. $\tilde{\mathbf{C}}_{\text{eve}} \leq \xi \mathbf{C}_{\text{opt}}$.

Proof. Note that $\pi_{\mathcal{A}_{\text{late}}} \leq \mathbf{C}_{\text{opt}} - R(\sigma_{\min})$. Also, note that since τ_s was not put on a slow machine immediately upon arrival, we must have $\pi_s \leq \sigma_s/\xi \leq \mathbf{C}_{\text{opt}}/\xi^2$. Then, applying [Claim 4.2.5](#) and [Claim 4.2.6](#) we have

$$\begin{aligned} \tilde{\mathbf{C}}_{\text{eve}} &\leq \pi_s + t_* + \pi_{\mathcal{F}_{\text{big}}'} + \pi_{\mathcal{A}_{\text{early}}} + \pi_{\mathcal{A}_{\text{late}}} \\ &\leq \mathbf{C}_{\text{opt}}/\xi^2 + (\sigma_{\min} + R(\sigma_{\min}))/\xi + \mathbf{C}_{\text{opt}} - R(\sigma_{\min}) \\ &\leq (1/\xi^2 + (1 + \xi - 1)/\xi + 1 - (\xi - 1))\mathbf{C}_{\text{opt}} \\ &= (3 + 1/\xi^2 - \xi)\mathbf{C}_{\text{opt}} \\ &= \xi \mathbf{C}_{\text{opt}}. \end{aligned}$$

□

□

Remark 4.3. The simple nature of the lower bound in [Proposition B.2](#), along with the fact that eve gets a competitive ratio quite close to ϕ might leave the impression that ϕ is clearly the correct competitive ratio, and a slightly better analysis of (a natural variant of) eve would be ϕ -competitive. However, this is not the case: in [Appendix A](#), we show that no *non-procrastinating* eventually-committing scheduler can achieve competitive ratio better than ξ , where a scheduler is called non-procrastinating if, whenever tasks are present, it always runs at least one task. Thus, if the competitive ratio of eve can be improved upon, doing so will require a substantially different scheduler: one which occasionally decides to do nothing at all despite work being present.

5 A 1.5-Competitive Never-Committing Scheduler

In this section we analyze never-committing schedulers. First we give a simple lower bound.

Proposition 5.1. *Fix $\varepsilon > 0$. There is no deterministic $((1 + \sqrt{3})/2 - \varepsilon)$ -competitive never-committing scheduler.*

Proof. We may assume $\varepsilon < .001$. Let $\psi = (\sqrt{3} - 1)/2$. Let $\mathcal{T} = ((1, 2\psi, 0), (\infty, 1 - \psi, \psi))$; that is, τ_1 has $\sigma_1 = 1, \pi_1 = 2\psi, t_1 = 0$, and τ_2 has $\sigma_2 = \infty, \pi_2 = 1 - \psi, t_2 = \psi$. Let $\mathcal{T}' = ((1, 2\psi, 0))$ be the same TAP without the second task. We have $C_{\text{opt}}^{\mathcal{T}'} = 2\psi$, since **opt** just runs the single task on the fast machine, and $C_{\text{opt}}^{\mathcal{T}} = 1$, since **opt** can run τ_1 on a slow machine and τ_2 on the fast machine.

Suppose that **alg** is a $(\psi + 1 - \varepsilon)$ -competitive scheduler. On TAP \mathcal{T}' , at time $\psi - \varepsilon/2$, we claim **alg** must be running τ_1 on the fast machine. If not, then **alg**'s completion time must be at least $\min(\sigma_1, \psi - \varepsilon/2 + \pi_1) = 1$, with the branch of the min depending on whether τ_1 is ever moved to the fast machine — but this gives competitive ratio $1/(2\psi) = 1 + \psi$.

Before time ψ , it is impossible to distinguish between \mathcal{T} and \mathcal{T}' . Thus, **alg** must be running τ_1 on the fast machine at time $\psi - \varepsilon/2$ on TAP \mathcal{T} . Now, we have $C_{\text{alg}}^{\mathcal{T}} \geq \min(\sigma_1 + \psi - \varepsilon/2, \pi_1 + \pi_2) = 1 + \psi - \varepsilon/2$, with the branch of the min depending on whether τ_1 is ever moved to a slow machine — but this gives competitive ratio $(1 + \psi - \varepsilon/2)/1$. Thus, **alg** is not actually $(\psi + 1 - \varepsilon)$ -competitive. \square

Now we give a 1.5-competitive never-cancelling scheduler, which we call **nev**. Note that this competitive ratio is smaller than the lower bound of $\phi \approx 1.618$ known for eventually committing schedulers, so this demonstrates a separation between the strengths of schedulers in the two models.

Scheduler 5.2. At time t :

- If task τ_i has $\sigma_i + t \leq 1.5C^t$ but is not currently running on a slow machine, start τ_i on a slow machine, cancelling its fast machine implementation if necessary.
- Let \mathcal{P} be the set of τ_i that have arrived and are not running on a slow machine. Choose $\tau_i \in \mathcal{P}$ maximizing $\sigma_i + t_i$, breaking ties by choosing the task with the smaller i . Run τ_i on the fast machine during this time step.

Theorem 5.3. *nev is a 1.5-competitive never-committing scheduler.*

Proof. Fix TAP \mathcal{T} . Let \tilde{C}_{nev} denote the final time when **nev** has work on the fast machine. Observe that if **nev** ever runs τ_i on a slow machine, then **nev** finishes τ_i before time $1.5C_{\text{opt}}$. Thus, to show that **nev** is 1.5-competitive it suffices to show $\tilde{C}_{\text{nev}} \leq 1.5C_{\text{opt}}$.

Let $\mathcal{A} = \{\tau_i : \sigma_i + t_i > C_{\text{opt}}\}$ be the set of tasks that **opt** *actually* runs on the fast machine.

Claim 5.3.1. *nev never runs a task $\tau \in \mathcal{A}$ on the fast machine after time C_{opt} .*

Proof. **nev** always allocates the fast machine to the present task with the largest value of $\sigma_i + t_i$ among tasks that aren't running on slow machines. Thus, whenever there are tasks from \mathcal{A} that aren't running on slow machines, **nev** will run one such task on the fast machine. **opt** is able to complete all tasks in \mathcal{A} on the fast machine by time C_{opt} . Thus, **nev** completes or starts on slow machines all tasks $\tau \in \mathcal{A}$ before time C_{opt} . \square

This means that the only way to have $\tilde{C}_{\text{nev}} > C_{\text{opt}}$ is if there are tasks with $\sigma_i + t_i \leq C_{\text{opt}}$ that have yet to be completed at time C_{opt} ; we assume that this is the case for the remainder of the proof. Let $\Pi(x)$ be the total amount of work **nev** performs on the fast machine after time C_{opt} across all tasks with $\sigma_i + t_i \in [x, 1.5x]$. For any $x \leq C_{\text{opt}}$, let $R(x) = \inf\{t \mid C^t \geq x\}$ be the first time an online algorithm becomes aware that the optimal schedule requires x completion time; the following key claim allows us to bound this left-over work $\Pi(x)$ in terms of $R(x)$.

Claim 5.3.2. *For all x , we have $\Pi(x) \leq x - R(x)$.*

Proof. Let \mathcal{J}_x denote the set of tasks with $\sigma_i + t_i \leq 1.5x$ that **nev** runs on the fast machine at some time after C_{opt} . First, note that all $\tau_i \in \mathcal{J}_x$ must have $t_i < R(x)$, or else τ_i would be placed on a slow machine upon arrival. Choose ε sufficiently small, such that no tasks arrive between times $R(x) - \varepsilon$ and $R(x)$. Since $R(x) - \varepsilon < R(x)$, we know $C^{R(x) - \varepsilon} < x$, and so **opt**($R(x) - \varepsilon$) must run all tasks with $\sigma_i + t_i \geq x$ on the fast machine. In order for **opt**($R(x) - \varepsilon$) to finish these tasks before time $C^{R(x) - \varepsilon} < x$, **opt**($R(x) - \varepsilon$)

must have at most $x - R(x) + \varepsilon$ fast work remaining across all such tasks.

Now, by the same argument as in [Claim 5.3.1](#), because `nev` prioritizes tasks with $\sigma_i + t_i \geq x$ over tasks with $\sigma_i + t_i < x$ on the fast machine whenever they are present (and not yet started on slow machines), `nev` has at most $x - R(x) + \varepsilon$ work remaining on tasks in \mathcal{J}_x at time $R(x) - \varepsilon$. Because no more tasks from \mathcal{J}_x arrive after this time, we have $\Pi(x) \leq x - R(x) + \varepsilon$ as well. The claim held for all $\varepsilon > 0$, and so taking $\varepsilon \rightarrow 0$ we have $\Pi(x) \leq x - R(x)$. \square

We now give an observation to control $R(x)$. Let τ_{i_*} be the task, among all tasks that `nev` runs on the fast machine after time C_{opt} , with the smallest value of $\sigma_i + t_i$. Let $\lambda = \sigma_{i_*} + t_{i_*}$.

Claim 5.3.3. *For all $x \geq \lambda$, we have $R(x) > 1.5x - \lambda$.*

Proof. First, note that $t_{i_*} \leq R(\lambda) \leq \lambda$ or else `nev` would start τ_{i_*} on a slow machine upon arrival. Now, because `nev` doesn't start τ_{i_*} on a slow machine at time $R(x) > t_{i_*}$, we have $R(x) + \sigma_{i_*} > 1.5x$. \square

To prove the theorem, it will now suffice to branch into two cases, based on how large λ is.

Case 1: $\lambda \geq (2/3)C_{\text{opt}}$.

In this case, since $1.5\lambda \geq C_{\text{opt}}$, by [Claim 5.3.1](#) all left-over work at time C_{opt} comes from tasks with $\sigma_i + t_i \in [\lambda, 1.5\lambda]$. By [Claim 5.3.2](#), the total amount of such work is at most $\lambda - R(\lambda)$. Then, by [Claim 5.3.3](#), we know $R(\lambda) > .5\lambda$. Together, this implies that the total amount of leftover work is at most $.5\lambda \leq .5C_{\text{opt}}$.

Case 2: $\lambda < 2C_{\text{opt}}/3$.

First note that $\lambda \geq C_{\text{opt}}/2$ or else τ_{i_*} would be started on a slow machine at time C_{opt} . So, all left-over work at time C_{opt} comes either from tasks with $\sigma_i + t_i \in [\lambda, 1.5\lambda]$, or from tasks with $\sigma_i + t_i \in [1.5\lambda, 1.5^2\lambda]$. By [Claim 5.3.2](#), we can therefore bound the total amount of leftover work by $(\lambda - R(\lambda)) + (1.5\lambda - R(1.5\lambda))$. Now, by [Claim 5.3.3](#), this quantity can be at most $(\lambda - .5\lambda) + (1.5\lambda - 1.25\lambda) = .75\lambda$. Since $\lambda < 2C_{\text{opt}}/3$, this is at most $.5C_{\text{opt}}$. \square

Remark 5.4. In [Appendix A](#) we show that [Scheduler 5.2](#) is optimal among never-committing schedulers that never cancel implementations running on slow machines. This shows that improving on [Scheduler 5.2](#) will require a substantially different scheduler.

6 Extending Beyond the Massively Parallel Regime

In [Theorem 3.3](#), we have shown that [Scheduler 3.1](#) is a 2-competitive instantly-committing scheduler in the Massively Parallel regime of the SPDP. In this section, we will show that in fact, [Scheduler 3.1](#) is a 2-competitive scheduler even outside of the Massively Parallel regime, although the analysis is slightly more complicated. This result is interesting in its own right, resolving an open question from [\[KW24\]](#). However, we think that the main virtue of this proof is that it serves as a proof-of-concept that results from the conceptually simpler Massively Parallel regime can be adapted to apply to the general SPDP: we conjecture that all upper bounds holding in the massively parallel regime should also hold in the general SPDP.

[Scheduler 3.1](#) is not a defined scheduler in the SPDP, because we specify the decisions for which tasks to run, but do not specify how to schedule the tasks. We extend `ins` to the general SPDP as follows:

Scheduler 6.1. When task τ_i arrives:

- If $\sigma_i + t_i > 2C^{t_i}$ parallelize τ_i .
- Otherwise, serialize τ_i .

At every timestep, if there are x serial jobs present, then `ins` schedules the jobs by allocating a processor to each of the $\min(p, x)$ serial jobs with the most remaining work, and then allocating any remaining processors to an arbitrary parallel job (if a parallel job is present).

Now we analyze `ins`. We say `ins` is *saturated* at time t if `ins` has no idle processors at time t .

Lemma 6.2. *If `ins` is unsaturated right before finishing, then $C_{\text{ins}} \leq 2C_{\text{opt}}$.*

Proof. We claim that if `ins` is unsaturated at time t , then for each task τ_i present at time t , τ_i has been run on every time step since it arrived. Suppose this is not the case. Then, there must have been some time step before time t when there were at least p serial jobs with at least as much remaining work as τ_i . But then τ_i will finish at the same time as these other jobs, contradicting the fact that `ins` is unsaturated at time t . Thus, if `ins` is unsaturated at time t , then $t \leq \sigma_i + t_i$ for some i such that `ins` ran τ_i in serial. Thus, $t \leq 2C_{\text{opt}}$, as desired. \square

By virtue of [Lemma 6.2](#) it suffices to consider the case that `ins` is saturated immediately before finishing. Let t_* be the final time in $[0, C_{\text{ins}}]$ when `ins` is unsaturated (we set $t_* = 0$ if `ins` is always saturated). Let $i_* \in [n]$ be the smallest i such that $t_i \geq t_*$; in fact we will have $t_{i_*} = t_*$, since in order to transition from being unsaturated to being saturated, some tasks must arrive. For integer $i \in [i_*, n]$, let K_{alg}^i denote the sum of π_j for each τ_j with $i_* \leq j \leq i$ that `alg` runs in parallel; If `alg` is an instantly-committing scheduler then K_{alg}^i can be computed at time t_i . Now we prove an analogue of [Lemma 3.2](#).

Lemma 6.3. *Fix a length n TAP. For all $i \in [i_*, n]$, and for all instantly-committing schedulers `alg`,*

$$K_{\text{ins}}^i \leq (C_{\text{alg}}^i - t_*)p + K_{\text{alg}}^i. \quad (9)$$

Proof. We prove (9) by induction on i . The base case is $i = i_*$; here the claim trivially holds. Subsequently, note that if `alg` takes at least as much work as `ins` on τ_i (i.e., either `alg` runs τ_i in parallel or `ins` runs τ_i in serial) then $K_{\text{alg}}^i \geq K_{\text{ins}}^i$, and $C_{\text{alg}}^i \geq t_*$, so (9) is true. In the case that `alg` serializes τ_i while `ins` parallelizes τ_i we now have:

$$C_{\text{alg}}^i \geq \sigma_i + t_i \geq 2C^{t_i}.$$

One consequence of this is that `opt`(t_i) parallelizes τ_i , and so we have (9) for $(i, \text{opt}(t_i))$. Thus,

$$K_{\text{ins}}^i \leq (C^{t_i} - t_*)p + K_{\text{opt}(t_i)}^i \leq (2C^{t_i} - 2t_*)p \leq (C_{\text{alg}}^i - t_*)p.$$

\square

In the Massively Parallel regime [Theorem 3.3](#) followed immediately from [Lemma 3.2](#). Slightly more work is required in the general setting, but [Lemma 6.3](#) is still very useful.

Theorem 6.4. *`ins` is a 2-competitive instantly-committing scheduler in the SPDP.*

Proof. Recall from [Lemma 6.2](#) that we need only consider the case that `ins` ends saturated, and recall the definition of t_* . For any scheduler `alg`, let B_{alg} denote the work that `alg` has left immediately before time t_* , and let K_{alg} be work that `alg` takes on tasks τ_i with $t_i \geq t_*$. Because `ins` ends saturated, we have

$$C_{\text{ins}} = t_* + (K_{\text{ins}} + B_{\text{ins}})/p.$$

Applying [Lemma 6.3](#) gives

$$t_* + (K_{\text{ins}} + B_{\text{ins}})/p \leq C_{\text{opt}} + (K_{\text{opt}} + B_{\text{ins}})/p. \quad (10)$$

So, to conclude, it suffices to show that $B_{\text{ins}} + K_{\text{opt}} \leq pC_{\text{opt}}$. Let S be the set of tasks that `ins` has present immediately before time t_* . Let $W = \sum_{\tau_i \in S} \sigma_i$. Clearly $B_{\text{ins}} \leq W$. On the other hand, `opt` must take at least W work on the tasks S , and can have made at most pt_* progress on these tasks by time t_* . Thus,

$$B_{\text{ins}} \leq W \leq pt_* + B_{\text{opt}}.$$

Therefore,

$$B_{\text{ins}} + K_{\text{opt}} \leq pt_* + B_{\text{opt}} + K_{\text{opt}} \leq pC_{\text{opt}}.$$

Using this in (10) gives $C_{\text{ins}} \leq 2C_{\text{opt}}$. \square

References

- [AGM11] S Anand, Naveen Garg, and Nicole Megow. “Meeting deadlines: How much speed suffices?” In: *Automata, Languages and Programming: 38th International Colloquium, ICALP 2011, Zurich, Switzerland, July 4-8, 2011, Proceedings, Part I 38*. Springer, 2011, pp. 232–243.

- [Asp+97] James Aspnes et al. “On-line routing of virtual circuits with applications to load balancing and machine scheduling”. In: *Journal of the ACM (JACM)* 44.3 (1997), pp. 486–504.
- [Awe+95] Baruch Awerbuch et al. “Load balancing in the L_p norm”. In: *Proceedings of IEEE 36th Annual Foundations of Computer Science*. IEEE. 1995, pp. 383–391.
- [BS83] Brenda S Baker and Jerald S Schwarz. “Shelf algorithms for two-dimensional packing problems”. In: *SIAM Journal on Computing* 12.3 (1983), pp. 508–525.
- [Car08] Ioannis Caragiannis. “Better bounds for online load balancing on unrelated machines”. In: *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*. 2008, pp. 972–981.
- [DLR23] Shichuan Deng, Jian Li, and Yuval Rabani. “Generalized unrelated machine scheduling problem”. In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2023, pp. 2898–2916.
- [DM07] Richard A. Dutton and Weizhen Mao. “Online scheduling of malleable parallel jobs”. In: *Proceedings of the 19th IASTED International Conference on Parallel and Distributed Computing and Systems*. PDCS '07. USA: ACTA Press, 2007, pp. 136–141. ISBN: 978-0-88986-704-8. (Visited on 02/02/2023).
- [GK10] Shouwei Guo and Liying Kang. “Online scheduling of malleable parallel jobs with setup times on two identical machines”. en. In: *European Journal of Operational Research* 206.3 (Nov. 2010), pp. 555–561. ISSN: 0377-2217. DOI: 10.1016/j.ejor.2010.03.005. URL: <https://www.sciencedirect.com/science/article/pii/S037722171000175X> (visited on 02/09/2023).
- [Gra69] R. L. Graham. “Bounds on Multiprocessing Timing Anomalies”. In: *SIAM Journal on Applied Mathematics* 17.2 (1969), pp. 416–429. DOI: 10.1137/0117039. eprint: <https://doi.org/10.1137/0117039>. URL: <https://doi.org/10.1137/0117039>.
- [Gup+17] Varun Gupta et al. “Stochastic online scheduling on unrelated machines”. In: *Integer Programming and Combinatorial Optimization: 19th International Conference, IPCO 2017, Waterloo, ON, Canada, June 26-28, 2017, Proceedings 19*. Springer. 2017, pp. 228–240.
- [Gup+20] Varun Gupta et al. “Greed works—online algorithms for unrelated machine stochastic scheduling”. In: *Mathematics of operations research* 45.2 (2020), pp. 497–516.
- [Gup+21] Anupam Gupta et al. “Stochastic load balancing on unrelated machines”. In: *Mathematics of Operations Research* 46.1 (2021), pp. 115–133.
- [HP08] Johann L Hurink and Jacob Jan Paulus. “Online algorithm for parallel job scheduling and strip packing”. In: *Approximation and Online Algorithms: 5th International Workshop, WAOA 2007, Eilat, Israel, October 11-12, 2007. Revised Papers 5*. Springer. 2008, pp. 67–74.
- [HS76] Ellis Horowitz and Sartaj Sahni. “Exact and Approximate Algorithms for Scheduling Nonidentical Processors”. In: *J. ACM* 23.2 (Apr. 1976), pp. 317–327. ISSN: 0004-5411. DOI: 10.1145/321941.321951. URL: <https://doi.org/10.1145/321941.321951>.
- [IL23] Sungjin Im and Shi Li. “Improved approximations for unrelated machine scheduling”. In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2023, pp. 2917–2946.
- [KW24] William Kuszmaul and Alek Westover. “Scheduling Jobs with Work-Inefficient Parallel Solutions”. In: *Proceedings of the 36th ACM Symposium on Parallelism in Algorithms and Architectures*. 2024, pp. 101–111.
- [LST90] Jan Karel Lenstra, David B Shmoys, and Éva Tardos. “Approximation algorithms for scheduling unrelated parallel machines”. In: *Mathematical programming* 46 (1990), pp. 259–271.

- [LT94] Walter Ludwig and Prasoon Tiwari. “Scheduling malleable and nonmalleable parallel tasks”. In: *Proceedings of the fifth annual ACM-SIAM symposium on Discrete algorithms*. 1994, pp. 167–176.
- [Mol] Marco Molinaro. “Stochastic ℓ_p Load Balancing and Moment Problems via the L-Function Method”. In: *Proceedings of the 2019 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 343–354. DOI: 10.1137/1.9781611975482.22. eprint: <https://epubs.siam.org/doi/pdf/10.1137/1.9781611975482.22>. URL: <https://epubs.siam.org/doi/abs/10.1137/1.9781611975482.22>.
- [MRT99] Gregory Mounie, Christophe Rapine, and Dennis Trystram. “Efficient approximation algorithms for scheduling malleable tasks”. In: *Proceedings of the eleventh annual ACM symposium on Parallel algorithms and architectures*. SPAA ’99. New York, NY, USA: Association for Computing Machinery, 1999, pp. 23–32. ISBN: 978-1-58113-124-6. DOI: 10.1145/305619.305622. URL: <https://doi.org/10.1145/305619.305622> (visited on 02/02/2023).
- [PSM20] Daniel R. Page, Roberto Solis-Oba, and Marten Maack. “Makespan minimization on unrelated parallel machines with simple job-intersection structure and bounded job assignments”. In: *Theoretical Computer Science* 809 (2020), pp. 204–217. ISSN: 0304-3975. DOI: <https://doi.org/10.1016/j.tcs.2019.12.009>. URL: <https://www.sciencedirect.com/science/article/pii/S0304397519307844>.
- [SS02] Andreas S. Schulz and Martin Skutella. “Scheduling Unrelated Machines by Randomized Rounding”. In: *SIAM Journal on Discrete Mathematics* 15.4 (2002), pp. 450–469. DOI: 10.1137/S0895480199357078. eprint: <https://doi.org/10.1137/S0895480199357078>. URL: <https://doi.org/10.1137/S0895480199357078>.
- [Tur+94a] John Turek et al. “Scheduling parallel tasks to minimize average response time”. In: *Proceedings of the fifth annual ACM-SIAM symposium on Discrete algorithms*. 1994, pp. 112–121.
- [Tur+94b] John Turek et al. “Scheduling parallelizable tasks to minimize average response time”. In: *Proceedings of the sixth annual ACM symposium on Parallel algorithms and architectures*. SPAA ’94. New York, NY, USA: Association for Computing Machinery, 1994, pp. 200–209. ISBN: 978-0-89791-671-4. DOI: 10.1145/181014.181331. URL: <https://doi.org/10.1145/181014.181331> (visited on 02/09/2023).
- [TWY92] John Turek, Joel L Wolf, and Philip S Yu. “Approximate algorithms scheduling parallelizable tasks”. In: *Proceedings of the fourth annual ACM symposium on Parallel algorithms and architectures*. 1992, pp. 323–332.
- [VHW14] Nodari Vakhania, Jose Hernandez, and Frank Werner. “Scheduling Unrelated Machines with Two Types of Jobs”. In: *International Journal of Production Research* 52 (Feb. 2014), pp. 1–9. DOI: 10.1080/00207543.2014.888789.
- [Yao77] Andrew Chi-Chin Yao. “Probabilistic computations: Toward a unified measure of complexity”. In: *18th Annual Symposium on Foundations of Computer Science (sfcs 1977)*. IEEE Computer Society. 1977, pp. 222–227.
- [YCZ18] Deshi Ye, Danny Z. Chen, and Guochuan Zhang. “Online scheduling of moldable parallel tasks”. en. In: *Journal of Scheduling* 21.6 (2018). Publisher: Springer, pp. 647–654. URL: https://ideas.repec.org/a/spr/jsched/v21y2018i6d10.1007_s10951-018-0556-2.html (visited on 02/03/2023).
- [YHZ09] Deshi Ye, Xin Han, and Guochuan Zhang. “A note on online strip packing”. In: *Journal of Combinatorial Optimization* 17.4 (2009), pp. 417–423.
- [Zha+22] Xiaoyan Zhang et al. “Randomized selection algorithm for online stochastic unrelated machines scheduling”. In: *Journal of Combinatorial Optimization* (2022), pp. 1–16.

A Barriers Against Improved Schedulers

In this section we show that the schedulers of [Section 4](#) and [Section 5](#) are optimal among natural restricted classes of schedulers. This highlights what changes must be made to the schedulers in order to have hopes of achieving better competitive ratios.

First we show that among *non-procrastinating* eventually-committing schedulers (i.e., eventually-committing schedulers with the property that whenever tasks are present, they will run at least one task), the scheduler [Scheduler 4.1](#) is optimal.

Proposition A.1. *Fix $\varepsilon > 0$. Let $\xi \approx 1.677$ denote the real root of the polynomial $2x^3 - 3x^2 - 1$. There is no deterministic $(\xi - \varepsilon)$ -competitive non-procrastinating eventually-committing scheduler.*

Proof. It suffices to consider the case that $\varepsilon < .001$. Fix a non-procrastinating eventually-committing scheduler \mathbf{alg} . Assume towards contradiction that \mathbf{alg} is $(\xi - \varepsilon)$ -competitive. We now describe a TAP \mathcal{T} on which $C_{\mathbf{alg}} \geq \xi C_{\text{opt}}$. The TAP starts with $\tau_1 = (1/\xi, 1/\xi^2, 0)$. Next, let $\tau_2 = (1 - \varepsilon^2, 1/\xi - \varepsilon^2, \varepsilon^2)$. Finally, at each time $t \in [\xi + 1/\xi - 2, 1 - \varepsilon^2] \cap (\mathbb{N}\varepsilon^2)$, give a task $\tau = (\infty, \varepsilon^2, t)$.

We now argue that \mathbf{alg} must run all the tasks on the fast machine. Because \mathbf{alg} is a non-procrastinating $(\xi - \varepsilon)$ -competitive scheduler, \mathbf{alg} must instantly start τ_1 on the fast machine (in case there are no tasks after τ_1). Now we argue that \mathbf{alg} runs τ_2 on the fast machine as well. Suppose that \mathbf{alg} starts τ_2 on a slow machine at some time t with

$$1 - \varepsilon^2 + t > (\xi - \varepsilon)C^t. \quad (11)$$

Then, \mathbf{alg} would not be $(\xi - \varepsilon)$ -competitive on the truncated TAP \mathcal{T}^t . Thus, \mathbf{alg} must not start τ_2 on a slow machine at any time t satisfying (11). We now show that (11) holds for all $t \geq \varepsilon$, thus proving that \mathbf{alg} must run τ_2 on the fast machine. For $t \in [\varepsilon^2, \xi + 1/\xi - 2)$ we have $C^t \leq 1/\xi$, and $1 - \varepsilon^2 + t \geq 1$, so (11) holds. For $t \geq \xi + 1/\xi - 2$ we have

$$C^t \leq \min(1, t - \xi + 2 + \varepsilon^2).$$

Thus, it suffices to show:

$$1 - \varepsilon^2 + t > (\xi - \varepsilon) \min(1, t - \xi + 2 + \varepsilon^2). \quad (12)$$

To show (12) it suffices to check (12) for $t = \xi - 1 - \varepsilon^2$ (by monotonicity of the inequality on either side of $t = \xi - 1$). At $t = \xi - 1 - \varepsilon^2$ (12) is:

$$\xi - 2\varepsilon^2 > \xi - \varepsilon,$$

which is true because $\varepsilon < .001$.

We have now shown that \mathbf{alg} runs all tasks in \mathcal{T} on the fast machine. Thus, (by definition of ξ)

$$C_{\mathbf{alg}} \geq 1/\xi^2 + 1/\xi - \varepsilon^2 + 1 - (1/\xi + \xi - 2) - \varepsilon^2 = \xi - 2\varepsilon^2.$$

However, $C_{\text{opt}} \leq 1$. This contradicts the assumption that \mathbf{alg} is $(\xi - \varepsilon)$ -competitive. \square

Now, we show that the 1.5-competitive scheduler of [Section 5](#) is optimal among never-committing schedulers that don't cancel tasks on slow machines.

Proposition A.2. *Let \mathbf{alg} be a deterministic never-committing scheduler that never cancels serial tasks. Then, for any $\varepsilon > 0$, there is a TAP \mathcal{T} with $n \leq O(1)$ on which \mathbf{alg} has is not $(1.5 - \varepsilon)$ -competitive.*

Proof. It suffices to consider the case that $\varepsilon < .001$. The TAP is defined as follows. First, $\tau_1 = (2, 1, 0)$. Then, for each time $t \in [\varepsilon^2, 1 - \varepsilon^2] \cap \mathbb{N}\varepsilon^2$, a task $\tau = (\infty, \varepsilon^2, t)$ arrives. We will show that if \mathbf{alg} starts τ_1 on a slow machine at any time t then \mathbf{alg} is not $(1.5 - \varepsilon)$ -competitive on \mathcal{T}^t . We show this by considering two cases.

Case 1: \mathbf{alg} starts τ_1 on a slow machine at time $t \in [0, 1]$.

If \mathbf{alg} does this, then $C_{\mathbf{alg}}^{\mathcal{T}^t} \geq 2 + t$. However, $C^t \leq t + 1 + \varepsilon^2$. Thus,

$$C_{\mathbf{alg}}^{\mathcal{T}^t}/C^t \geq \frac{2+t}{t+1+\varepsilon^2} \geq \frac{3}{2+\varepsilon^2} > 1.5 - \varepsilon.$$

So \mathbf{alg} cannot start τ_1 on a slow machine at this time.

Case 2: \mathbf{alg} starts τ_1 on a slow machine at time $t \geq 1$.

If \mathbf{alg} does this, then $C_{\mathbf{alg}} \geq 2 + t$. However, $C_{\text{opt}} \leq 2$. Thus,

$$C_{\mathbf{alg}}/C_{\text{opt}} \geq 1.5.$$

In conclusion, \mathbf{alg} must run τ_1 on the fast machine. But then

$$C_{\mathbf{alg}} \geq 3 - \varepsilon^2 > (1.5 - \varepsilon)C_{\text{opt}} = (1.5 - \varepsilon)2,$$

a contradiction. \square

B Lower Bounds from [KW24]

In this section we state, for the reader's convenience, the lower bounds from [KW24] against instantly- and eventually- committing schedulers.

Proposition B.1 (Kuzmaul, Westover [KW24]). *Fix $\varepsilon > 0$. There is no deterministic $(2 - \varepsilon)$ -competitive instantly-committing scheduler.*

Proof. Consider an n -task TAP where for each $i \in [n]$, the i -th task has $\sigma_i = 2^i, \pi_i = 2^{i-1}$, and the arrival times are all very close to 0. For each $i \in [n]$, it is possible to handle the first i tasks in the TAP with completion time 2^{i-1} . Thus, a $(2 - \varepsilon)$ -competitive scheduler cannot afford to run task τ_i on a slow machine. So, a $(2 - \varepsilon)$ -competitive scheduler must run all tasks on the fast machine, giving completion time at least $2^n - 1$ on this TAP, while $C_{\text{opt}} \leq 2^{n-1}$. For large enough n this implies that the scheduler is not actually $2 - \varepsilon$ competitive. \square

Proposition B.2 (Kuzmaul, Westover [KW24]). *Fix $\varepsilon > 0$. There is no deterministic $(\phi - \varepsilon)$ -competitive eventually-committing scheduler, where $\phi \approx 1.618$ is the golden ratio.*

Proof. Suppose that **alg** is a $(\phi - \varepsilon)$ -competitive eventually-committing scheduler. Let $\tau_1 = (\phi, 1, 0)$; if there are no further tasks, **alg** must run τ_1 on the fast machine, starting at some time $t_0 \leq 1/\phi$. Let $\tau_2 = (\infty, \phi - t_0, t_0)$. On this TAP, $C_{\text{opt}} = \phi$, while $C_{\text{alg}} \geq \phi + 1 = \phi^2$. So **alg** is not $(\phi - \varepsilon)$ -competitive. \square

C Randomized Lower Bounds

In this section we give lower bounds against randomized schedulers. Our main tool is Yao's minimax principle [Yao77], which allows us to prove a lower bound on the competitive ratio by exhibiting a distribution over TAPs, and showing that any deterministic scheduler has poor expected cost on a random TAP drawn from the distribution.

Proposition C.1. *For any $\varepsilon > 0.03$, there is no $(5/3 - \varepsilon)$ -competitive DoA scheduler, even with randomization.*

Proof. Fix $N = 25$. For $k \in \mathbb{N}$, define \mathcal{T}_k to be a length k TAP with $\sigma_i = 2^{i+1}, \pi_i = 2^i$. Let \mathcal{D} denote the following distribution over TAPs: choose $k \in [N]$ uniformly randomly, and then output TAP \mathcal{T}_k . By brute force enumeration of all possible deterministic instantly-committing strategies, one can show that no such strategy is 1.637-competitive on this TAP. \square

Proposition C.2. *For any $\varepsilon > 0$, there is no $((1 + \sqrt{3})/4 - \varepsilon)$ -competitive eventually-committing scheduler, even with randomization.*

Proof. In the proof of [Proposition 5.1](#) we defined two TAPs, and showed that no deterministic eventually-committing scheduler is $((1 + \sqrt{3})/2 - \varepsilon)$ -competitive on both of the TAPs. One can show that if we randomly choose between the two TAPs of [Proposition 5.1](#), there is no deterministic eventually-committing scheduler with expected competitive ratio $(1 + \sqrt{3})/4 - \varepsilon$. \square