

Decentralized Welfare-Optimizing Mechanisms for Scheduling Games*

Richard Cole [†] José R. Correa [‡] Vasilis Gkatzelis [§] Vahab Mirrokni [¶]
Neil Olver ^{||}

Abstract

Game Theory and its algorithmic counterpart, Mechanism Design, are by now standard tools for studying massive decentralized engineering systems such as the Internet. In this context, one of the most prominent issues concerns the design of mechanisms inducing socially efficient outcomes, such as that of Vickrey. Unfortunately these *socially optimal* mechanisms more often than not require full information and prohibitively large computational resources, something hopelessly impractical in large distributed environments.

In this work we study simple mechanisms that require only *local* information to operate properly. The specific setting in which we describe our work is a classic computationally hard scheduling problem in computer science (namely that of unrelated machine scheduling to minimize the total weighted completion time; here, selfish jobs choose a machine on which to be processed in order to minimize the time at which they are released), though we expect that our results extend to other settings with similar characteristics. Our findings indicate that one can design simple and local mechanisms inducing outcomes whose social cost is bounded by a small constant times that of the socially efficient solution. Furthermore, we exhibit a somewhat counter-intuitive phenomenon by noting that mechanisms yielding Pareto dominated outcomes may in fact enhance the overall performance of the system. The latter is explained by observing that the additional costs are actually well aligned with negative externalities that players impose on others. We also find that the use of randomization in the design of local mechanisms can lead to yet more improvement in the total welfare.

Most of the policies we consider have attractive properties—notably, they induce exact potential games. This fact, along with the other game-theoretic insights gained, allow us to give a new combinatorial approximation algorithm for the underlying optimization problem. This constitutes an example where ideas from economics are key to the development of algorithms for computationally difficult combinatorial optimization problems.

*This work was supported in part by NSF grants CCF0830516 and CCF1115849, by FONDECYT grant 1090050, and by Andreas Mentzelopoulos Scholarships for the University of Patras.

[†]cole@cims.nyu.edu. Courant Institute, New York University, N.Y.

[‡]jcorrea@dii.uchile.cl. Departamento de Ingeniería Industrial, Universidad de Chile.

[§]gkatz@cims.nyu.edu. Courant Institute, New York University, N.Y.

[¶]mirrokni@google.com. Google Research, New York, N.Y.

^{||}olver@math.mit.edu. Department of Mathematics, MIT.

1 Introduction

1.1 Multiagent resource allocation

The efficient allocation of resources has long been one of the main goals of Economics, Operations Research, and Computer Science. Recently, the study of *multiagent resource allocation* [14], which aims to achieve efficient allocation of resources among a set of self-interested agents, has been receiving increasing attention. Different types of resources, agent preferences, measures of efficiency, and allocation procedures give rise to different settings. One important distinction among allocation procedures, which will also play an important role in our work, is whether they are *centralized* or *decentralized*.

It is the centralized approach that has been traditionally followed, especially in Computer Science and Operations Research. Here, all related information is gathered (or elicited) by one entity which then aims to compute an allocation that optimizes some objective. Research along these lines has attempted to understand the computational complexity of such processes, drawing the boundary between what is tractable and what is not. For settings where this information is private and needs to be elicited from the users the main concern has been truthful implementation.

On the other hand, as the world becomes increasingly interconnected, a much more (geographically) distributed set of resources has become available, leading to a rising need for decentralized allocation processes. One of the main reasons is that centralized control in such distributed systems, examples of which include the Grid [32] and PlanetLab [46], severely impacts their scalability. The obvious drawback of decentralization, besides the limited communication among the allocation processes which may lead to inefficient solutions, is that such systems are more prone to strategic manipulation by their users. Our goal in this paper is to measure the inefficiency of the resource allocations that arise as equilibrium points of the induced game and to apply decentralized mechanisms aiming to minimize it.

1.2 Price of anarchy and coordination mechanisms

Given a game and a social cost function, the *price of anarchy* [44] has become a standard way to measure the deterioration due to selfish behavior. The price of anarchy of an instance of a game is simply the worst case ratio of the social cost at an equilibrium to that at the social optimum. The price of anarchy of a game is then the worst-case value of the price of anarchy of any instance of the game.

The games that arise within massive distributed systems like the ones mentioned above can take different forms but it is well known that strategic behavior by the users often leads to significant inefficiencies in the final allocation, i.e., a high price of anarchy. Since there is no centralized “benevolent dictator” to enforce the good behavior of the participating users, incentive design becomes a critical tool for achieving efficiency. To this end, several approaches have been proposed, including some enforcing strategies on a fraction of users as a Stackelberg strategy [7, 43, 48, 57] and others using monetary transfers [8, 20, 30, 13]. The primary drawback of these methods is the need for global knowledge of the system: the mechanism itself is still centralized.

To alleviate this problem Chistodoulou, Koutsoupias and Nanavati [18] proposed a different approach, which they called *coordination mechanisms*. More specifically, they consider settings where users can choose (part of) which resource they are going to be allocated. A strategy profile then corresponds to an assignment of users to resources and a coordination mechanism aims to provide the incentives that lead to more efficient equilibrium allocations. What makes coordination mechanisms a purely decentralized solution for this setting is the fact that they consist of several independent and decentralized allocation processes, one for each of the available resources. That is,

each resource has its own decentralized policy which decides how the resource will be allocated to the users that seek it. These policies eschew any form of side payment and thus the manipulation of the users' incentives is instead achieved only through appropriate regulation of how much of their chosen resource they are allocated. This means that, for example, instead of requiring a player to contribute some form of payment, the policy would choose to provide a less valuable allocation of its resource. We stress again that these policies only consider local information, i.e. the allocation choice is a function only of the users that demand it, thus remaining completely oblivious to the state of the rest of the system.

Given this decentralized regulation of the resource allocation, different coordination mechanisms induce different games, and thus lead to different sets of equilibrium allocations. Even for a fixed set of users and resources, two different coordination mechanisms may yield very different user costs for exactly the same strategy profile. It is these mechanism-specific costs that determine the social cost of a given equilibrium allocation. In evaluating the efficiency of these equilibrium points, one needs a well defined benchmark to compare this social cost against. The definition of the price of anarchy of the induced game would point to the social optimum w.r.t. these same costs (specific to the mechanism). The price of anarchy of a coordination mechanism is instead defined not as the price of anarchy of the induced game, but as the worst case ratio of the social cost at an equilibrium to the optimal social cost that could possibly be achieved by the centralized optimization approach. We will sometimes use the synonym *coordination ratio* as a reminder of this distinction.

1.3 Machine scheduling

In this paper, we study methods for the efficient allocation of a type of resource aiming to model machines or servers of some distributed system (such as the Internet) which offer a service that its users need. All the user tasks or jobs need to be scheduled on such machines and competition for these machines may lead to delays in servicing the users' demands. This model of *machine scheduling* [47] has been extensively studied since the 1950's, and it is generally considered a canonical model for studying settings related to job scheduling and processing.

In this model there are n jobs and m machines; each job must be assigned to a single machine. The processing time of a job can differ depending on the machine it is executed on; let p_{ij} denote the processing time of job j on machine i . Each job can also have an associated weight, which may be interpreted as a measure of importance. A resource allocation for such a problem instance (a *schedule*) consists of an assignment of jobs to machines, and a specification of the order according to which the jobs on each machine will be processed. For any such assignment and ordering, the *completion time* of each job is now determined; if on machine i , the assigned jobs are, in order, j_1, \dots, j_r , then the completion time of job j_1 is p_{ij_1} , the completion time of job j_2 is $p_{ij_1} + p_{ij_2}$, and so on.

The most basic model is that of *identical* machines, where the processing time of any job is the same on all of the machines. In the more general model of *related* machines each machine has a speed, and the processing time of a job on a machine is proportional to the speed of that machine. The main scheduling model that we study is *unrelated* machine scheduling in which the processing times are arbitrary, thus capturing all the above models as special cases.

In this work we consider the *scheduling game* that is induced due to the lack of centralized control. Each job is a fully informed player wanting to minimize its individual completion time, and its set of strategies correspond to the set of machines. A job's completion time on a machine depends not only on the strategies chosen by other players (in particular, which other players chose that machine), but also on the order that the jobs are run on the machine¹; in other words, this is

¹Actually, we will allow schedules that are more general than just executing the jobs in some order, but this

a setting with *externalities*. The cost of a job will be its weighted completion time; its completion time multiplied by its weight. The objective function that research on this setting has mostly focused on is the *makespan*, i.e. the maximum completion time over all jobs, which, for unweighted jobs, corresponds to the egalitarian social cost (see Section 1.5). This paper studies the utilitarian social cost instead, i.e. the unweighted or weighted sums of completion times.

A coordination mechanism for this setting is a set of *local policies*, one per machine, specifying how the jobs assigned to that machine are scheduled. We will actually consider the slightly more restrictive *strongly local policies*. For such policies the schedule on machine i must be a function of only the processing times p_{ij} and weights w_j of the jobs assigned to the machine. In contrast, in simply *local policies* the schedule on a machine may in addition depend on the full vector $\mathbf{p}_j = (p_{1j}, p_{2j}, \dots, p_{im})$ of processing times of each job j assigned to the machine. This is a somewhat weaker notion of locality, providing the policies with more information about the given problem instance.

1.4 Our Results

We begin by studying Smith’s Rule, a policy according to which machines process jobs in increasing order of their processing time to weight ratio. This is a natural first candidate to analyze since it is known that, for any given assignment of jobs to machines, this is the policy that minimizes our social cost function [56]. We prove that the coordination ratio for this policy is exactly 4, improving upon a result by Correa and Queyranne [23], who showed the same bound but for the less general model of restricted related machines (see Section 2). The constant coordination ratio for the weighted sum of completion times is in sharp contrast to the mostly negative results that have been shown for the makespan objective discussed above; in particular, no coordination mechanism is known that achieves a constant coordination ratio for makespan.

We also show that if we restrict ourselves to deterministic policies which always run jobs one after the other in some order, regardless of how this ordering depends on the weights and processing times of the assigned jobs—then this factor of 4 cannot be improved. We overcome this barrier in two ways; the first is deterministic, and adds artificial delays; the second is randomized, and achieves an even better total welfare.

The deterministic policy is most naturally described as a *preemptive* one. In our context, preemption really refers to time multiplexing: the machine runs the jobs “in parallel”, dividing its processing resources amongst the active jobs. The preemptive policy that we consider (**ProportionalSharing**) splits the processing capacity of a machine among its uncompleted jobs in proportion to their weights. This generalizes the **EqualSharing** policy [26], which splits the processing capacity equally amongst the jobs; what **ProportionalSharing** does for the unweighted case. We uncover a close connection of this policy to Smith’s Rule, allowing us to apply a similar proof strategy, but yielding a significantly improved coordination factor of 2.618. This improvement using preemption is somewhat counter-intuitive if one considers the fact that, for any given assignment of jobs to machines, any preemptive schedule is Pareto dominated by the schedule that a non-preemptive policy produces. This result is also in contrast to the makespan case, where even in the unweighted case the **EqualSharing** policy achieves a coordination ratio of $\Theta(m)$ [26], no better than Smith’s rule. To make sense of this phenomenon we show that, for a fixed assignment, the cost that each job suffers according to this very natural preemptive policy actually equals the cost that it would suffer if Smith’s Rule were instead being used, plus the externalities that this job would cause to the jobs that would have been scheduled after it. Each job is therefore forced to

simplifies the discussion for now.

internalize externalities that it causes to jobs on the same machine, leading to improved incentives. We also show an improved bound of 2.5 for the coordination ratio in the unweighted case.

On the other hand, we show that, under some restrictions, no deterministic policy can achieve a factor better than 2.166. To break this new barrier we consider a policy we call **Rand**, in which jobs are randomly (but non-uniformly) ordered, based on their processing time to weight ratio. This randomized policy also forces jobs which have a higher priority according to Smith’s Rule to suffer delays due to externalities that they cause, again leading to better incentives and even more efficient equilibrium allocations. One of the benefits of randomization is that although the jobs are made to suffer for (part of) their externalities, yet the schedule that the policy produces is always Pareto efficient. We give a bound of $32/15 \approx 2.133$ for the coordination ratio of this policy, a significant improvement over **ProportionalSharing**. In addition, in the case where the weighted sum of processing times is negligible compared to the total cost, our randomized policy has a much better coordination ratio of $\pi/2$, which is tight. The proofs here are perhaps the most interesting, involving a connection to the classical *Hilbert matrix*.

We prove all of the upper bound results in a common framework that brings out the structure in the scheduling games we consider. Once the framework has been set up, our proofs become short and elegant, and we anticipate that the approach may prove useful elsewhere too. We are able to relate the games induced by each of the policies we consider to certain inner product spaces. Proving upper bounds on the price of anarchy then becomes much simpler, in most cases involving an application of Cauchy-Schwartz and some form of “norm distortion” inequality to relate back to the Smith’s rule cost. While we present our proofs for pure strategies and pure Nash equilibria, we observe that all the results can be stated within the smoothness framework of Roughgarden [49] (see Section 2). This implies that all the bounds hold for more general equilibrium concepts including mixed Nash equilibria and correlated equilibria.

The game obtained when using Smith’s rule as the policy has a defect: it does not necessarily possess pure Nash equilibria [23]. Nevertheless, we show that the other policies we consider all induce exact potential games, giving another indication that **ProportionalSharing** and **Rand** are very natural policies. In fact, we can use these properties, along with other game-theoretic insights we have gained to give a result for the underlying *centralized* optimization problem.

From a purely centralized optimization perspective, the problem of minimizing the weighted sum of completion times has been extensively studied. The problem is APX-hard [38] on unrelated machines, and the current best polynomial time algorithm has an approximation factor of $\frac{3}{2}$ [53, 54]. All previous constant-factor approximation algorithms are based on rounding linear or convex programs. Complementing all these known non-combinatorial approximation algorithms, we design a new *combinatorial* $(2+\epsilon)$ -approximation algorithm for optimizing the weighted sum of completion times on unrelated machines.

In designing our approximation algorithm we take advantage of the fact that coordination mechanisms are related to local search algorithms. Starting from an initial solution, a local search algorithm iteratively moves to neighboring solutions which improve the global objective. This is based on a neighborhood relation that is defined on the set of solutions. Now, if one considers the strategy profiles of the game induced by the coordination mechanism as solutions, the best-response moves of the users in this game implicitly define the set of possible local moves. The speed of convergence and the approximation factor of local search algorithms for scheduling problems have been studied mainly for the makespan objective function [25, 27, 29, 40, 50, 52, 58, 1, 6]. Our combinatorial approximation algorithm for the weighted sum of completion times is the first local search algorithm for this problem, and is different from the previously studied algorithms for the makespan objective. The neighborhood implicitly defined by the coordination mechanism at hand

is non-trivial and it seems unlikely that such a simple algorithm could be designed without the initial game-theoretic intuition.

1.5 Related Work

Previous work on scheduling games mainly concerned the makespan social cost. One of the first such games to be considered was the one induced by the **Makespan** policy [44], according to which all jobs are released at the same time; each job’s completion time is equal to the sum of the processing times of the jobs assigned to its machine. This scheduling game gathered significant attention, eventually leading to a sequence of tight price of anarchy bounds for different machine models [24, 33, 4]. The games induced by **Makespan** are also known as *load balancing games*. In their paper introducing coordination mechanism design [18], Christodoulou, Koutsoupias and Nanavati analyzed mechanisms for identical machines using the **ShortestFirst** and **LongestFirst** policies, which process jobs in non-decreasing and non-increasing order of their processing times respectively. Immorlica et al. [41] studied these three coordination mechanisms, along with a randomized one which orders jobs randomly in a uniform fashion, for several different machine scheduling models. They surveyed the known results for these settings, uncovering connections of these local policies with greedy and local search algorithms [40, 29, 50, 25, 1, 6, 9, 58]. Apart from price of anarchy related results, they also studied the speed of convergence to equilibria and the existence of pure Nash equilibria for the **ShortestFirst** and **LongestFirst** policies. Azar, Jain, and Mirrokni [5] showed that the **ShortestFirst** policy and in fact any *strongly* local fixed ordering policy (defined in Section 2) does not achieve a coordination ratio better than $\Omega(m)$. Additionally, they presented a non-preemptive local policy that achieves a coordination ratio of $O(\log m)$ and a policy that induces potential games and gives a coordination ratio of $O(\log^2 m)$. Caragiannis [11], among other results, showed an alternative coordination mechanism that guaranteed a coordination ratio of $O(\log m)$ for unrelated machines, while still inducing potential games. Fleischer and Svitkina [31] showed a lower bound of $\Omega(\log m)$ for all local fixed ordering policies, thus proving that Caragiannis’ mechanism is optimal w.r.t. the price of anarchy within this class. This bound had though already been overcome by Caragiannis [11] who presented a local preemptive policy with approximation factor $O(\log m / \log \log m)$. More recently, Dürr and Thang [26] proved that the **EqualSharing** policy results in potential games and achieves a coordination ratio of $\Theta(m)$ for the unrelated machines model. Whether there exists a coordination mechanism (even preemptive or randomized) achieving a constant coordination ratio for the makespan social cost function remains an interesting open question.

Our work concerns the utilitarian social cost, or (weighted) sum of completion times. For this objective, Correa and Queyranne [23] studied Smith’s rule for the restricted related machine model and they exhibited an instance for which the induced game does not possess a pure Nash equilibrium. They also presented bounds for the price of anarchy of **SmithRule** in this model. Finally, Hoeksma and Uetz [37] showed better price of anarchy bounds for the less general setting of unweighted jobs and related machines using **ShortestFirst**; the unweighted variant of **SmithRule**.

2 Preliminaries

Throughout this paper, let J be a set of n jobs to be scheduled on a set I of m machines. Let p_{ij} denote the processing time of job $j \in J$ on machine $i \in I$ and w_j denote its weight (or importance). The shorthand notation ρ_{ij} will be used for the ratio p_{ij}/w_j . Jobs that have both the same processing time and the same weight can be distinguished from one another only if they have been assigned a unique ID; otherwise, the jobs are called *anonymous*.

We will refer to the following standard scheduling models:

Identical machines. All machines are identical, meaning each job needs the same processing time on each machine: $p_{ij} = p_{i'j}$ for all $i, i' \in I$. The model of *restricted* identical machines is a variant according to which each job can be run only on some specified subset of machines.

Related machines. The machines may have different speeds, and the processing time of a job is inversely proportional to the speed: $p_{ij} = p_j/\sigma_i$, where σ_i represents the speed of machine i , and p_j the processing requirement of job j . The *restricted* related machines variant is again obtained by possibly restricting the set of machines to which each job can be assigned.

Unrelated machines. The processing times are arbitrary. This is the most general of these models. There is no need to distinguish between restricted and unrestricted variants, since we allow specifying that a job takes infinite time on a machine.

A *coordination mechanism* for this setting is a set of local policies, one for each machine. Each such policy determines how to schedule the set of jobs assigned to the machine it controls, thus defining the *completion time* c_j of each job j in that set. A coordination mechanism thereby gives rise to a *scheduling game* in which there are n agents (jobs) and each agent's strategy set is the set of machines I . A strategy profile corresponds to an assignment of jobs to machines, represented by a vector \mathbf{x} , where x_j gives the machine to which job j is assigned. Given such an assignment \mathbf{x} , the cost of job j is its weighted completion time, as determined by the policy on the machine x_j . We let $w_j c_j^\alpha(\mathbf{x})$ and $C^\alpha(\mathbf{x})$ denote the cost for player j and the social cost respectively, where $\alpha \in \{SR, PS, SF, ES, R\}$ denotes the policy, namely SmithRule, ProportionalSharing, ShortestFirst, EqualSharing and Rand, respectively.² The agent controlling each job aims to choose a strategy (i.e., a machine) that minimizes its cost or, in the case of randomized policies, its expected cost. The mechanisms that we analyze are designed with the goal of minimizing the utilitarian social cost, i.e. $C^\alpha(\mathbf{x}) = \sum_{j \in J} w_j c_j^\alpha(\mathbf{x})$.

A strategy profile \mathbf{x} of a scheduling game instance is a *pure Nash equilibrium* (PNE) if no player has an incentive to unilaterally change its strategy. Formally, if this instance is induced by coordination mechanism α , then for all $j \in J$ and all $i \in I$ we get $w_j c_j^\alpha(\mathbf{x}) \leq w_j c_j^\alpha(\mathbf{x}_{-j}, i)$, where the notation (\mathbf{x}_{-j}, i) denotes the assignment \mathbf{x} , except modified so that job j is assigned to machine i . Also, slightly abusing notation, let $X_i = \{j \in J \mid x_j = i\}$ denote the set of jobs that have chosen machine i in configuration \mathbf{x} , and define X_i^* analogously for \mathbf{x}^* .

In order to measure the efficiency of a coordination mechanism α for a given scheduling game instance, we study its social cost in PNE assignments. We are interested in the worst case ratio of the social cost in a PNE assignment divided by the optimal social cost achievable from a centralized optimization approach. It is known that the optimal solution of the centralized optimization problem schedules jobs on machines according to SmithRule [56], so the optimal social cost can be expressed as $C^{SR}(\mathbf{x}^*)$, where $\mathbf{x}^* = \arg \min_{\mathbf{x}'} C^{SR}(\mathbf{x}')$. If we also let $\mathbf{x} = \arg \max_{\mathbf{x}' \in E} C^\alpha(\mathbf{x}')$ be the worst equilibrium assignment w.r.t. the social cost induced by α , then this ratio is equal to $C^\alpha(\mathbf{x})/C^{SR}(\mathbf{x}^*)$ for the given game instance. Following the definition of [18] the (*pure*) *price of anarchy* or *coordination ratio* of coordination mechanism α is defined to be the maximum such ratio, taken over all the scheduling game instances that the mechanism may induce.

Adapting the work of Roughgarden [49] to this setting, we define a coordination mechanism α to be (λ, μ) -smooth if for any two configurations \mathbf{x} and \mathbf{x}^* ,

$$\sum_{j \in J} w_j c_j^\alpha(\mathbf{x}_{-j}, x_j^*) \leq \lambda C^{SR}(\mathbf{x}^*) + \mu C^\alpha(\mathbf{x}).$$

²The coordination mechanisms we study in this paper use the same local policy on each machine, so henceforth we refer to a coordination mechanism using the name of the policy.

If a coordination mechanism is (λ, μ) smooth, then this yields an upper bound of $\frac{\lambda}{1-\mu}$ which applies not only to its pure price of anarchy, but also to its *robust price of anarchy*. This, among other things, implies that it is not only the social cost of PNE that is bound to be at most $\frac{\lambda}{1-\mu}$ times the optimal social cost, but also the social cost of any correlated equilibrium [49].

A game is a *potential game* if there exists a potential function over strategy profiles such that any player’s deviation leads to a drop of the potential function if and only if its cost drops. A potential game is *exact* if after each move, the changes to the potential function and to the player’s cost are equal. It is easy to see that a potential game always possesses a pure Nash equilibrium, corresponding to a local minimum of the potential function.

2.1 Classification of policies.

It will be useful (particularly for discussing lower bounds) to identify the main classes of strongly local policies that will concern us in this work.

Fixed ordering policies³. These policies assign an order on *all* jobs, based on the jobs’ characteristics on the machine (processing time, weight, and possibly ID). Then, for a given assignment, the jobs assigned to the machine are executed according to this order. One motivation for these policies is that they satisfy the *independence of irrelevant alternatives (IIA)* property: for any pair of jobs, their relative ordering is independent of which other jobs are assigned to the machine. This property appears as an axiom in voting theory, bargaining theory and logic.

Flexible ordering policies. In this class, policies still execute jobs in some fixed order, but that order may depend arbitrarily on the set of jobs assigned to the machine. Here we require that the jobs on a machine are executed consecutively in some deterministic order. Moreover, we require that there be no idle time between jobs, and that jobs are released immediately upon completion. The reason for this restriction is to distinguish from the next case.

Preemptive. Preemption refers to the ability to suspend a job before it completes in order to execute another job.⁴ The initial job can then be resumed later. Preemption allows for *time multiplexing*: by switching between a number of jobs very quickly, the illusion is given that the jobs are being run simultaneously on the machine. In fact, the schedules that preemptive policies yield are Pareto dominated by those of policies that do not use preemption. Thus, as we explain in more detail in Section 4, such policies can equally well be considered as deterministic policies without preemption, but where jobs may be held back after completion; hence the restriction in our definition of flexible ordering policies.

Randomized. Here, the policy may schedule the jobs randomly, according to a distribution depending only on the processing times and weights of the jobs on the machine. While more general schedules are possible, it’s helpful to think of simply a random ordering of the jobs.

We also call, e.g., a coordination mechanism consisting of fixed ordering policies a “fixed ordering coordination mechanism”.

³These were called simply *ordering policies* in [5], but we wish to emphasise the distinction with flexible ordering policies, defined next.

⁴Note that, unlike in some literature on machine scheduling, preemption here does *not* imply that a job can be processed on a different machine after it is suspended.

3 Smith's Rule

Smith's rule is a fixed ordering policy that schedules jobs on machine i in increasing order of $\rho_{ij} = p_{ij}/w_j$. In the unweighted case, this reduces to the **ShortestFirst** policy. It is known that given an assignment of jobs to machines, in order to minimize the weighted sum of completion times, Smith's rule is optimal [56]. It is therefore natural to consider this policy as a good first candidate to study. Our first theorem shows that using this rule will result in Nash equilibria with social cost at most a constant factor of 4 away from the optimum.

All our proofs will use a common framework. The proof for Smith's rule is the simplest, but will introduce a number of aspects of this framework. We begin by defining the map $\varphi : I^J \rightarrow L_2([0, \infty))^I$, which maps a configuration to a vector of functions, as follows. If $\mathbf{f} = \varphi(\mathbf{x})$, then

$$f_i(y) = \sum_{j \in X_i: \rho_{ij} \geq y} w_j \quad (\text{recall that } \rho_{ij} = p_{ij}/w_j).$$

We let $\langle g, h \rangle := \int_0^\infty g(y)h(y)dy$ denote the usual inner product on L_2 , and in addition define $\langle \mathbf{f}, \mathbf{g} \rangle := \sum_{i \in I} \langle f_i, g_i \rangle$. In both cases, $\|\cdot\|$ refers to the induced norm. We also define

$$\eta(\mathbf{x}) = \sum_{j \in J} w_j p_{x_j j}.$$

We then have

Lemma 3.1. *For any configuration \mathbf{x} , $C^{SR}(\mathbf{x}) = \frac{1}{2}\|\varphi(\mathbf{x})\|^2 + \frac{1}{2}\eta(\mathbf{x})$.*

Proof. Let $\mathbf{f} = \varphi(\mathbf{x})$. We have

$$\begin{aligned} \|\varphi(\mathbf{x})\|^2 &= \sum_{i \in I} \int_0^\infty f_i(y)^2 dy \\ &= \sum_{i \in I} \sum_{j \in X_i} \sum_{k \in X_i} w_j w_k \int_0^\infty \mathbf{1}_{\rho_{ij} \geq y} \mathbf{1}_{\rho_{ik} \geq y} dy \\ &= \sum_{i \in I} \sum_{j \in X_i} \sum_{k \in X_i} w_j w_k \min\{\rho_{ij}, \rho_{ik}\} \\ &= \sum_{i \in I} \sum_{j \in X_i} w_j \left(2 \sum_{\substack{k \in X_i \\ \rho_{ik} \leq \rho_{ij}}} p_{ik} - p_{ij} \right) \\ &= 2C^{SR}(\mathbf{x}) - \eta(\mathbf{x}). \end{aligned} \tag{1}$$

The result follows. □

Theorem 3.2. *The price of anarchy of SmithRule for unrelated machines is at most 4.*

Proof. Let \mathbf{x} and \mathbf{x}^* be two assignments, with \mathbf{x} a Nash equilibrium, and write $\mathbf{f} = \varphi(\mathbf{x})$, $\mathbf{f}^* = \varphi(\mathbf{x}^*)$. We assume for simplicity that all jobs have distinct ratios (of processing time to weight). (This assumption is just for simplicity; alternatively, we could introduce a tie breaking rule.)

We first calculate a job j 's completion time according to \mathbf{x} , and use the Nash condition:

$$c_j^{SR} = \sum_{\substack{k: x_k = x_j \\ \rho_{x_k k} < \rho_{x_j j}}} p_{x_k k} + p_{x_j j} \leq \sum_{\substack{k: x_k = x_j^* \\ \rho_{x_k k} < \rho_{x_j^* j}}} p_{x_k k} + p_{x_j^* j}.$$

So

$$\begin{aligned}
C^{SR}(\mathbf{x}) &= \sum_j w_j c_j^{SR} \leq \sum_{i \in I} \sum_{j \in X_i^*} \left(\sum_{\substack{k \in X_i \\ \rho_{ik} < \rho_{ij}}} w_j w_k \frac{p_{ik}}{w_k} + p_{ij} w_j \right) \\
&\leq \sum_{i \in I} \sum_{j \in X_i^*} \sum_{k \in X_i} w_j w_k \min\{\rho_{ik}, \rho_{ij}\} + \sum_{i \in I} \sum_{j \in X_i^*} p_{ij} w_j \quad (2) \\
&= \sum_{i \in I} \sum_{j \in X_i^*} \sum_{k \in X_i} w_j w_k \int_0^\infty \mathbf{1}_{\rho_{ij} \geq y} \mathbf{1}_{\rho_{ik} \geq y} dy + \eta(\mathbf{x}^*) \\
&= \langle \mathbf{f}^*, \mathbf{f} \rangle + \eta(\mathbf{x}^*). \quad (3)
\end{aligned}$$

Now applying Cauchy-Schwartz, followed by the inequality $ab \leq a^2 + b^2/4$ for $a, b \geq 0$, we obtain

$$\begin{aligned}
C^{SR}(\mathbf{x}) &\leq \|\mathbf{f}\| \|\mathbf{f}^*\| + \eta(\mathbf{x}^*) \\
&\leq \|\mathbf{f}^*\|^2 + \frac{1}{4} \|\mathbf{f}\|^2 + \eta(\mathbf{x}^*) \\
&\leq 2C^{SR}(\mathbf{x}^*) + \frac{1}{2} C^{SR}(\mathbf{x}) \quad \text{by Lemma 3.1.}
\end{aligned}$$

Hence $C^{SR}(\mathbf{x}) \leq 4C^{SR}(\mathbf{x}^*)$. □

Notice that in this proof, the cost $C^{SR}(\mathbf{x})$ of an assignment \mathbf{x} is closely related to the norm of $\varphi(\mathbf{x})$, and the inequality obtained from the Nash condition is bounded by a term involving the inner product $\langle \varphi(\mathbf{x}), \varphi(\mathbf{x}^*) \rangle$. This will be a common feature of all our proofs.

For simplicity, the proof above was written as a pure price of anarchy bound; \mathbf{x} was taken to be a pure Nash equilibrium. However, it is clear that the proof in fact yields a robust price of anarchy bound, as defined by Roughgarden [49]. More precisely, the above proof shows that SmithRule is $(2, 1/2)$ -smooth.

The following result, proved in Appendix A, shows that no fixed ordering coordination mechanism can do better than SmithRule. (In fact, the result can be extended to all *flexible* ordering coordination mechanisms, but we will not discuss this here.) This also implies that the bound of Theorem 3.2 is tight.

Theorem 3.3. *The pure price of anarchy of any set of fixed ordering policies is at least 4. This is true even for the case of restricted identical machines with unweighted jobs.*

We note that for the unit weight case, a constant upper bound on the coordination ratio of Smith's rule can be obtained via a reduction from the *priority routing model* of Farzad et al. [28]. However, the resulting bound is not optimal.

4 Improvements with Preemption and Randomization

4.1 Preemptive Coordination Mechanism

In this section, we study the power of preemption (or equivalently, delays) and present the following preemptive policy, named **ProportionalSharing**. Jobs are scheduled in parallel using time-multiplexing, and, at any moment in time, each uncompleted assigned job receives a fraction of the processor time equal to its weight divided by the total weight of uncompleted jobs on the machine. In the unweighted case, this gives the **EqualSharing** policy.

We will show that **ProportionalSharing** has a better coordination ratio than any fixed ordering policy. These results create a clear dichotomy between such policies and **ProportionalSharing**. This

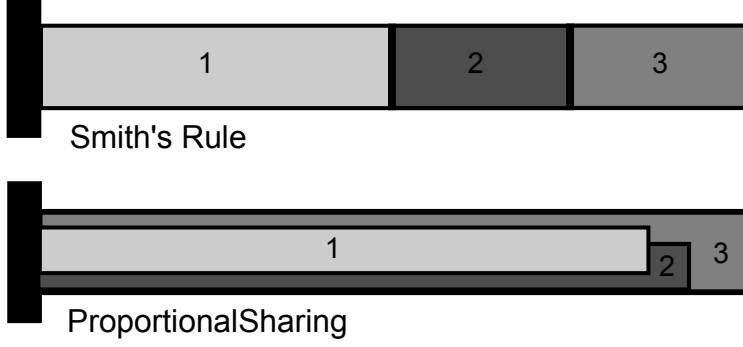


Figure 1: Three jobs with $p_1 = 4$ and $w_1 = 7$ for the first, $p_2 = 2$ and $w_2 = 3$ for the second, and $p_3 = 2$ and $w_3 = 2$ for the third, being scheduled using **SmithRule** and **ProportionalSharing**.

may seem counter-intuitive at first, since, given an assignment of jobs to machines, the schedule produced by **ProportionalSharing** is Pareto dominated by that of **SmithRule**. To be more precise, on each machine, every job apart from the one that **SmithRule** would schedule last strictly prefers **SmithRule** to **ProportionalSharing**; the one scheduled last is indifferent between the two schedules. This can also be seen in Figure 1 which compares how the two policies would schedule a given set of jobs.

Lemma 4.1. *Given an assignment \mathbf{x} , the weighted completion time of a job j on some machine i using **ProportionalSharing** (whether currently assigned there or not) is*

$$w_j c_j^{PS}(\mathbf{x}_{-j}, i) = \sum_{k \in X_i \setminus \{j\}} w_j w_k \min\{\rho_{ij}, \rho_{ik}\} + w_j p_{ij}. \quad (4)$$

Proof. First, observe that for two jobs k and k' , if $\rho_{ik} \leq \rho_{ik'}$ then job k will complete before (or at the same time as) job k' . To see this, consider the situation at the time when the earlier of the two jobs is completed. Let q and q' be the amount of processing time that has been allocated to k and k' by this time. Then $q' = \frac{w_{k'}}{w_k} q$. If k is not completed, then $q < w_k \rho_{ik}$, and so $q' < w_{k'} \rho_{ik} \leq p_{ik'}$, and k' is not completed either.

Let t be the time when job j is completed. All jobs k with $\rho_{ik} \leq \rho_{ij}$ have completed by this time; thus each such job has received p_{ik} units of processing time. On the other hand, all jobs k with $\rho_{ik} > \rho_{ij}$ are not yet complete at time t , and for each w_j units of processing time job j receives, job k receives w_k units. Thus by time t , the processing time spent on any such job k will be exactly $\frac{p_{ij} w_k}{w_j}$. Since the total processing time is the sum of the processing times allocated to all the jobs, we thus have that

$$t = \sum_{\substack{k \in X_i \setminus \{j\} \\ \rho_{ik} \leq \rho_{ij}}} p_{ik} + \sum_{\substack{k \in X_i \\ \rho_{ik} > \rho_{ij}}} \frac{w_k}{w_j} p_{ij} + p_{ij}$$

Thus

$$\begin{aligned} w_j c_j^{PS}(\mathbf{x}_{-j}, i) &= \sum_{\substack{k \in X_i \setminus \{j\} \\ \rho_{ik} \leq \rho_{ij}}} w_j p_{ik} + \sum_{\substack{k \in X_i \\ \rho_{ik} > \rho_{ij}}} w_k p_{ij} + w_j p_{ij} \\ &= \sum_{k \in X_i \setminus \{j\}} w_j w_k \min\{\rho_{ij}, \rho_{ik}\} + w_j p_{ij}. \end{aligned}$$

□

A better understanding of why `ProportionalSharing` performs better despite the Pareto inefficiency of the schedules it produces can be obtained by observing the following corollary of Lemma 4.1.

Corollary 4.2. *Given an assignment \mathbf{x} , the weighted completion time of a job j on some machine i using `ProportionalSharing` is*

$$w_j c_j^{PS}(\mathbf{x}_{-j}, i) = w_j c_j^{SR}(\mathbf{x}_{-j}, i) + \sum_{\substack{k \in X_i \\ \rho_{ik} > \rho_{ij}}} w_k p_{ij}$$

This corollary precisely quantifies what cost, in addition to the `SmithRule` cost, this job is forced to suffer. A closer look reveals that this additional cost (the rightmost term) is exactly equal to externalities that job j would cause if the assignment of jobs to machines remained the same but `SmithRule` was used instead. That is, the sum for each job k that would have been scheduled after job j ($\rho_{ik} > \rho_{ij}$), of the cost increase that job j causes to that job ($w_k p_{ij}$). From this perspective, `ProportionalSharing` can be thought of (and also implemented) as using `SmithRule` to determine the processing order, but then delaying the release of each job after it is completed until the additional cost equals these externalities. Since we already know that, for any given assignment, `SmithRule` would produce the social welfare maximizing schedule, one may expect that our preemptive policy exactly “internalizes the externalities” of the players and should therefore lead to the optimal assignment in equilibrium. The reason why this is not the case is that the participation of job j in the game does not cause externalities only to jobs that are assigned to its machine. Nevertheless, our policies are necessarily oblivious to what the state of the system is beyond the machine they control, so these “local externalities” may be the best possible alternative. By taking these local externalities into consideration, `ProportionalSharing` better aligns the interests of a player with those of the system, leading not only to better assignments than `SmithRule` but also to a better social cost, despite the (otherwise unnecessary) delays suffered. Other advantages of this coordination mechanism include that, unlike `SmithRule`, it can handle anonymous jobs, and that the games it induces always possess PNE.

From Lemma 4.1 we also immediately obtain the following corollary (note the factor 2 difference compared to the main term in the cost of `Smith`’s rule in Lemma 3.1), which will be used in proving the two subsequent theorems:

Corollary 4.3. *For any assignment \mathbf{x} , $C^{PS}(\mathbf{x}) = \|\varphi(\mathbf{x})\|^2$.*

Theorem 4.4. *The price of anarchy of `ProportionalSharing` for unrelated machines is at most $\phi + 1 = \frac{3+\sqrt{5}}{2} \approx 2.618$. Moreover, this bound is tight even for the restricted related machines model.*

Proof. Let \mathbf{x} be an equilibrium assignment, and \mathbf{x}^* any arbitrary assignment. From the Nash condition,

$$\begin{aligned} C^{PS}(\mathbf{x}) &\leq \sum_{j \in J} w_j c_j^{PS}(\mathbf{x}_{-j}, \mathbf{x}^*) \\ &\leq \sum_{i \in I} \sum_{j \in X_i^*} \left(\sum_{k \in X_i} w_j w_k \min\{\rho_{ij}, \rho_{ik}\} + p_{ij} \right) && \text{by Lemma 4.1} \\ &= \langle \varphi(\mathbf{x}), \varphi(\mathbf{x}^*) \rangle + \eta(\mathbf{x}^*) && \text{by (1).} \end{aligned} \tag{5}$$

Following the same method of analysis as for Smith's Rule, we obtain

$$\begin{aligned}
C^{PS}(\mathbf{x}) &\leq \|\varphi(\mathbf{x})\| \|\varphi(\mathbf{x}^*)\| + \eta(\mathbf{x}^*) \\
&\leq \alpha \|\varphi(\mathbf{x}^*)\|^2 + \frac{1}{4\alpha} \|\varphi(\mathbf{x})\|^2 + \eta(\mathbf{x}^*) \\
&\leq 2\alpha C^{SR}(\mathbf{x}^*) + \frac{1}{4\alpha} C^{PS}(\mathbf{x}) + (1 - \alpha)\eta(\mathbf{x}^*) \\
&\leq (1 + \alpha)C^{SR}(\mathbf{x}^*) + \frac{1}{4\alpha} C^{PS}(\mathbf{x}),
\end{aligned}$$

using that $\eta(\mathbf{x}^*) \leq C^{SR}(\mathbf{x}^*)$. Setting $\alpha = (1 + \sqrt{5})/4$ yields $C^{PS}(\mathbf{x})/C^{SR}(\mathbf{x}^*) \leq \frac{3+\sqrt{5}}{2}$.

The tightness of this bound follows from a construction in [12], where in fact they show that even if C^{PS} is used for the cost of *OPT*, i.e., we consider the ratio $C^{PS}(\mathbf{x})/C^{PS}(\mathbf{x}^*)$, this can be arbitrarily close to $1 + \phi$. \square

The reader will observe how similar the proof above was to the proof of Theorem 3.2, once the relevant costs have been described in terms of the inner product. In particular, (5) is precisely identical to (3) used in the proof for Smith's rule. In the proof of Theorem 3.2, (2) can be interpreted as a kind of symmetrization step, which is needed since the inner product $\langle \varphi(\mathbf{x}), \varphi(\mathbf{x}^*) \rangle$ is symmetric. *ProportionalSharing* is already symmetric in the appropriate sense, and so there is a tighter connection between the Nash condition and the inner product. This same symmetry property will be shared in the randomized policy we consider in the next section.

Notice also that since *SmithRule* and *ProportionalSharing* were described in terms of the same mapping and inner product, it was very easy to relate $\|\varphi(\mathbf{x}^*)\|$ back to the cost of *SmithRule*. This will be less straightforward for the randomized policy discussed in the next section.

The coordination ratio obtained may remind the reader of similar bounds for weighted congestion games [2]. It is important to stress that since our bounds do not follow from these results. What *can* be deduced from Azar et al. [2] is that $C^{PS}(\mathbf{x})/C^{PS}(\mathbf{x}^*) \leq \phi + 1$ for any Nash assignment \mathbf{x} . Our result shows that the ratio does not get any worse when we compare against the stronger benchmark of $C^{SR}(\mathbf{x}^*)$.

In the case of equal weights, we obtain a slightly improved bound using the following lemma instead of the Cauchy-Schwartz inequality. This is a tighter version of an inequality initially used by Christodoulou and Koutsoupias [17]:

Lemma 4.5. *For every pair of nonnegative integers k and k^* ,*

$$k^*(k + 1) \leq \frac{1}{3}k^2 + \frac{5}{6}k^*(k^* + 1).$$

Proof. This translates to showing that for all nonnegative integers k and k^* ,

$$5k^{*2} + 2k^2 - 6k^*k - k^* \geq 0.$$

Rewriting the left hand side as $2(k - \frac{3}{2}k^*)^2 + \frac{1}{2}k^{*2} - k^*$, we see immediately that the inequality holds for $k^* \geq 2$ and $k^* = 0$. In the case $k^* = 1$, the required inequality simplifies to $k^2 - 3k + 2 \geq 0$ which is true for all integral k . \square

Theorem 4.6. *The price of anarchy of *EqualSharing* for unrelated machines is at most 2.5. This bound is tight even for the restricted related machines model.*

Proof. Let \mathbf{x} be any some assignment, and let $\mathbf{f} = \varphi(\mathbf{x})$. Since $w_j = 1$ for all j ,

$$f_i(y) = |\{j \in J : x_j = i \text{ and } p_{ij} \geq y\}|,$$

and also

$$\begin{aligned}
\eta(\mathbf{x}) &= \sum_i \sum_{j \in X_i} p_{ij} \\
&= \sum_{i \in I} \int \sum_{j \in X_i} \mathbf{1}_{y \leq p_{ij}} dy \\
&= \sum_{i \in I} \int f_i(y) dy.
\end{aligned}$$

Thus from Corollary 4.3 and Lemma 3.1,

$$C^{ES}(\mathbf{x}) = \int f_i^2(y) dy \quad \text{and} \quad C^{SF}(\mathbf{x}) = \frac{1}{2} \int f_i(y)(f_i(y) + 1) dy. \quad (6)$$

Now suppose \mathbf{x} is a Nash equilibrium; take \mathbf{x}^* to be any assignment, with $\mathbf{f}^* = \varphi(\mathbf{x}^*)$. We continue from (5):

$$\begin{aligned}
C^{ES}(\mathbf{x}) &\leq \langle \mathbf{f}, \mathbf{f}^* \rangle + \eta(\mathbf{x}^*) \\
&= \int f_i(y)(f_i^*(y) + 1) dy \\
&\leq \int \frac{1}{3} f_i^2(y) + \frac{5}{6} f_i^*(y)(f_i^*(y) + 1) dy && \text{by Lemma 4.5} \\
&= \frac{1}{3} C^{ES}(\mathbf{x}) + \frac{5}{3} C^{SF}(\mathbf{x}^*) && \text{by (6)}.
\end{aligned}$$

This gives a price of anarchy bound of 2.5.

The tightness of the bound follows from Theorem 3 of [12]. The authors present a load balancing game lower bound, which is equivalent to assuming that all jobs have unit processing times and the machines are using `EqualSharing`; thus the same proof yields a (pure) price of anarchy lower bound for restricted related machines and unweighted jobs. \square

Once again, all of the upper bounds also hold for the robust price of anarchy. On the negative side, we have the following (the proof of which can be found in Appendix A). Call a coordination mechanism *prompt* if on any machine, the completion time of all jobs assigned to the machine are never larger than the sum of processing times of jobs on the machine. Equivalently, each machine uses its full capacity and does not delay the release of a job after its completion.

Proposition 4.7. *When jobs are anonymous, the coordination ratio of any deterministic prompt coordination mechanism is at least 13/6.*

4.2 Randomized Coordination Mechanism

In this section we examine the power of randomization and present `Rand`, a randomized policy that satisfies the following property: if two jobs j and j' are assigned to machine i , then

$$\mathbb{P}\{j \text{ precedes } j' \text{ in the ordering}\} = \frac{\rho_{ij'}}{\rho_{ij} + \rho_{ij'}}. \quad (7)$$

Recall $\rho_{ij} = p_{ij}/w_j$. A distribution over orderings with this property can be constructed as follows. Starting from the set of jobs X_i assigned to machine $i \in I$, select job $j \in X_i$ with probability $\rho_{ij}/\sum_{k \in X_i} \rho_{ik}$, and schedule j at the end. Then remove j from the list of jobs, and repeat this

process. Note that this policy is different from a simple randomized policy that orders jobs uniformly at random. In fact, this simpler policy is known to give an $\Omega(m)$ price of anarchy bound for the makespan objective [41], and the same family of examples developed in [41] gives an $\Omega(m)$ lower bound for this policy in our setting.

As we show below, this randomized policy outperforms any deterministic strongly local policy that has the “prompt” property defined above. In an attempt to explain this success, it is straightforward to verify that, unlike `ProportionalSharing`, this policy produces Pareto efficient schedules. One can actually show that it Pareto dominates `ProportionalSharing`. Yet, contrary to `SmithRule`, for any pair of jobs assigned to the same machine, there is positive probability that any one of the two is scheduled later, thus suffering a delay because of the other. In this sense, `Rand` gives high priority jobs the incentive to avoid crowded machines if they have better alternatives, but it does so without introducing very long delays.

Theorem 4.8. *The price of anarchy when using the `Rand` policy is at most $32/15 = 2.133\dots$. Moreover, if the sum of the processing times of the jobs is negligible compared to the social cost of the optimal solution, this bound improves to $\pi/2$, which is tight.*

The high level approach for obtaining these upper bounds is in exactly the same spirit as the previous section: find an appropriate mapping φ from an assignment into a convenient inner product space. To make the mapping and inner product space easier to describe, we assume in this section that the processing times have been scaled such that the ratios ρ_{ij} are all integral. We also take κ large enough so that, except for infinite processing times, $\rho_{ij} \leq \kappa$ for all $i \in I, j \in J$. These assumptions are inessential and easily removed.

An inner product space. The map φ we use gives the *signature* for each machine: in the unweighted case, this simply describes how many jobs of each size are assigned to the machine.

Definition 4.9. Given an assignment \mathbf{x} , its *signature* $\varphi(\mathbf{x}) \in \mathbb{R}_+^{m \times \kappa}$ is a vector indexed by a machine i and a processing time over weight ratio r ; we denote this component by $\varphi(\mathbf{x})_r^i$. Its value is then defined as

$$\varphi(\mathbf{x})_r^i := \sum_{\substack{j \in X_i \\ \rho_{ij}=r}} w_j.$$

We also let $\varphi(\mathbf{x})^i$ denote the vector $(\varphi(\mathbf{x})_0^i, \varphi(\mathbf{x})_1^i, \dots, \varphi(\mathbf{x})_\kappa^i)$.

Let M be the $\kappa \times \kappa$ matrix given by

$$M_{rs} = \frac{rs}{r+s}.$$

Lemma 4.10. *Let \mathbf{x} be some assignment, and let $\mathbf{u} = \varphi(\mathbf{x})$. If job j is assigned to machine i , its expected completion time is given by*

$$c_j^R = (M\mathbf{u}^i)_{\rho_{ij}} + \frac{1}{2}p_{ij}.$$

If j is not assigned to i , then its expected completion time upon switching to i would be

$$c_j^R = (M\mathbf{u}^i)_{\rho_{ij}} + p_{ij}.$$

Proof. We consider case (i); (ii) is similar. So $x_j = i$. The expected completion time of job j on machine i is

$$\begin{aligned} c_j^R &= \sum_{k \in X_i \setminus \{j\}} p_{ik} \mathbb{P}\{\text{job } k \text{ ahead of job } j\} + p_{ij} \\ &= \sum_{k \in X_i \setminus \{j\}} p_{ik} \frac{\rho_{ij}}{\rho_{ij} + \rho_{ik}} + p_{ij} \\ &= \sum_{k \in X_i} p_{ik} \frac{\rho_{ij}}{\rho_{ij} + \rho_{ik}} + \frac{1}{2} p_{ij}. \end{aligned}$$

We can rewrite this in terms of the signature as

$$c_j^R = \sum_s u_s^i M_{\rho_{ij}s} + \frac{1}{2} p_{ij} = (M\mathbf{u}^i)_{\rho_{ij}} + \frac{1}{2} p_{ij}. \quad \square$$

A crucial observation is the following:

Lemma 4.11. *The matrix M is positive definite.*

Proof. Let D be the diagonal matrix with $D_{rr} = r$. Then we have $M = DHD$, where the $\kappa \times \kappa$ matrix H is given by $H_{rs} = \frac{1}{r+s}$. This is a submatrix of the infinite Hilbert matrix $\left(\frac{1}{r+s-1}\right)_{r,s \in \mathbb{N}}$. The Hilbert matrix has the property that it is *totally positive* [16], meaning that the determinant of any submatrix is positive. It follows immediately that H is positive definite, and hence so is M . \square

Thus we may define an inner product by

$$\langle \mathbf{u}, \mathbf{v} \rangle_R := \sum_{i \in I} (\mathbf{u}^i)^T M \mathbf{v}^i, \quad (8)$$

with an associated norm $\|\cdot\|_R$. In addition, the total cost $\sum_j w_j c_j^R(\mathbf{x})$ of an assignment \mathbf{x} may be written in the convenient form

$$C^R(\mathbf{x}) = \|\varphi(\mathbf{x})\|_R^2 + \frac{1}{2} \eta(\mathbf{x}). \quad (9)$$

Competitiveness of Rand on a single machine. An interesting extra complication that occurs with this policy is that, unlike with `ProportionalSharing`, the inner product describing the cost of `Rand` is quite different to the one describing `SmithRule`. Since we ultimately need to compare against `OPT`, we need to relate the cost of `Rand` and `SmithRule`. For this reason, the performance of `Rand` on even a *single* machine, compared to `SmithRule`, plays an important role.

So suppose we have n jobs with processing times p_j and weight w_j , for $j \leq n$. The signature \mathbf{u} is given by just $u_r = \sum_{j: p_j/w_j=r} w_j$. By considering (1), it follows that the weighted sum of completion times according to `SmithRule` is

$$\mathbf{u}^T S \mathbf{u} + \frac{1}{2} \sum_j w_j p_j,$$

where $S_{rs} = \frac{1}{2} \min\{r, s\}$. Compare this to the corresponding formula for `Rand`:

$$\mathbf{u}^T M \mathbf{u} + \frac{1}{2} \sum_j w_j p_j.$$

The extra $\sum_j w_j p_j$ terms only help, and in fact turn out to be negligible in the worst case example; ignoring them, the goal is to determine $\sup_{\mathbf{u} \geq \mathbf{0}} \frac{\mathbf{u}^T M \mathbf{u}}{\mathbf{u}^T S \mathbf{u}}$. So the question is closely related to the worst-case distortion between two norms (not quite, because of the nonnegativity constraint).

Interestingly, it turns out that this problem has been considered, and solved, in a different context. In [19], Chung, Hajela and Seymour consider the problem of *self-organizing sequential search*. In order to prove a tight bound on the performance of the “move-to-front” heuristic compared to the optimal ordering, they show:

Theorem 4.12 ([19]). *For any sequence u_1, u_2, \dots, u_k with $u_r > 0$ for all r ,*

$$\sum_{r,s} u_r u_s \frac{rs}{r+s} < \frac{\pi}{4} \sum_{r,s} u_r u_s \min\{r, s\}.$$

(We also present a quite different proof of the theorem in Appendix B.) Thus on a single machine, Rand costs at most a factor $\frac{\pi}{2}$ more than SmithRule. Moreover, this is tight [34] (take $p_j = 1/j^2$, $w_j = 1$, and let $n \rightarrow \infty$). Of course, it follows immediately that for any number of machines and any assignment \mathbf{x} ,

$$C^R(\mathbf{x}) \leq \frac{\pi}{2} C^{SR}(\mathbf{x}). \quad (10)$$

All in all, we find that $\pi/2$ is a tight upper bound on the competitiveness of Rand on a single machine. The following lemma (which may also be cast as a norm distortion question), is much more easily demonstrated:

Lemma 4.13. *For any assignment \mathbf{x} , we have $C^R(\mathbf{x}) \leq 2C^{SR}(\mathbf{x}) - \eta(\mathbf{x})$.*

Proof. Consider a particular machine i . We have

$$\begin{aligned} \sum_{j,k \in X_i} w_j w_k \frac{\rho_{ij} \rho_{ik}}{\rho_{ij} + \rho_{ik}} &= \sum_{j \neq k \in X_i} w_j w_k \frac{\rho_{ij} \rho_{ik}}{\rho_{ij} + \rho_{ik}} + \frac{1}{2} \sum_{j \in X_i} w_j p_{ij} \\ &\leq \sum_{j \neq k \in X_i} w_j w_k \min\{\rho_{ij}, \rho_{ik}\} + \frac{1}{2} \sum_{j \in X_i} w_j p_{ij} \\ &= \sum_{j,k \in X_i} w_j w_k \min\{\rho_{ij}, \rho_{ik}\} - \frac{1}{2} \sum_{j \in X_i} w_j p_{ij}. \end{aligned}$$

Summing over all machines gives

$$C^R(\mathbf{x}) - \frac{1}{2} \eta(\mathbf{x}) \leq 2(C^{SR}(\mathbf{x}) - \frac{1}{2} \eta(\mathbf{x})) - \frac{1}{2} \eta(\mathbf{x})$$

from which the bound is immediate. \square

The upper bound. We are now ready to prove the main theorem of this section.

Proof of Theorem 4.8. Let \mathbf{x} be the assignment at a Nash equilibrium, and \mathbf{x}^* the assignment of the optimal solution, and let $\mathbf{u} = \varphi(\mathbf{x})$ and $\mathbf{u}^* = \varphi(\mathbf{x}^*)$.

From the Nash condition and Lemma 4.10, we obtain

$$\begin{aligned} C^R(\mathbf{x}) &\leq \sum_{j \in J} w_j c_j^R(\mathbf{x}_{-j}, \mathbf{x}_j^*) \\ &\leq \sum_{i \in I} \sum_{j \in X_i^*} w_j M(\mathbf{u}^i)_{\rho_{ij}} + \eta(\mathbf{x}^*) \\ &= \sum_{i \in I} (\mathbf{u}^{*i})^T M \mathbf{u}^i + \eta(\mathbf{x}^*) \\ &= \langle \mathbf{u}^*, \mathbf{u} \rangle_R + \eta(\mathbf{x}^*). \end{aligned}$$

Applying Cauchy-Schwartz,

$$\begin{aligned} C^R(\mathbf{x}) &\leq \|\mathbf{u}^*\|_R \|\mathbf{u}\|_R + \eta(\mathbf{x}^*) \\ &\leq \frac{2}{3} \|\mathbf{u}^*\|_R^2 + \frac{3}{8} \|\mathbf{u}\|_R^2 + \eta(\mathbf{x}^*), \end{aligned} \tag{11}$$

Now recalling the definition of φ and applying Lemma 4.13, we obtain

$$\begin{aligned} C^R(\mathbf{x}) &\leq \frac{2}{3}(C^R(\mathbf{x}^*) - \frac{1}{2}\eta(\mathbf{x}^*)) + \frac{3}{8}(C^R(\mathbf{x}) - \frac{1}{2}\eta(\mathbf{x})) + \eta(\mathbf{x}^*) \\ &\leq \frac{2}{3}(2C^{SR}(\mathbf{x}^*) - \frac{3}{2}\eta(\mathbf{x}^*)) + \frac{3}{8}(C^R(\mathbf{x}) - \frac{1}{2}\eta(\mathbf{x})) + \eta(\mathbf{x}^*) \\ &\leq \frac{4}{3}C^{SR}(\mathbf{x}^*) + \frac{3}{8}C^R(\mathbf{x}). \end{aligned}$$

This gives a coordination ratio of 32/15.

In the case where $\eta(\mathbf{x}^*)$ is negligible, we continue from (11):

$$C^R(\mathbf{x}) \leq \|\mathbf{u}^*\|_R \|\mathbf{u}\|_R \leq \sqrt{C^R(\mathbf{x}^*)} \cdot \sqrt{C^R(\mathbf{x})}.$$

Thus $C^R(\mathbf{x}) \leq C^R(\mathbf{x}^*)$ —the Nash equilibria are in fact optimal assignments when measured by C^R ! From (10) we obtain $C^R(\mathbf{x})/C^{SR}(\mathbf{x}^*) \leq \pi/2$. \square

As noted in Appendix A, a slight modification of the construction used to prove Proposition 4.7 can be used to show that the worst-case price of anarchy of Rand is at least 5/3.

5 Potential Games and an Algorithmic Application

Potential games. Under SmithRule it may happen that no pure Nash equilibrium exists [23]. Here we show that ProportionalSharing and Rand both induce exact potential games, which hence always have pure Nash equilibria. It also follows from this that certain natural best response dynamics very quickly converge to solutions which cost at most slightly larger than the price of anarchy bound; we discuss this in more detail at the end of this section.

The following theorem generalizes [26, Theorem 3], which addresses EqualSharing (i.e., the unweighted case).

Theorem 5.1. *The ProportionalSharing mechanism induces exact potential games, with potential*

$$\Phi^{PS}(\mathbf{x}) = \frac{1}{2}C^{PS}(\mathbf{x}) + \frac{1}{2}\eta(\mathbf{x}). \tag{12}$$

Likewise, the Rand mechanism yields exact potential games with potential

$$\Phi^R(\mathbf{x}) = \frac{1}{2}C^R(\mathbf{x}) + \frac{1}{2}\eta(\mathbf{x}). \tag{13}$$

Proof. We give the proof for ProportionalSharing; the proof for Rand is similar.

Consider an assignment \mathbf{x} and a job $j \in J$, and let i be the machine to which j is assigned. Define \mathbf{x}' as the assignment differing from \mathbf{x} only in that job j moves to some machine $i' \neq i$.

We may write the change in the potential function as

$$\Phi^{PS}(\mathbf{x}') - \Phi^{PS}(\mathbf{x}) = \sum_{k \in J} D_k + \frac{1}{2}w_j(p_{i'j} - p_{ij}), \tag{14}$$

where

$$D_k = \frac{1}{2}w_k (c_k^{PS}(\mathbf{x}') - c_k^{PS}(\mathbf{x})).$$

Consider a job $k \neq j$ on machine i . Since only job j left the machine, we have from Lemma 4.1 that

$$c_k^{PS}(\mathbf{x}') - c_k^{PS}(\mathbf{x}) = -w_j \min\{\rho_{ij}, \rho_{ik}\}.$$

Thus

$$\begin{aligned} \sum_{k \in X_i \setminus \{j\}} D_k &= -\frac{1}{2}w_j \sum_{k \in X_i \setminus \{j\}} w_k \min\{\rho_{ij}, \rho_{ik}\} \\ &= -\frac{1}{2}w_j(c_j^{PS}(\mathbf{x}) + p_{ij}). \end{aligned}$$

Similarly, considering jobs on i' yields

$$\begin{aligned} \sum_{k \in X_{i'}} D_k &= \frac{1}{2}w_j \sum_{k \in X_{i'}} w_k \min\{\rho_{i'j}, \rho_{i'k}\} \\ &= \frac{1}{2}w_j(c_j^{PS}(\mathbf{x}') - p_{i'j}). \end{aligned}$$

All other jobs are unaffected by the change, and so do not contribute to (14). Summing all terms (including D_j), we obtain

$$\Phi^{PS}(\mathbf{x}') - \Phi^{PS}(\mathbf{x}) = w_j(c_j^{PS}(\mathbf{x}') - c_j^{PS}(\mathbf{x})),$$

exactly the change in the cost of job j . □

A combinatorial approximation algorithm. Minimizing the *unweighted* sum of completion times is polynomial time solvable, even for unrelated machines [39, 10]. For identical parallel machines, the `ShortestFirst` policy leads to an optimal schedule at any pure Nash equilibrium [22, 41]. On the other hand, minimizing the weighted sum of completion times is NP-complete even for identical machines [45]. This special case does admit a *polynomial time approximation scheme* (PTAS) however: for any $\epsilon > 0$, a solution only a factor $1 + \epsilon$ more expensive than the optimal one can be found in polynomial time [55].

By contrast, the general case of unrelated machines case is APX-hard [38]—no PTAS is possible. A sequence of papers gave improving constant-factor approximation algorithms, all based on rounding a linear or convex programming relaxation. The first was a $16/3$ -approximation algorithm [35], based on rounding an appropriate linear programming relaxation. An improvement to $\frac{3}{2} + \epsilon$, again based on linear programming, was given in [51]. Finally the best currently known factor, a $\frac{3}{2}$ -approximation, was obtained based on a convex quadratic relaxation of the problem [53, 54].

In this section, we will give a very simple and *combinatorial* approximation algorithm. While it does not quite match the best factor of $\frac{3}{2}$, it achieves a factor of $2 + \epsilon$, for any $\epsilon > 0$.

The basic idea is as follows. Consider some policy (say, `Rand`). If we could find a Nash equilibrium to the game using this policy, this gives us an assignment which costs at most ρ times the optimum, where ρ is the price of anarchy bound that we have obtained for this policy. So if we can always find some Nash equilibrium in polynomial time, we obtain a factor ρ approximation algorithm. Of course, there is no longer any need to keep to the suboptimal scheduling given by the policy; once we have the Nash assignment \mathbf{x} , we can switch back to using `SmithRule`. By choosing the policy carefully, we are able to guarantee a gain in this step, and achieve a factor that is better than the best price of anarchy bound that we managed to achieve.

The policy we use, which we call `Approx`, is a variation of `ProportionalSharing` with some additional delays. Schedule the jobs exactly as in `ProportionalSharing`, but hold each job j back by an

amount equal to its processing time. In other words, the completion time of any job j under an assignment \mathbf{x} is

$$c_j^A(\mathbf{x}) = c_j^{PS}(\mathbf{x}) + p_{x_j j}.$$

Comparing Lemma 3.1 and Corollary 4.3, we see that

$$C^{PS}(\mathbf{x}) = 2C^{SR}(\mathbf{x}) - \eta(\mathbf{x}).$$

Thus

$$C^A(\mathbf{x}) = C^{PS}(\mathbf{x}) + \eta(\mathbf{x}) = 2C^{SR}(\mathbf{x}).$$

This will give us a saving of a factor of 2 when we switch from using the **Approx** policy to **SmithRule**. It turns out that **Approx** has a coordination ratio of 4; thus for any Nash equilibrium \mathbf{x} with respect to this policy, $C^{SR}(\mathbf{x}) \leq 2OPT$.

Unfortunately we do not know how to compute an equilibrium allocation to this game (similarly for **ProportionalSharing** and **Rand**, in fact). Despite this, we will show that a natural best response dynamic will converge very quickly to some assignment of cost at most $2OPT$; by running this dynamic for long enough, we can obtain a solution of cost at most $(2 + \epsilon)OPT$ in a polynomial number of steps, for any $\epsilon > 0$. This will follow from general results on the robust price of anarchy proved by Roughgarden [49], drawing on work by Awerbuch et al. [3] and Chien and Sinclair [15]. We will actually prove everything we need, primarily in order to be able to give a precise stopping condition for our algorithm, something which is not quite explicit in [49]. It also demonstrates that there is no difficulty in extending the proofs to price of anarchy bounds on coordination mechanisms rather than games, although this is fairly immediate from a consideration of the original proofs.

Consider the following natural best response dynamics: simply pick the job which can improve its disutility (weighted completion time) the most by deviating, and allow that job to move. Following [49] we will call this *maximum-gain* best response dynamics. Given some coordination mechanism α and configuration \mathbf{x} , let

$$\Delta_j^\alpha = w_j (c_j^\alpha(\mathbf{x}) - c_j^\alpha(\mathbf{x}_{-j}, x'_j)) \quad \text{for any } j \in J,$$

where x'_j is the best response move for player j . Also let

$$\Delta^\alpha(\mathbf{x}) = \sum_{j \in J} \Delta_j^\alpha.$$

Definition 5.2. An assignment \mathbf{x} is an ϵ -approximate equilibrium if $\Delta^\alpha(\mathbf{x}) < \epsilon C^\alpha(\mathbf{x})$.

The full algorithm using these dynamics is described in Algorithm 1.

Algorithm 1: A factor $2 + \epsilon$ approximation algorithm for minimizing the weighted sum of completion times.

- 1 Let \mathbf{x}_0 be an assignment where each job is assigned to a machine on which it has minimum processing time.
 - 2 Using the **Approx** policy, run basic dynamics until an $\epsilon/4$ -approximate equilibrium \mathbf{x} is obtained.
 - 3 Return assignment \mathbf{x} , scheduled according to **SmithRule**.
-

In order to bound the running time of our local-search algorithm we use the following theorem, slightly adapted from [49, Proposition 2.6], which is in turn based on [3].

Proposition 5.3. [49] *Let α be a (λ, μ) -smooth coordination mechanism, let \mathbf{x}^0 be any initial configuration, and let $\hat{\mathbf{x}}$ be the global minimizer of Φ^α . Then for any $\epsilon > 0$, maximum-gain best response dynamics generates an ϵ -approximate equilibrium \mathbf{x} in at most $O\left(\frac{n}{\epsilon} \log\left(\frac{\Phi^\alpha(\mathbf{x}^0)}{\Phi^\alpha(\hat{\mathbf{x}})}\right)\right)$ steps, and this assignment satisfies $C^\alpha(\mathbf{x}) \leq \frac{\lambda}{1-\mu-\epsilon} OPT$.*

Proof. By the definition of (λ, μ) -smooth, we have

$$\sum_{j \in J} w_j c_j^\alpha(\mathbf{x}_{-j}, x_j^*) \leq \mu C^\alpha(\mathbf{x}) + \lambda OPT.$$

Thus $\Delta^\alpha(\mathbf{x}) \geq (1 - \mu)C^\alpha(\mathbf{x}) - \lambda OPT$, and so if \mathbf{x} is an ϵ -approximate equilibrium,

$$\epsilon C^\alpha(\mathbf{x}) \geq \Delta^\alpha(\mathbf{x}) \geq (1 - \mu)C^\alpha(\mathbf{x}) - \lambda OPT,$$

implying the required cost bound.

Let \mathbf{x}^t be the assignment after t steps of basic dynamics. Suppose that \mathbf{x}^t is not an ϵ -approximate equilibrium, so $\Delta^\alpha(\mathbf{x}^t) > \epsilon C^\alpha(\mathbf{x}^t)$. Then if j is the player which can improve the most, we must have $\Delta_j^\alpha(\mathbf{x}^t) \geq \frac{\epsilon}{n} C^\alpha(\mathbf{x}^t)$. Then since $C^\alpha(\mathbf{x}^t) \leq \Phi^\alpha(\mathbf{x}^t)$ and $\Phi^\alpha(\mathbf{x}^{t+1}) = \Phi^\alpha(\mathbf{x}^t) - \Delta_j^\alpha(\mathbf{x}^t)$, we have

$$\Phi^\alpha(\mathbf{x}^{t+1}) \leq (1 - \frac{\epsilon}{n})\Phi^\alpha(\mathbf{x}^t).$$

Thus if no ϵ -approximate equilibrium is found in the first T steps,

$$\Phi^\alpha(\hat{\mathbf{x}}) \leq \Phi^\alpha(\mathbf{x}^T) \leq (1 - \frac{\epsilon}{n})^T \Phi^\alpha(\mathbf{x}^0).$$

This yields the required bound on the number of steps. \square

We omit the proofs of the following two lemmas, which are essentially identical to those of Theorem 5.1 and Theorem 4.4 respectively:

Lemma 5.4. *The Approx coordination mechanism induces an exact potential game, with potential function*

$$\Phi^A(\mathbf{x}) = \frac{1}{2}C^A(\mathbf{x}) + \eta(\mathbf{x}).$$

Lemma 5.5. *The Approx coordination mechanism is $(3, 1/4)$ -smooth.*

We may now prove the main result of this section.

Theorem 5.6. *For any $0 < \epsilon < 1$, Algorithm 1 runs in polynomial time, and returns a schedule of cost at most $(2 + \epsilon)OPT$.*

Proof. We first argue that $\Phi^A(\mathbf{x}^0) \leq (n + 1)\Phi^A(\hat{\mathbf{x}})$, where $\hat{\mathbf{x}}$ is a global minimizer of Φ^A . Consider two jobs j and k assigned to some machine i under \mathbf{x}^0 , such that j is processed before k under Smith's rule. Then $\rho_{ij} \leq \rho_{ik}$. The total contribution to the cost $C^{SR}(\mathbf{x}^0)$ due to the delay of job k by job j is $w_k p_{ij}$. But we have

$$w_k p_{ij} = w_j w_k \rho_{ij} \leq \frac{1}{2}(w_j^2 + w_k^2)\rho_{ij} \leq \frac{1}{2}(w_j p_{ij} + w_k p_{ik}).$$

Summing over all pairs of jobs processed on the same machine, the total contribution to $C^{SR}(\mathbf{x}^0)$ due to delays is at most

$$\sum_{i \in I} \sum_{j \neq k \in X_i^0} \frac{1}{2}(w_j p_{ij} + w_k p_{ik}) \leq (n - 1) \sum_{i \in I} \sum_{j \in X_i^0} w_j p_{ij} = (n - 1)\eta(\mathbf{x}^0).$$

Hence $C^{SR}(\mathbf{x}^0) \leq n\eta(\mathbf{x}^0)$, and so

$$\begin{aligned} \Phi^A(\mathbf{x}^0) &\leq C^{SR}(\mathbf{x}^0) + \eta(\mathbf{x}^0) \\ &\leq (n+1)\eta(\mathbf{x}^0) \\ &\leq (n+1)\eta(\hat{\mathbf{x}}) && \text{(since } \mathbf{x}^0 \text{ minimizes } \eta) \\ &\leq (n+1)\Phi^A(\hat{\mathbf{x}}). \end{aligned}$$

Since **Approx** is $(3, 1/4)$ -smooth, by Lemma 5.3 the algorithm returns an assignment \mathbf{x} of cost $C^A(\mathbf{x}) \leq \frac{3}{1-1/4-\epsilon/4} OPT$ in $O(\frac{n \log n}{\epsilon})$ steps of best response dynamics. Thus the algorithm runs in polynomial time, and simplifying, $C^{SR}(\mathbf{x}) \leq (2 + \epsilon) OPT$. \square

Since we also have robust price of anarchy bounds for **ProportionalSharing** and **Rand**, and these both induce exact potential games, fast convergence statements can be made for these policies as well. As well as the maximum-gain best response dynamics used in Algorithm 1, best response dynamics where a random player is chosen at each round also leads to convergence, with high probability [49].

6 Concluding remarks

On mapping machines to edges of a parallel link network, the machine scheduling problem for the case of related machines becomes a special case of general selfish routing games. In this context, the ordering policies on machines correspond to local queuing policies at the edges of the network. From this perspective, it would be interesting to generalize our results to network routing games. Designing such local queuing policies would be an important step toward more realistic models of selfish routing games when the routing happens over time [36, 28, 42, 21]. We hope that our new technique along with the policies proposed in this paper could serve as a building block toward this challenging problem.

All the mechanisms discussed here are strongly local. For the case of the makespan objective, one can improve the coordination ratio from $\Theta(m)$ to $\Theta(\log m)$ by using local policies instead of just strongly local policies. It remains open whether there are local policies that perform even better than our strongly local ones. In particular, we do not know of any local policy that does better than **Rand**.

Acknowledgements. We thank Tanmoy Chakraborty and Andy Skrzypacz for helpful discussions and Ioannis Caragiannis for pointing out related literature. The fourth author also thanks Yossi Azar for initial discussions about the subject of study of this paper. Part of this work was done while the second and last authors were visiting EPFL; we thank Fritz Eisenbrand for his hospitality.

References

- [1] J. Aspnes, Y. Azar, A. Fiat, S.A. Plotkin, and O. Waarts. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *J. ACM*, 44(3):486–504, 1997.
- [2] B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *STOC*, pages 57–66, 2005.

- [3] B. Awerbuch, Y. Azar, A. Epstein, V.S. Mirrokni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *ACM Conference on Electronic Commerce*, pages 264–273, 2008.
- [4] B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in worst-case equilibria. *Theor. Comput. Sci.*, 361(2-3):200–209, 2006.
- [5] Y. Azar, K. Jain, and V.S. Mirrokni. (Almost) optimal coordination mechanisms for unrelated machine scheduling. In *SODA*, pages 323–332, 2008.
- [6] Y. Azar, J. Naor, and R. Rom. The competitiveness of on-line assignments. *Journal of Algorithms*, 18:221–237, 1995.
- [7] A. Bagchi. Stackelberg differential games in economic models. *Springer-Verlag*, 1984.
- [8] M. Beckman, C.B. McGuire, and C.B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
- [9] A. Borodin, M. Nielsen, and C. Rackoff. (incremental) priority algorithms. In *SODA*, pages 752 – 761, 2002.
- [10] J. Bruno, E.G. Coffman, and R. Sethi. Scheduling independent tasks to reduce mean finishing time. *Communications of the ACM*, 17:382–387, 1974.
- [11] I. Caragiannis. Efficient coordination mechanisms for unrelated machine scheduling. In *SODA*, pages 815–824, 2009.
- [12] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight bounds for selfish and greedy load balancing. In *ICALP (1)*, pages 311–322, 2006.
- [13] I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos. Taxes for linear atomic congestion games. *ACM Transactions on Algorithms*, 7(1):13, 2010.
- [14] Y. Chevaleyre, P.E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J.A. Padget, S. Phelps, J.A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica (Slovenia)*, 30(1):3–31, 2006.
- [15] S. Chien and A. Sinclair. Convergence to approximate nash equilibria in congestion games. In *SODA*, pages 169–178, 2007.
- [16] M.-D. Choi. Tricks or treats with the hilbert matrix. *The American Mathematical Monthly*, 90(5):301–312, May 1983.
- [17] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *STOC*, pages 67–73, 2005.
- [18] G. Christodoulou, E. Koutsoupias, and A. Nanavati. Coordination mechanisms. *Theor. Comput. Sci.*, 410(36):3327–3336, 2009.
- [19] F. Chung, D. Hajela, and P. D. Seymour. Self-organizing sequential search and hilbert’s inequalities. *Journal of Computer and System Sciences*, 36(2):148–157, April 1988.
- [20] R. Cole, Y. Dodis, and T. Roughgarden. How much can taxes help selfish routing? *J. Comput. Syst. Sci.*, 72(3):444–467, 2006.

- [21] R. Cominetti, J.R. Correa, and O. Larre. Existence and uniqueness of equilibria for flows over time. In *ICALP (part II)*, pages 552–563, 2011.
- [22] R.W. Conway, W.L. Maxwell, and L.W. Miller. *Theory of Scheduling*. Addison-Wesley, Reading, MA, 1967.
- [23] J.R. Correa and M. Queyranne. Efficiency of equilibria in restricted uniform machine scheduling with minsum social cost. 2010 (manuscript).
- [24] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *SODA*, pages 413–420, 2002.
- [25] E. Davis and J.M. Jaffe. Algorithms for scheduling tasks on unrelated processors. *J. ACM*, 28(4):721–736, 1981.
- [26] C. Dürr and N.K. Thang. Non-clairvoyant scheduling games. In *SAGT*, pages 135–146, 2009.
- [27] E. Even-dar, A. Kesselman, and Y. Mansour. Convergence time to nash equilibria. In *ICALP*, pages 502–513, 2003.
- [28] B. Farzad, N. Olver, and A. Vetta. A priority-based model of routing. *Chicago Journal of Theoretical Computer Science*, 2008(1).
- [29] G. Finn and E. Horowitz. A linear time approximation algorithm for multiprocessor scheduling. *BIT*, 19:312–320, 1979.
- [30] L. Fleischer, K. Jain, and M. Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *FOCS*, pages 277–285, 2004.
- [31] L. Fleischer and Z. Svitkina. Preference-constrained oriented matching. In *ANALCO*, pages 56–65, 2010.
- [32] I. Foster and C. Kesselman. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann Publishers, November 1998.
- [33] M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. Computing nash equilibria for scheduling on restricted parallel links. *Theory Comput. Syst.*, 47(2):405–432, 2010.
- [34] G.H. Gonnet, J.I. Munro, and H. Suwanda. Exegesis of self-organizing linear search. *SIAM Journal on Computing*, 10(3):613–637, 1981.
- [35] L.A. Hall, A.S. Schulz, D.B. Shmoys, and J. Wein. Scheduling to minimize average completion time: Off-line and on-line approximation algorithms. *Math. Oper. Res.*, 22(3):513–544, 1997.
- [36] M. Hoefer, V. S. Mirrokni, H. Röglin, and S.-H. Teng. Competitive routing over time. In *WINE*, pages 18–29, 2009.
- [37] R. Hoeksma and M. Uetz. The price of anarchy for minsum related machine scheduling. 2011 (manuscript).
- [38] H. Hoogeveen, P. Schuurman, and G.J. Woeginger. Non-approximability results for scheduling problems with minsum criteria. In *IPCO*, pages 353–366, 1998.
- [39] W.A. Horn. Minimizing average flow time with parallel machines. *Operations Research*, 21(3):846–847, 1973.

- [40] O.H. Ibarra and C.E. Kim. Heuristic algorithms for scheduling independent tasks on nonidentical processors. *J. ACM*, 24(2):280–289, 1977.
- [41] N. Immorlica, L. Li, V.S. Mirrokni, and A.S. Schulz. Coordination mechanisms for selfish scheduling. *Theor. Comput. Sci.*, 410(17):1589–1598, 2009.
- [42] S. Koch and M. Skutella. Nash equilibria and the price of anarchy for flows over time. *Theory of Computing Systems*, to appear.
- [43] Y.A. Korilis, A.A. Lazar, and A. Orda. Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Transactions on Networking*, 5(1):161–173, 1997.
- [44] E. Koutsoupias and C.H. Papadimitriou. Worst-case equilibria. *Computer Science Review*, 3(2):65–69, 2009.
- [45] J.K. Lenstra, A.H.G. Rinnooy Kan, and P. Brucker. Complexity of machine scheduling problems. *Ann. Discrete Math.*, 4:281–300, 1977.
- [46] L.L. Peterson, T.E. Anderson, D.E. Culler, and T. Roscoe. A blueprint for introducing disruptive technology into the internet. *Computer Communication Review*, 33(1):59–64, 2003.
- [47] M. Pinedo. *Scheduling: Theory, Algorithms and Systems*. Springer, third edition, 2008.
- [48] T. Roughgarden. Stackelberg scheduling strategies. In *STOC*, pages 104–113, 2001.
- [49] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC*, pages 513–522, 2009.
- [50] S. Sahni and Y. Cho. Bounds for list schedules on uniform processors. *Siam J. of Computing*, 9:91–103, 1980.
- [51] A.S. Schulz and M. Skutella. Scheduling unrelated machines by randomized rounding. *SIAM J. Discrete Math.*, 15(4):450–469, 2002.
- [52] P. Schuurman and T. Vredeveld. Performance guarantees of local search for multiprocessor scheduling. In *IPCO*, pages 370–382, 2001.
- [53] J. Sethuraman and M.S. Squillante. Optimal scheduling of multiclass parallel machines. In *SODA*, pages 963–964, 1999.
- [54] M. Skutella. Convex quadratic and semidefinite programming relaxations in scheduling. *J. ACM*, 48(2):206–242, 2001.
- [55] M. Skutella and G.J. Woeginger. A ptas for minimizing the total weighted completion time on identical parallel machines. *Math. Oper. Res.*, 25(1):63–75, 2000.
- [56] W. Smith. Various optimizers for single stage production. *Naval Res. Logist. Quart.*, 3(1-2):59–66, 1956.
- [57] H. von Stackelberg. Marktform und Gleichgewicht. *Springer-Verlag*, 1934. English translation entitled *The Theory of the Market Economy*.
- [58] T. Vredeveld. *Combinatorial approximation algorithms. Guaranteed versus experimental performance*. 2002. Ph.D. thesis.

A Lower Bounds

Proof of Theorem 3.3. We begin by presenting the family of game instances that leads to a pure price of anarchy approaching 4 for games induced by **SmithRule** in the restricted identical machines model [23], and then show how to generalize the lowerbound based on this construction.

There are m machines and k groups of jobs g_1, \dots, g_k , where group g_r has m/r^2 jobs. We assume that m is such that all groups have integer size and let j_{rs} denote the s -th job of the r -th group. A job j_{rs} can be assigned to machines $1, \dots, s$, and we assume that for two jobs j_{rs} and $j_{r's'}$ with $s < s'$, then $j_{r's'}$ has higher priority than j_{rs} (if $s = s'$, the ordering can be arbitrary).

If every job j_{rs} is assigned to machine s , there are exactly m/r^2 jobs with completion time r ($1 \leq r \leq k$), which leads to a total cost of $m \sum_{r=1}^k 1/r$. On the other hand, assigning each job to the machine with smallest index among all the ones that minimize its completion time gives a PNE assignment whose total cost is $\Omega(4m \sum_{r=1}^k 1/r)$ [23].

Fixed ordering policies. Let \mathbf{x} be the Nash equilibrium described above, and \mathbf{x}^* the optimal assignment. Note that $x_j \leq x_j^*$ for every job j . Since the job ordering under the optimal assignment does not affect the cost, we only need to make sure that for any two jobs j, j' with $x_{j'}^* < x_j^*$, j gets higher priority than j' on x_j^* .

Given a specific lower bound instance for **SmithRule**, we have n job *slots*, each defined by the pair of machines x_j and x_j^* . Given a set of ordering policies, each machine has its own strictly ordered list of all n jobs. What we need to do is assign a specific job to each slot so that the ordering restrictions as specified in the previous paragraph comply with the lists. We start from the slot j with the greatest index x_j , and we assign the first job of machine x_j 's list to this slot. We then erase this job from all lists and repeat. In case of a tie, that is if there is more than one slot with the same x_j , we first consider the slots with greater x_j^* machine index. This ensures that, given the PNE assignment, any job that deviates back to its OPT machine will suffer cost at least as much as in the **SmithRule** instance, while its cost in the PNE is the same as in the given instance. Therefore, \mathbf{x} is a PNE for this set of ordering policies. □

Proof of Proposition 4.7. The construction is a slight variant of one given in Caragiannis et al. [12] for load balancing games. We define the construction in terms of the *game graph*; a directed graph, with nodes corresponding to machines, and arcs corresponding to jobs. The interpretation of an arc (i^*, i) is that the corresponding machine is run on i at the Nash equilibrium, and i^* in the optimal solution (all jobs can only be run on at most two machines in the instance we construct).

Our graph consists of a binary tree of depth ℓ , with a path of length ℓ appended to each leaf of the tree. In addition, there is a loop at the endpoint of each path. All arcs are directed towards the root; the root is considered to be at depth zero. In the binary tree, on a machine at depth i , the processing time of any job that can run on that machine is $(3/2)^{\ell-i}$. In the chain, on a machine at distance k from the tree leaves all processing times are $(1/2)^k$.

By slightly perturbing the processing times of jobs on different machines, it is easily checked that if every job is run on the machine pointed to by its corresponding arc, the assignment is a pure NE. The latter holds for arbitrary prompt strongly local coordination mechanisms so long as jobs are anonymous. On the other hand, if all jobs choose their alternative strategy, we obtain the optimal solution. A straightforward calculation shows that, in the limit $\ell \rightarrow \infty$, the ratio of the cost of the NE to the optimal cost converges to $13/6 > 2.166$. □

Rand. The previous instance can be easily modified to give a lower bound on the performance of **Rand**. Simply take the same instance but replace $3/2$ by $4/3$ and $1/2$ by $2/3$. The same assignment

then gives a PNE, and in this case the ratio of interest approaches $5/3$.

B The performance of Rand on a single machine

Proof of Theorem 4.12. We want to prove that for any sequence u_1, \dots, u_k , $u_i \geq 0$, the following inequality holds:

$$\sum_i \sum_j u_i u_j \frac{ij}{i+j} \leq \frac{\pi}{4} \sum_i \sum_j u_i u_j \min\{i, j\}.$$

We will in fact prove that for any sequence x_1, x_2, \dots, x_n , $x_i \in \mathbb{N}$,

$$\sum_i \sum_j \frac{x_i x_j}{x_i + x_j} < \frac{\pi}{4} \sum_i \sum_j \min\{x_i, x_j\}. \quad (15)$$

This implies the inequality in the statement, for the choice $u_r = |\{i : x_i = r\}|$, and hence clearly for any integer sequence (u_i) . An obvious scaling argument then gives it for general nonnegative u_i .

Since both summations in (15) are symmetric, we may assume without loss of generality that $x_1 \geq \dots \geq x_n \geq 0$. Then, we note that $\sum_{i=1}^n \sum_{j=1}^n \min\{x_i, x_j\} = 2 \sum_{i=1}^n x_i(i-1/2)$. Also, observe that the inequality is homogeneous so that proving the inequality is equivalent to proving that the optimal value of the following concave optimization problem is less than $\pi/2$:

$$z = \max \left\{ \sum_{i=1}^n \sum_{j=1}^n \frac{x_i x_j}{x_i + x_j} : \text{s.t. } \sum_{i=1}^n x_i(i-1/2) = 1, x_1 \geq \dots \geq x_n \geq 0 \right\}.$$

Clearly $z \leq z'$, where

$$z' = \max \left\{ \sum_{i=1}^n \sum_{j=1}^n \frac{x_i x_j}{x_i + x_j} : \text{s.t. } \sum_{i=1}^n x_i(i-1/2) = 1, x_i \geq 0 \text{ for all } i = 1, \dots, n \right\}.$$

Furthermore, we may assume that in an optimal solution all variables satisfy $x_i > 0$. Otherwise, we could consider the problem in a smaller dimension. Thus, the KKT optimality conditions state that for all $i = 1, \dots, n$ we have

$$\mu(i-1/2) = 2 \sum_{j=1}^n \left(\frac{x_j}{x_i + x_j} \right)^2. \quad (16)$$

Multiplying by x_i , summing over all i , and using $\sum_{i=1}^n x_i(i-1/2) = 1$, we obtain:

$$\mu = 2 \sum_{i=1}^n \sum_{j=1}^n x_i \left(\frac{x_j}{x_i + x_j} \right)^2 = \sum_{i=1}^n \sum_{j=1}^n \frac{x_i x_j}{(x_i + x_j)^2} (x_i + x_j) = z'.$$

Now consider (16) with $i^* = \arg \max_i x_i(i-1/2)^2$. We have that

$$z' = \frac{2}{i^* - 1/2} \sum_{j=1}^n \left(\frac{x_j}{x_{i^*} + x_j} \right)^2 \leq 2(i^* - 1/2)^3 \sum_{j=1}^{\infty} \left(\frac{1}{(i^* - 1/2)^2 + (j - 1/2)^2} \right)^2.$$

Using standard complex analysis it can be shown that the latter summation equals

$$(\pi/2)((i^* - 1/2)\pi \tanh(\pi(i^* - 1/2))^2 + \tanh(\pi(i^* - 1/2)) - \pi(i^* - 1/2)),$$

which is less than $\pi/2$. □