

Approximate Strong Equilibrium in Job Scheduling Games

Michal Feldman* Tami Tamir†

February 17, 2008

Abstract

A Nash Equilibrium (NE) is a strategy profile that is resilient to unilateral deviations, and is predominantly used in analysis of competitive games. A downside of NE is that it is not necessarily stable against deviations by coalitions. Yet, as we show in this paper, in some cases, NE does exhibit stability against coalitional deviations, in that the benefits from a joint deviation are bounded. In this sense, NE approximates *strong equilibrium* (SE) [6].

We provide a framework for quantifying the stability and the performance of various assignment policies and solution concept in the face of coalitional deviations. Within this framework we evaluate a given configuration according to three measurements: (i) IR_{min} : the maximal number α , such that there exists a coalition in which the minimum improvement ratio among the coalition members is α (ii) IR_{max} : the maximum improvement ratio among the coalition's members. (iii) DR_{max} : the maximum possible damage ratio of an agent outside the coalition.

This framework can be used to study the proximity between different solution concepts, as well as to study the existence of approximate SE in settings that do not possess any such equilibrium. We analyze these measurements in job scheduling games on identical machines. In particular, we provide upper and lower bounds for the above three measurements for both NE and the well-known assignment rule *Longest Processing Time* (LPT) (which is known to yield a NE). Most of our bounds are tight for any number of machines, while some are tight only for three machines. We show that both NE and LPT configurations yield small constant bounds for IR_{min} and DR_{max} . As for IR_{max} , it can be arbitrarily large for NE configurations, while a small bound is guaranteed for LPT configurations. For all three measurements, LPT performs strictly better than NE.

With respect to computational complexity aspects, we show that given a NE on $m \geq 3$ identical machines and a coalition, it is NP-hard to determine whether the coalition can deviate such that every member decreases its cost. For the unrelated machines settings, the above hardness result holds already for $m \geq 2$ machines.

*School of Business Administration and Center for the Study of Rationality, Hebrew University of Jerusalem. *E-mail*: mfeldman@cs.huji.ac.il. Research partially supported by a grant of the Israel Science Foundation, BSF, Lady Davis Fellowship, and an IBM faculty award.

†School of Computer Science, The Interdisciplinary Center, Herzliya, Israel. *E-mail*: tami@idc.ac.il.

1 Introduction

We consider job scheduling problems, in which n jobs are assigned to m identical machines and incur a cost which is equal to the total load on the machine they are assigned to¹. These problems have been widely studied in recent years from a game theoretic perspective [20, 3, 10, 11, 14]. In contrast to the traditional setting, where a central designer determines the allocation of jobs into machines and all the participating entities are assumed to obey the protocol, in distributed settings, the situation may be different. Different machines and jobs may be owned by different *strategic* entities, who will typically attempt to optimize their own objective rather than the global objective. Game theoretic analysis provides us with the mathematical tools to study such situations, and indeed has been extensively used in recent years by computer scientists. This trend is motivated in part by the emergence of the Internet, which is composed of distributed computer networks managed by multiple administrative authorities and shared by users with competing interests [23].

Most game theoretic models applied to job scheduling problems, as well as other network games (e.g., [13, 2, 24, 4]), use the solution concept of *Nash equilibrium* (NE), in which the strategy of each agent is a best response to the strategies of all other agents. While NE is a powerful tool for predicting outcomes in competitive environments, its notion of stability applies only to unilateral deviations. However, even when no single agent can profit by a unilateral deviation, NE might still not be stable against a group of agents *coordinating* a joint deviation, which is profitable to *all the members* of the group. This stronger notion of stability is exemplified in the *strong equilibrium* (SE) solution concept, coined by Aumann (1959). In a strong equilibrium, no coalition can deviate and improve the utility of *every* member of the coalition.




Figure 1: An example of a configuration (a) that is a Nash equilibrium but is not resilient against coordinated deviations, since the jobs of load $\{5, 5, 2, 2\}$ all profit from the deviation demonstrated in (b).

As an example, consider the configuration depicted in Figure 1(a). It is a NE since no job can reduce its cost through a unilateral deviation (recall that the cost of each job is defined to be the load on the machine it is assigned to, as assumed in many job scheduling models). One may think that a NE on identical machines is also sustainable against joint deviations.

¹This cost function characterizes systems in which jobs are processed in parallel, or when all jobs on a particular machine have the same single pick-up time, or need to share some resource simultaneously.

Yet, as was already observed in [3], this may not be true². For example, the configuration above is not resilient against a coordinated deviation of the coalition $\Gamma = \{5, 5, 2, 2\}$ deviating to configuration (b), where the jobs of load 5 decrease their costs from 10 to 8, and the jobs of load 2 improve from 5 to 4. Note that the cost of the two jobs of load 3 (which are not members of the coalition) increases.

In the example above, every member of the coalition improves its cost by a (multiplicative) factor of $\frac{5}{4}$. By how much more can a coalition improve? Is there a bound on the *improvement ratio*? As it will turn out, this example is in fact the most extreme one in a sense that will be clarified below. Thus, while NE is not completely stable against coordinated deviations, in some settings, it does provide us with some notion of approximate stability to coalitional deviations (or *approximate strong equilibrium*).

In this paper we provide a framework for studying the notion of approximate stability to coalitional deviations. In our analysis, we consider three different measurements. The first two measure the stability of a configuration, and the third measures the worst possible effect on the non-deviating jobs.

1. Minimum Improvement Ratio: This notion is discussed in Section 3, and refers to configurations from which no coalition of agents can deviate such that *every* member of the coalition improves by a large factor³. Formally, the improvement ratio of a job in the coalition is the ratio between its pre- and post-deviation cost. We say that a configuration s forms an α -SE if there is no coalition in which each agent can improve by a factor of more than α . This notion was also studied by [1] in the context of SE existence. There, the author showed that for a sufficiently large α , an α -SE always exists. The justification behind this concept is that agents may be willing to deviate only if they improve by a sufficiently high factor (due to, for example, some overhead associated with the migration).

For three machines, we show that every NE is a $\frac{5}{4}$ -SE. That is, there is no coalition that can deviate such that every member improves by a factor larger than $\frac{5}{4}$. For this case, we also provide a matching lower bound (recall Figure 1 above), that holds for any $m \geq 3$. For arbitrary m , we show that every NE is a $(2 - \frac{2}{m+1})$ -SE. Our proof technique draws a connection between makespan approximation⁴ and approximate stability.

We also consider a subclass of NE, produced by the *Longest Processing Time* (LPT) rule [17]. The LPT rule sorts the jobs in a non-increasing order of their loads and greedily assigns each job to the least loaded machine. It is easy to verify that every configuration produced by LPT is a NE [16]. Is it also a SE? Note that for the instance depicted in Figure 1, LPT would

²This statement holds for $m \geq 3$. For 2 identical machines, every NE is also a SE [3].

³Throughout this paper, we define approximation by a *multiplicative* factor. Since the improvement and damage ratios for all the three measurements presented below are constants greater than one (as will be shown below), the *additive* ratios are unbounded. Formally, for any value a it is possible to construct instances (by scaling the instances we provide for the multiplicative ratio) in which the cost of all jobs is reduced, or the cost of some jobs is increased, by at least an additive factor of a .

⁴makespan is defined as the maximum load on any machine in the configuration.

have produced a SE. However, as we show, this is not always the case. Yet, for $m = 3$, every LPT-based configuration is a $\frac{2}{\sqrt{34}-4}$ -SE (≈ 1.092), and we also provide a matching lower bound, that holds for any $m \geq 3$. For arbitrary m , we show an upper bound of $\frac{4}{3} - \frac{1}{3m}$. These results indicate that LPT is more stable than NE with respect to coalitional deviations.

2. Maximum Improvement Ratio: In Section 4 we study an alternative notion of approximate stability, in which there is no coalition such that *some* agent improves by a factor of more than α . This notion is similar in spirit to stability against a large *total* improvement. Interestingly, we find out that given a NE configuration, the improvement ratio of a single agent may not be bounded, for any $m \geq 3$. In contrast, for LPT-based configurations on three machines, no agent can improve by a factor of $\frac{5}{3}$ or more and this bound is tight. Thus, with respect to maximum IR, the relative stability of LPT compared to NE is significant. For arbitrary m , we provide a lower bound of $2 - \frac{1}{m}$, which we believe to be tight.

3. Maximum Damage Ratio: As is the case for the jobs of load 3 in Figure 1, some jobs might be hurt from a coalitional deviation. The third measurement that we consider is the worst possible effect of a deviation on these naive jobs. Formally, the *maximum damage ratio* is the maximal ratio between the pre- and post-deviation cost of a job. Note that it does not measure the stability of a configuration – we assume that an agent’s motivation to deviate is not influenced by the potential damage it will cause others. However, this measurement is important since it guarantees a bound on the maximal damage that any agent can experience. In Section 5, we prove that the maximum damage ratio is less than 2 for any NE configuration, and less than $\frac{3}{2}$ for any LPT-based configuration. Both bounds hold for any $m \geq 3$, and for both we provide matching lower bounds. Note that the minimum damage ratio is of no practical interest.

In summary, our results in Sections 3-5 (see Table 1) indicate that NE-based configurations are approximately stable with respect to the IR_{min} measurement. Moreover, the performance of jobs outside the coalition would not be hurt by much as a result of a coalitional deviation. It would be interesting to study in what families of games NE are guaranteed to provide approximate SE. As for IR_{max} , our results provide an additional benefit of the LPT rule, which is already known to possess attractive properties (with respect to, e.g., makespan approximation and stability against unilateral deviations).

	IR_{min}			IR_{max}		DR_{max}	
	upper bound		lower bound	upper bound	lower bound	upper bound	lower bound
	$m = 3$	$m \geq 3$					
NE	$\frac{5}{4}$	$2 - \frac{2}{m+1}$	$\frac{5}{4}$	unbounded		2	2
LPT	$\frac{2}{\sqrt{34}-4}$	$\frac{4}{3} - \frac{1}{3m}$	$\frac{2}{\sqrt{34}-4}$	$\frac{5}{3}$ ($m=3$)	$2 - \frac{1}{m}$	$\frac{3}{2}$	$\frac{3}{2}$

Table 1: Our results for the three measurements. Unless specified otherwise, the results hold for arbitrary m .

In Section 6, we study computational complexity aspects of coalitional deviations. We find

that it is NP-hard to determine whether a NE configuration on $m \geq 3$ identical machines is a SE. Moreover, given a particular configuration and a set of jobs, it is NP-hard to determine whether this set of jobs can engage in a coalitional deviation. For unrelated machines (i.e., where each job incurs a different load on each machine), the above hardness results hold already for $m = 2$ machines. These results might have implications on coalitional deviations with computationally restricted agents.

Related work: NE is shown in this paper to provide approximate stability against coalitional deviations. A related body of work studies how well NE approximates the optimal outcome of competitive games. The Price of Anarchy was defined in [23, 20] as the ratio between the worst-case NE and the optimum solution, and has been extensively studied in various settings, including job scheduling [20, 10, 11], network design [2, 4, 5, 13], network routing [24, 7, 9], and more.

The notion of strong equilibrium (SE) [6] expresses stability against coordinated deviations. The downside of SE is that most games do not admit any SE, in contrast to NE which always exists (in mixed strategies). Various recent works have studied the existence of SE in particular families of games. [3] showed that in every job scheduling game and (almost) every network creation game, a SE exists. In addition, [12, 18, 19, 25] provided a topological characterization for the existence of SE in different congestion games, including routing and cost-sharing connection games. The vast literature on SE [18, 19, 22, 8] concentrate on pure strategies and pure deviations, as is the case in our paper. In job scheduling settings, [3] showed that if mixed deviations are allowed, it is often the case that no SE exists. When a SE exists, aside from its robustness, it has other appealing properties. For example, in many cases, the price of anarchy with respect to SE (denoted the *strong price of anarchy* in [3]) is significantly better than the price of anarchy with respect to NE [3, 14, 21].

2 Model and Preliminaries

In our job scheduling setting there is a set of m identical machines, $M = \{M_1, \dots, M_m\}$, and n jobs, $N = \{1, \dots, n\}$, where job j has load p_j , and is controlled by a single agent (in the remainder of the paper, we use agents and jobs interchangeably). A schedule $s \in S : N \rightarrow M$ (also denoted a configuration) is an assignment of jobs into machines. The load of a machine M_i in a configuration $s \in S$, denoted $C_i(s)$, is the sum of the loads of the jobs assigned to M_i , that is $C_i(s) = \sum_{\{j|s(j)=M_i\}} p_j$. In our model, the individual cost of player $j \in N$, denoted $c_j(s)$, is the total load on the machine job j is assigned to, i.e., $c_j(s) = C_i(s)$, where $s(j) = M_i$. Note that the internal order of the jobs on a particular machine does not affect the jobs' individual costs.

A configuration $s \in S$ is a pure **Nash Equilibrium** if no player $j \in N$ can benefit from unilaterally migrating to another machine. A configuration $s \in S$ is a pure **Strong Equilibrium** if no coalition $\Gamma \subseteq N$ can form a coordinated deviation in a way that *every* member of the coalition reduces its cost.

Recall that $C_i(s)$ denotes the load on machine i in configuration s . Let s' denote the post-deviation configuration. Then, $C_i(s')$ denotes the load on machine i after the deviation. When clear in the context, we abuse notation and denote the load on machine i before and after the deviation by C_i and C'_i , respectively. In addition, we let P_{i_1, i_2} be a binary indicator whose value is 1 if some job in the coalition migrates from M_{i_1} to M_{i_2} , and 0 otherwise. Since jobs in the coalition improve their cost by definition, $P_{i_1, i_2} = 1$ implies that $C'_{i_2} < C_{i_1}$. The *improvement ratio* of a job $j \in \Gamma$, migrating from machine M_{i_1} (with initial load C_{i_1}) to machine M_{i_2} (with post-deviation load C'_{i_2}), is $IR(j) = C_{i_1}/C'_{i_2}$. Clearly, for any job j in the coalition, $IR(j) > 1$. The *damage ratio* of a job $j \notin \Gamma$, assigned on machine M_i is $DR(j) = C'_i/C_i$. Clearly, for any job j not in the coalition, $IR(j) \leq 1$ (else j is part of the coalition). Finally, we refer to coalitions deviating from NE or LPT-based configurations as *NE-based* and *LPT-based coalitions*, respectively.

Definition 2.1 *A configuration s is an α -strong equilibrium (α -SE) if for any deviation and any coalition Γ , it holds that $\min_{j \in \Gamma} IR(j) \leq \alpha$. We also say that for any Γ , $IR_{\min}(s, \Gamma) \leq \alpha$.*

For the maximum improvement ratio, we say that $IR_{\max}(s, \Gamma) \leq \alpha$ if for any deviation of a coalition Γ , it holds that $\max_{j \in \Gamma} IR(j) \leq \alpha$.

For the maximum damage ratio, we say that $DR_{\max}(s, \Gamma) \leq \alpha$ if for any deviation of a coalition Γ , it holds that $\max_{j \notin \Gamma} DR(j) \leq \alpha$.

We next provide several useful observations and claims that prove useful in our analysis below.

Observation 2.2 *At least one job leaves any machine participating in an NE-based coalition.*

Proof: Suppose that there exists a machine to which a job migrates but no job leaves. Then, the job that migrates to it would also migrate alone, contradicting the original schedule is a NE. \square

Definition 2.3 *Assume w.l.o.g that M_1 is the most loaded machine in a given configuration. We say that a coalition obeys the flower structure if for all $i > 1$, $P_{1,i} = P_{i,1} = 1$ and for all $i, j > 1$, $P_{i,j} = 0$ (See Figure 2).*




Figure 2: A graph representation of a coalition on 5 machines obeying the flower structure. There is an edge from M_i to M_j if and only if $P_{i,j} = 1$.

In particular, for $m = 3$, a coalition obeys the flower structure if $P_{1,2} = P_{2,1} = P_{1,3} = P_{3,1} = 1$ and $P_{2,3} = P_{3,2} = 0$.

Claim 2.4 *Any NE-based coalition on three machines obeys the flower structure.*

Proof: Let M_1 be the most loaded machine. We first show that $P_{2,3} = P_{3,2} = 0$. Assume first

that both $P_{2,3} = P_{3,2} = 1$. Thus, $C'_2 < C_3$ and $C'_3 < C_2$. Clearly, since jobs only migrate among the machines, $\sum_i C_i = \sum_i C'_i$. Therefore, it must be that $C'_1 > C_1$. However, in any action of a coalition the load on the most loaded machine does not increase. A contradiction. Therefore, at most one of $P_{2,3}, P_{3,2}$ can be 1. Assume w.l.o.g that $P_{2,3} = 1$. By Observation 2.2 some job leaves M_3 , and by the above it cannot be to M_2 . Thus, it must be that $P_{3,1} = 1$. Similarly, some job leaves M_1 . If $P_{1,2} = 1$, then we get that $C'_1 < C_3$, $C'_2 < C_1$, and $C'_3 < C_2$, contradicting $\sum_i C_i = \sum_i C'_i$. If $P_{1,3} = 1$ then we get that $C'_1 < C_3$, $C'_2 < C_2$ (no job is added to M_2), and $C'_3 < C_1$, contradicting $\sum_i C_i = \sum_i C'_i$ again. Thus, $P_{2,3} = 0$. The proof of $P_{3,2} = 0$ is symmetric.

In order to get the flower structure we need to show $P_{1,2} = P_{1,3} = P_{2,1} = P_{3,1} = 1$. We know that all three machines participate in the coalition action. By the above $P_{2,3} = P_{3,2} = 0$. By Claim 2.2 some job leaves each of M_2, M_3 , therefore, $P_{2,1} = P_{3,1} = 1$. Also, some job leaves M_1 , thus at least one of $P_{1,2}, P_{1,3}$ equals 1. Assume w.l.o.g that $P_{1,2} = 1$. We show that also $P_{1,3} = 1$. In particular, we show that $C'_3 > C_3$, and since $P_{2,3} = 0$ it must be that $P_{1,3} = 1$. Assume the opposite, that is $C'_3 \leq C_3$. We already know that $P_{1,2} = P_{2,1} = 1$. Thus, $C'_2 < C_1$, $C'_1 < C_2$, and by our assumption $C'_3 \leq C_3$. That is, $\sum_i C'_i < \sum_i C_i$. A contradiction. \square

It is known [3] that any NE-schedule on two identical machines is also a SE. By the above claim, at least four jobs participate in any coalition on three machines. Clearly, at least four jobs participate in any coalition on $m > 3$ machines. Therefore,

Corollary 2.5 *For every NE-based coalition Γ , it holds that $|\Gamma| \geq 4$.*

3 α -Strong Equilibrium

In this section, the stability of configurations is measured by $\min_{j \in \Gamma} IR(j)$. We first provide a complete analysis (i.e. matching upper and lower bounds) for $m = 3$ for both NE and LPT. For arbitrary m , we provide an upper bound for NE and LPT, and show that the lower bounds for $m = 3$ hold for any m .

Theorem 3.1 *Any NE schedule on three machines is a $\frac{5}{4}$ -SE.*

Proof: Given a NE configuration s on 3 machines, let Γ be a coalition, and $r = IR_{\min}(s, \Gamma)$. By Claim 2.4, Γ obeys the flower structure. Therefore: $C'_1 \leq C_2/r$; $C'_1 \leq C_3/r$; $C'_2 \leq C_1/r$; and $C'_3 \leq C_1/r$. Let $P = \sum_j p_j$ (also $= C_1 + C_2 + C_3$).

Summing up the above inequalities we get $r \leq (C_1 + P)/(C'_1 + P)$.

Claim 3.2 $C_1 \leq P/2$.

Proof: Let g be the larger between $C_1 - C_2$ and $C_1 - C_3$. By the flower structure, there are at least two jobs on M_1 , thus $g \leq C_1/2$ - since otherwise some job would benefit from leaving M_1 , contradicting the NE. By definition of g , we know that $2C_1 \leq C_2 + C_3 + 2g$, and since $2g \leq C_1$, we get that $C_1 \leq P/2$. \square

Distinguish between two cases:

1. $C'_1 \geq P/5$: in this case $r \leq (C_1 + P)/(C'_1 + P) \leq (3P/2)/(6P/5) = 5/4$.
2. $C'_1 < P/5$: It means that $C'_2 + C'_3 > 4P/5$ (M_2 and M_3 have the rest of the load), that is, at least one of $C'_2, C'_3 > 2P/5$. W.l.o.g. let it be M_2 . By the flower structure some job from M_1 migrates M_2 . This job's improvement ratio is C_1/C'_2 , which, by Claim 3.2, is less than $(P/2)/(2P/5) = 5/4$. Thus, again, $r < 5/4$.

□

The above analysis is tight as shown in Figure 1. Moreover, this lower bound can be extended to any $m > 3$ by adding $m - 3$ machines and $m - 3$ heavy jobs assigned to these machines. Thus,

Theorem 3.3 *For $m \geq 3$, there exists a NE schedule s and a coalition Γ s.t. $IR_{min}(s, \Gamma) = \frac{5}{4}$.*

For LPT-based configurations, the bound on the minimum improvement ratio is lower:

Theorem 3.4 *Any LPT-based schedule on three machines is a $(\frac{2}{\sqrt{34}-4} \approx 1.0924)$ -SE.*

Proof: Let M_1 be the most loaded machine in the schedule. Let x be the lightest (also last) job assigned to M_1 (we also denote its load by x). Let ℓ denote the load on M_1 before x is assigned to it, that is $\ell = C_1 - x$. For a give LPT-based schedule s and a coalition Γ let $r = IR_{min}(s, \Gamma)$.

By Claim 2.4, Γ obeys the flower structure. Therefore: $C'_1 \leq C_2/r$; $C'_1 \leq C_3/r$; $C'_2 \leq C_1/r$; and $C'_3 \leq C_1/r$. Let $P = \sum_j p_j$ (also = $C_1 + C_2 + C_3$). Summing up the above inequalities we get

$$r \leq \frac{C_1 + P}{C'_1 + P} \leq \frac{2\ell + 2x + C_2 + C_3}{C'_1 + \ell + x + C_2 + C_3} \leq \frac{4\ell + 2x}{C'_1 + 3\ell + x}. \quad (1)$$

The last inequality is due to the fact that $C_2 + C_3 \geq 2\ell$ (by the LPT rule, else x would not have been assigned to M_1), and since the middle term is decreasing with $C_2 + C_3$ (since $C'_1 < \ell + x$). We know that $C'_2 + C'_3 = P - C'_1$ (M_2 and M_3 have the rest of the load), thus, at least one of C'_2, C'_3 has load at least $(P - C'_1)/2$. W.l.o.g. let it be M_2 . Summing up $C'_1 \leq C_2/r$; $C'_1 \leq C_3/r$ we get $C'_1 \leq \frac{C_2 + C_3}{2r}$. Also, as explained above $C_2 + C_3 \geq 2\ell$. Therefore, since some job migrates from M_1 to M_2 we can bound r as follows.

$$r \leq \frac{C_1}{C'_2} \leq \frac{\ell + x}{(P - C'_1)/2} = \frac{2\ell + 2x}{P - C'_1} \leq \frac{2\ell + 2x}{C_1 + C_2 + C_3 - \frac{C_2 + C_3}{2r}} \leq \frac{2\ell + 2x}{\ell + x + 2\ell(1 - \frac{1}{2r})}.$$

Therefore, $\ell r + x r + 2\ell r - \ell \leq 2\ell + 2x$. Implying $r \leq \frac{3\ell + 2x}{3\ell + x}$. Note that this bound for r is decreasing with ℓ and independent of C'_1 , while the bound 1 for r is increasing with ℓ and decreasing with C'_1 . Therefore, the maximal possible value for r is achieved when $\frac{4\ell + 2x}{C'_1 + 3\ell + x} = \frac{3\ell + 2x}{3\ell + x}$. That is, $3\ell^2 + \ell x - 3\ell C'_1 - 2x C'_1 = 0$, or $C'_1 = \frac{3\ell^2 + \ell x}{3\ell + 2x}$. Distinguish between two cases:

1. $C'_1 \geq 3x$: In this case we get $6x^2 + 8\ell x - 3\ell^2 \leq 0$ implying $x \leq \frac{-8\ell + \sqrt{64\ell^2 + 72\ell^2}}{12}$. That is, $x \leq \frac{\ell}{6}(\sqrt{34} - 4)$. Therefore, $r \leq \frac{3\ell + 2x}{3\ell + x} \leq \frac{2}{\sqrt{34} - 4}$. Note that r is increasing with x .
2. $C'_1 < 3x$: In this case we get $x > \frac{\ell}{6}(\sqrt{34} - 4)$. Therefore, $r \leq \frac{2\ell + 2x}{3\ell + x - y} \leq \frac{2\ell + 2x}{3\ell - 2x} \leq \frac{2}{\sqrt{34} - 4}$. Note that r is decreasing with x .

□




Figure 3: An LPT-based coalition on 3 machines in which all migrating jobs improve by $\frac{2}{\sqrt{34}-4}$.

The above analysis is tight as shown in Figure 3. Moreover, as for NE, this lower bound can be extended to any $m > 3$ by adding dummy jobs and machines. Thus,

Theorem 3.5 *For any $m \geq 3$, there exists an LPT schedule s and a coalition Γ s.t. $IR_{min}(s, \Gamma) = \frac{2}{\sqrt{34}-4}$.*

We next provide upper bounds for arbitrary m .

Definition 3.6 *Let $I = \langle N, M \rangle$ be an instance of job scheduling with machines M and jobs N . For a given $M' \subseteq M$, let $N' \subseteq N$ be the set of jobs scheduled on M' , and consider the instance $I' = \langle N', M' \rangle$. An assignment method is said to be subset-preserving if for any I , and M' , it produces the same schedule on M' with input I and with input I' .*

Claim 3.7 *LPT is a subset-preserving method.*

Proof: The proof is by induction on the order of the jobs in I' . Note that since N' is a sublist of N , the jobs in N' are in the same order as in N . The first job is scheduled on the first empty machine. For any other job $j \in N'$, by the induction hypothesis, when j is scheduled, the load on each of the machines is identical to the load of the corresponding machines at the time j was scheduled as a member of N . This load is generated only by jobs in N' that come before j in N . Therefore, by LPT, j is scheduled on the least loaded machine among the machines M' . □

Also, note that NE has the subset-preserving property, that is, if a schedule on M machines is NE, then the sub-schedule on any subset of $M' \subseteq M$ of the machines is also NE.

Lemma 3.8 *Let A be an assignment method that is (i) subset-preserving, (ii) yields Nash equilibrium, and (iii) approximates the minimum makespan within a factor of r , where r is non-decreasing in m . Then, A produces an r -SE.*

Proof: Assume towards contradiction that there exists an instance I for A on m machines, such that in the schedule of I produced by A , there exists a coalition in which the improvement ratio of every member is greater than r . Let Γ be such a coalition of minimum size. For every machine from which a job J_j migrates, there must exist a job migrating to it, otherwise, $\Gamma \setminus \{J_j\}$ is also a coalition having $IR_{min} > r$, in contradiction to the minimality of Γ . Let M' denote the set of machines that are part of the coalition, let $N' \subseteq N$ be the set of jobs assigned to M' by A , and let $m' = |M'|$. Consider the instance $I' = \langle N', M' \rangle$. Since A is subset-preserving, the

jobs of N' are scheduled by A on M' in I' exactly as their schedule on M' when scheduled as part of I . In particular, when I' is scheduled by A , the coalition Γ exists, and all the machines M' take part in it. Moreover, each of the jobs improves by a factor of more than r . In other words, for any pair of machines i, j , such that $P_{i,j} = 1$, we have $\frac{C_i}{C'_j} > r(m) \geq r(m')$, where $r(m)$ is the approximation ratio of A on m machines. On the other hand, since A produces an $r(m')$ -approximation, for any machine i , $C_i \leq r(m')OPT(I')$, where $OPT(I')$ is the minimum possible makespan of I' on M' machines. Therefore, if $P_{i,j} = 1$ then $r(m') < \frac{C_i}{C'_j} \leq \frac{r(m')OPT(I')}{C'_j}$. In other words, for any machine j that receives at least one job, $C'_j < OPT(I')$.

However, since at least one job has migrated to each of the m' participating machines, after the deviation the machines M' are assigned all the jobs of N' and they all have load less than $OPT(I')$. A contradiction. \square

The following two results are direct corollaries of Lemma 3.8, Claim 3.7, the observation that NE has the subset preserving property, and the fact that LPT and NE provide the respective approximation ratios of $\frac{4}{3} - \frac{1}{3m}$ [17] and $2 - \frac{2}{m+1}$ [15, 26] to the minimum makespan. These bounds are not tight, but the gap between the lower and upper bounds is only a small constant. We believe that the tight upper bound provided for $m = 3$ is a tight bound for arbitrary m for both NE and LPT.

Corollary 3.9 *Any schedule produced by LPT on m identical machines is a $(\frac{4}{3} - \frac{1}{3m}) - SE$.*

Corollary 3.10 *Any NE schedule on m identical machines is a $(2 - \frac{2}{m+1}) - SE$.*

4 Maximum Improvement Ratio

In this section, the stability of a configuration is measured by $\max_{j \in \Gamma} IR(j)$. We provide a complete analysis for NE configurations and any $m \geq 3$, and for LPT configurations on three machines. The lower bound for LPT on three machines can be extended to arbitrary m . Our results show a significant difference between NE in general and LPT. While the improvement ratio of NE-based coalition can be arbitrarily high, for LPT-based coalition, the highest possible improvement ratio of any participating job is less than $\frac{5}{3}$.

Theorem 4.1 *For any $m \geq 3$ machines, the maximum improvement ratio of a NE-based coalition on m machines is not bounded.*

Proof: Given r , consider the NE-schedule on 3 machines given in 4(a). The coalition consists of $\{1, 1, 2r, 2r\}$. Their improved schedule is given in Figure 4(b). The improvement ratio of the jobs of load 1 is $2r/2 = r$. For $m > 3$, dummy machines and jobs can be added. \square

In contrast to NE-based deviations, for LPT-based deviations we are able to bound the maximum improvement ratio by a small constant:

Theorem 4.2 *For any LPT schedule on three machines, the maximum improvement ratio of any coalition is less than $\frac{5}{3}$.*




Figure 4: An NE-based coalition in which the jobs of load 1 have improvement ratio r .

Proof: Assume, w.l.o.g, that M_1 is the most loaded machine. Let x be the lightest (also last) job assigned to M_1 (we also denote its load by x). Denote by S_i the set (and also the total load) of jobs that remain on M_i and do not participate in the coalition. Denote by $p_{i,j}$ the set (and also total load) of jobs migrating from M_i to M_j . For $i = 1$, we consider x as a different set, that is, the sets $S_1, p_{1,2}, p_{1,3}$ do not include x .

Claim 4.3 $C'_2 > C_2$ and $C'_3 > C_3$.

Proof: To show that $C'_2 > C_2$, assume the opposite, that is $C'_2 \leq C_2$. By Claim 2.4, $P_{1,3} = P_{3,1} = 1$. Thus, $C'_3 < C_1$, $C'_1 < C_3$, and by our assumption $C'_2 \leq C_2$. That is, $\sum_i C'_i < \sum_i C_i$. A contradiction. The proof of $C'_3 > C_3$ is symmetric. \square

Claim 4.4 $C'_1 < \min(C'_2, C'_3)$.

Proof: By Claim 2.4, $P_{1,2} = 1$, and thus $C'_1 < C_2$. By the above claim $C_2 < C'_2$. Thus, $C'_1 < C'_2$. The proof of $C'_1 < C'_3$ is symmetric. \square

The job x is assigned to M_1 by LPT, meaning that the load on M_2 and M_3 is at least $S_1 + p_{1,2} + p_{1,3}$ at that time. Since the load on M_2, M_3 could only increase after the time x is assigned, we get that

$$S_1 + p_{1,2} + p_{1,3} \leq S_2 + p_{2,1} \quad \text{and} \quad S_1 + p_{1,2} + p_{1,3} \leq S_3 + p_{3,1}. \quad (2)$$

Therefore (sum up the two):

$$2(S_1 + p_{1,2} + p_{1,3}) \leq S_2 + S_3 + p_{2,1} + p_{3,1}. \quad (3)$$

Distinguish between two cases:

(i) x remains on M_1 . In This case, $C_1 = S_1 + p_{1,2} + p_{1,3} + x$; $C_2 = S_2 + p_{2,1}$; $C_3 = S_3 + p_{3,1}$, while after the coalition is active $C'_1 = S_1 + p_{2,1} + p_{3,1} + x$; $C'_2 = S_2 + p_{1,2}$; $C'_3 = S_3 + p_{1,3}$.

Since the jobs in $p_{1,2}$ and $p_{1,3}$ are part of the coalition, $C'_2 + C'_3 < 2C_1$. Deducing $p_{1,2}$ and $p_{1,3}$ from both sides we get $S_2 + S_3 < p_{1,2} + p_{1,3} + 2S_1 + 2x$. Combining with Equation 3, we get:

$$p_{1,2} + p_{1,3} < p_{2,1} + p_{3,1} + 2x. \quad (4)$$

Claim 4.5 $x \leq \min(p_{2,1}, p_{3,1})$

Proof: We show $x \leq p_{3,1}$, the proof of $x \leq p_{2,1}$ is symmetric. Assume $x > p_{3,1}$, it means that when x is assigned to M_1 , the load on M_3 is composed of jobs that are a subset of S_3 only. Therefore, by the LPT rule, $S_3 \geq p_{1,2} + p_{1,3} + S_1$. Also, given that the jobs of $p_{1,3}$ are part of the coalition, we know that $S_3 < p_{1,2} + S_1 + x$. Combining these two inequalities, we get that $x > p_{1,3}$. However, x is the lightest job on M_1 and by Claim 2.4, $p_{1,3}$ is not empty and must consist of at least one job – having load at least x . A contradiction. \square

By Claim 4.4, the improvement ratio of x , which equals C_1/C'_1 , is the largest among the coalition. This ratio can now be bounded as follows:

$$\frac{C_1}{C'_1} = \frac{S_1 + p_{1,2} + p_{1,3} + x}{S_1 + p_{2,1} + p_{3,1} + x} < \frac{S_1 + p_{2,1} + p_{3,1} + 3x}{S_1 + p_{2,1} + p_{3,1} + x} \leq \frac{5}{3}.$$

The left inequality follows from Equation 4. The right one follows from Claim 4.5 and from the fact that S_1 might be empty.

(ii) x leaves M_1 . We assume w.l.o.g that x moves to M_2 . In This case, $C_1 = S_1 + p_{1,2} + p_{1,3} + x$; $C_2 = S_2 + p_{2,1}$; $C_3 = S_3 + p_{3,1}$, while after the coalition is active $C'_1 = S_1 + p_{2,1} + p_{3,1}$; $C'_2 = S_2 + p_{1,2} + x$; $C'_3 = S_3 + p_{1,3}$.

Since the jobs in $p_{1,2}$ and $p_{1,3}$ are part of the coalition, $C'_2 + C'_3 < 2C_1$. Deducing $x, p_{1,2}$ and $p_{1,3}$ from both sides we get $S_2 + S_3 < p_{1,2} + p_{1,3} + 2S_1 + x$. Combining with Equation 3, we get:

$$p_{1,2} + p_{1,3} < p_{2,1} + p_{3,1} + x. \quad (5)$$

Claim 4.6 $x \leq \min(p_{2,1}, p_{3,1})$

Proof: We first show $x \leq p_{2,1}$. Assume $x > p_{2,1}$, it means that when x is assigned to M_1 , the load on M_2 is composed of jobs that are a subset of S_2 only. Therefore, by the LPT rule, $S_2 \geq p_{1,2} + p_{1,3} + S_1$. Also, given that the jobs of $p_{1,2}$ are part of the coalition, we know that $S_2 + x < p_{1,3} + S_1 + x$. Combining these two inequalities, we get that $p_{1,2} < 0$. A contradiction. The proof of $x \leq p_{3,1}$ is identical to this proof as given in Claim 4.5 for the case where x remains on M_1 . \square

Claim 4.7 $x \leq \min(S_2, S_3)$

Proof: We first show $x \leq S_2$. Assume $x > S_2$, it means that when x is assigned to M_1 , the load on M_2 is composed of jobs that are a subset of $p_{2,1}$ only. Therefore, by the LPT rule, $p_{2,1} \geq S_1 + p_{1,2} + p_{1,3}$. However, by Claim 4.3, $C'_2 > C_2$, therefore $p_{2,1} < p_{1,2} + x$. Thus, $S_1 + p_{1,3} < x$. However, $x \leq p_{1,3}$. A contradiction. To show $x \leq S_3$, note that if $x < S_3$ then by a similar argument to the above $p_{3,1} \geq p_{1,2} + p_{1,3}$. By Claim 4.3, $C'_3 > C_3$. Therefore $p_{1,3} > p_{3,1}$, implying $S_1 + p_{1,2} < 0$. A contradiction. \square

If S_1 is not empty then the jobs of S_1 have improvement ratio C_1/C'_1 which is, by Claim 4.4, the largest ratio among the coalition. This ratio can now be bounded as follows:

$$\frac{C_1}{C'_1} = \frac{S_1 + p_{1,2} + p_{1,3} + x}{S_1 + p_{2,1} + p_{3,1}} \leq \frac{S_1 + p_{2,1} + p_{3,1} + 2x}{S_1 + p_{2,1} + p_{3,1}} < \frac{5}{3}.$$

The left inequality follows from Equation 5. The right one follows from Claim 4.6, and from the fact that S_1 is not empty and includes at least one job of load at least x .

If S_1 is empty, then as we show below, the maximal improvement ratio is less than $3/2$. We bound separately the improvement ratio of $p_{1,2}, p_{1,3}$, and $p_{i,1} (i \in \{1, 2\})$. Denote by $r_{i,j}$ the IR of jobs moving from M_i to M_j . In addition to Equations 2 and 5, and to Claims 4.6 and 4.7, we also use below Claim 4.3. Specifically, $p_{2,1} < p_{1,2} + x$ and $p_{3,1} < p_{1,3}$. Finally, bear in mind that $S_1 = \emptyset$.

$$r_{1,2} = \frac{C_1}{C'_2} = \frac{p_{1,2} + p_{1,3} + x}{S_2 + p_{1,2} + x} \leq \frac{S_2 + p_{2,1} + x}{S_2 + p_{1,2} + x} < \frac{S_2 + p_{2,1} + x}{S_2 + p_{2,1}} < \frac{3}{2}.$$

$$r_{1,3} = \frac{C_1}{C'_3} = \frac{p_{1,2} + p_{1,3} + x}{S_3 + p_{1,3}} \leq \frac{S_3 + p_{3,1} + x}{S_3 + p_{1,3}} < \frac{S_3 + p_{3,1} + x}{S_3 + p_{3,1}} < \frac{3}{2}.$$

$$r_{i,1} = \frac{C_i}{C'_1} = \frac{S_i + p_{i,1}}{p_{2,1} + p_{3,1}} < \frac{p_{1,2} + p_{1,3}}{p_{2,1} + p_{3,1}} < \frac{p_{2,1} + p_{3,1} + x}{p_{2,1} + p_{3,1}} < \frac{3}{2}.$$

□




Figure 5: An LPT-based coalition on m machines in which the job of load $1 + \varepsilon$ assigned to M_1 has improvement ratio arbitrarily close to $2 - \frac{1}{m}$.

The above analysis is tight, as demonstrated in Figure 5 for $m = 3$ (where the improvement ratio is $2 - \frac{1}{m} = \frac{5}{3}$). Moreover, this figure shows that this lower bound can be generalized for any $m \geq 3$. The job of load $1 + \varepsilon$ that remains on M_1 improves its cost from $2m - 1 + \varepsilon$ to $m(1 + \varepsilon)$, that is, for this job, j , $IR(j) = \frac{2m-1+\varepsilon}{m(1+\varepsilon)} = 2 - \frac{1}{m} - \delta$. Formally,

Theorem 4.8 *For any $m \geq 3$, there exists an LPT-based configuration s and a coalition Γ such that $IR_{max}(s, \Gamma) = 2 - \frac{1}{m} - \delta$ for an arbitrarily small $\delta > 0$.*

Note that the coalitional deviation in Figure 5 obeys the flower structure. We conjecture that the upper bound of $\frac{5}{3}$ for $m = 3$ can be generalized for any m , i.e., that for any LPT-based configuration s , and coalition Γ it holds that which $IR_{max}(s, \Gamma) < 2 - \frac{1}{m}$.

5 Maximum Damage Ratio

In this section, the quality of a configuration is measured by $\max_{j \notin \Gamma} DR(j)$. Recall that $DR(j) = \frac{C'_i}{C_i}$, where i is the machine on which j is scheduled. For non-deviating jobs, this ratio might be larger than 1, and we would like to bound its maximal possible value. We provide a complete analysis for NE and LPT-based configurations and any $m \geq 3$. Once again, we find out that LPT provides a better performance guarantee compared to general NE: the cost of any job in an LPT schedule cannot increase by a factor $\frac{3}{2}$ or larger, while it can increase by a factor arbitrarily close to 2 for NE schedules.

Theorem 5.1 *For any m , the damage ratio caused by any NE-based coalition is less than 2.*

Proof: Let Γ be a coalition. Let M_1 be the most loaded machine participating in the coalition. There are at least two jobs on M_1 , since otherwise, M_1 has the longest job in the instance, and being the only job on a machine, it cannot benefit from being part of Γ . This implies that for any machine $i \neq 1$ participating in the coalition, $C_i > C_1/2$, since otherwise, at least one of the jobs on M_1 can benefit from moving to M_i – contradicting the fact that this is an NE. Also, $C'_i < C_1$, since otherwise the jobs migrating into M_i do not benefit. Combining the above bounds, we get that for any job on M_i , $DR(j) = \frac{C'_i}{C_i} < \frac{C_1}{C_i} < \frac{2C_1}{C_i} = 2$. The above is valid for any $i \neq 1$. To complete the proof note that jobs on M_1 can only benefit from the coalition action. \square

The above analysis is tight as shown in Figure 4: The damage ratio of the jobs of load $2r - 1$ is $(4r - 1)/(2r)$, which can be arbitrarily close to 2. Formally,

Theorem 5.2 *For any $m \geq 3$, there exists a NE-based configuration s and a coalition Γ such that $DR_{\max}(s, \Gamma) = 2 - \delta$ for an arbitrarily small $\delta > 0$.*

For LPT-based coalitions we obtain a smaller bound:

Theorem 5.3 *For any m , the damage ratio caused by any LPT-based coalition is less than $\frac{3}{2}$.*

Proof: Let M_1 be the most loaded machine in the coalition. M_1 must have at least 2 jobs. Let x be the load of the last job assigned to M_1 , and let $\ell = C_1 - x$. For every machine in the coalition, it must hold that $C_i \geq \ell$ (since else, x would not have been assigned to M_1), and $C'_i < \ell + x$ (since all jobs must improve).

case (a): $\ell \geq 2x$, and then for any machine M_i , $\frac{C'_i}{C_i} < \frac{\ell+x}{\ell} \leq \frac{3}{2}$.

case (b): $\ell < 2x$. We show that no coalition exists in this case. If $\ell < 2x$, then (by LPT) M_1 has exactly 2 jobs, of loads ℓ and x . By LPT, every other machine must have (i) one job of load at least ℓ (and possibly other small jobs), or (ii) two jobs of load at least x (and possible other small jobs). Let k and k' be the number of machines of type (i) and (ii), respectively (excluding M_1). Thus, there is a total of $k + 1$ jobs of load ℓ and $2k' + 1$ jobs of load x . After the deviation, no machine can have jobs of load ℓ and x together, nor can it have three jobs of load x . The $k + 1$ machines assigned with the $k + 1$ jobs of load ℓ after the deviation cannot be assigned any other job of load x . So, we end up with $2k' + 1$ jobs of load x that should be assigned to k'

machines. Thus, there must be a machine with at least three jobs of load x . Contradiction. \square




Figure 6: An LPT-based coalition, in which the damage ratio of the job of load $1+2\epsilon$ on M_3 is arbitrarily close to $\frac{3}{2}$.

The above analysis is tight as shown in Figure 6. Moreover, by adding dummy machines and jobs it can be extended to any $m \geq 3$. Formally,

Theorem 5.4 *For any $m \geq 3$, there exists an LPT-based configuration s and a coalition Γ such that $DR_{max}(s, \Gamma) = \frac{3}{2} - \delta$ for an arbitrarily small $\delta > 0$.*

6 Computational Complexity

It is easy to see that one can determine whether a given configuration is a NE in polynomial time. Yet, for SE, this task is more involved. In this section, we provide some hardness results about coalitional deviations.




Figure 7: Partition induces a coalition in a schedule on identical machines.

Theorem 6.1 *Given a NE schedule on $m \geq 3$ identical machines, it is NP-hard to determine if it is a SE.*

Proof: We give a reduction from *Partition*. Given a set A of n integers a_1, \dots, a_n with total size $2B$, and the question whether there is a subset of total size B , construct the schedule in Figure 7(a). In this schedule on three machines there are $n+4$ jobs of loads $a_1, \dots, a_n, B-2, B-2, B-1, B-1$. We assume w.l.o.g. that $\min_i a_i \geq 3$, else the whole instance can be scaled. Thus, schedule 7(a) is a NE. For $m \geq 3$, add $m-3$ machines each with a single job of load $2B$.

Claim 6.2 *The NE schedule in Figure 7(a) is a SE if and only if there is no partition.*

Proof: If there is a partition into S_1, S_2 , each having total size B , then the schedule in Figure 7(b) is better for the jobs originated from the partition instance and for the two $(B-2)$ -jobs.

All the partition jobs improved from cost $2B$ to cost $2B - 1$, and the $(B - 2)$ -jobs improved from $2B - 3$ to $2B - 4$.

Next, we show that if there is no partition then the initial schedule SE. By Theorem 2.5, in any action of a coalition on 3 machines, jobs must migrate to M_1 from both M_2 and M_3 . In order to decrease the load from $2B - 3$, the set of jobs migrating to M_1 must be the set of two jobs of load $B - 2$. Also, it must be that all the partition jobs move away from M_1 - otherwise, the total load on M_1 will be at least $2B - 4 + 1 = 2B - 3$, which is not an improvement for the $(B - 2)$ -jobs. This implies that the jobs of M_1 split between M_2 and M_3 . However, since there is no partition, one of the two subsets is of total load at least $B + 1$. These jobs will join a job of load $B - 1$ to get a total load of at least $2B$, which is not an improvement over the $2B$ -load in the initial schedule. \square

\square

A direct corollary of the above proof is the following:

Corollary 6.3 *Given a NE schedule and a coalition, it is NP-hard to determine whether the coalition can deviate.*

Theorem 6.1 holds for any $m \geq 3$ identical machines. For $m \leq 2$, a configuration is a NE if and only if it is a SE [3], and therefore it is possible to determine whether a given configuration is SE in polynomial time. Yet, the following theorem shows that for the case of unrelated machines, the problem is NP-hard already for $m = 2$.

Theorem 6.4 *Given a NE schedule on $m \geq 2$ unrelated machines, it is NP-hard to determine if it is a SE.*

Proof: We give a reduction from *Partition*. Given n integers a_1, \dots, a_n with total size $2B$, and the question whether there is a subset of total size B , construct the following instance for scheduling: there are 2 machines and $n + 1$ jobs with the following loads (for $\varepsilon < 1/(n - 1)$):

$$p_{i,1} = a_i + \varepsilon \text{ and } p_{i,2} = 2a_i + \varepsilon, \forall i \in \{1, \dots, n\}; p_{n+1,1} = B, \text{ and } p_{n+1,2} = 2B + n\varepsilon.$$

Consider the schedule in which all the jobs $1, \dots, n$ are on M_1 , and J_{n+1} is on M_2 . The completion times of both machines are $2B + n\varepsilon$. It is a NE.




Figure 8: Partition induces a coalition in a schedule on related machines.

Claim 6.5 *The NE schedule in Figure 8(a) is a SE if and only if there is no partition.*

Proof: If there is a partition into A_1, A_2 , each having total size B , then the schedule given in Figure 8(b) is better for everyone. The completion time of M_1 is $2B + |A_1|\varepsilon < 2B + n\varepsilon$ and the

completion time of M_2 is $2B + |A_2|\varepsilon < 2B + n\varepsilon$.

Next, we show that if there is no partition then the initial schedule is a SE. Since there is no partition, in any partition into A_1, A_2 , one of the two subsets, w.l.o.g., A_1 has total size at least $B + 1$. A_1 will only increase its load by migrating to M_2 even alone (bearing a load of at least $2B + 2 + |A_1|\varepsilon$ instead of $2B + n\varepsilon$). Therefore, A_1 will not leave M_1 . However, if A_1 stays at M_1 , job $n + 1$ is better-off staying at M_2 (since if it migrates, it bears a load of at least $2B + 1 + |A_1|\varepsilon$ which is not smaller than $2B + n\varepsilon$ for any $|A_1|$ and $\varepsilon \leq 1/(n - 1)$). \square

\square

A direct corollary of the above proof is the following:

Corollary 6.6 *Given an NE schedule on unrelated machines and a coalition, it is NP-hard to determine whether the coalition can deviate.*

It remains an open problem whether there exists a polynomial time approximation scheme that provides a $(1 + \varepsilon)$ -SE.

7 Conclusions and Open Problems

In this paper we studied the stability and performance of NE and LPT with respect to the three measurements of IR_{min} , IR_{max} , and DR_{max} . We proved upper and lower bounds for NE and LPT-based configurations. Some of the problems that remain open are:

1. For IR_{min} there is a gap between the upper and lower bound for $m > 3$. We believe that the lower bound of $\frac{5}{4}$ for NE and $\frac{2}{\sqrt{34}-4}$ for LPT that we proved to be tight for $m = 3$, are tight also for arbitrary m . Is there a NE-configuration (on $m > 3$ machines) that is not a $\frac{5}{4}$ -SE? Is there an LPT configuration that is not a $\frac{2}{\sqrt{34}-4}$ -SE?
2. For IR_{max} , and $m \geq 3$ we presented a lower bound of $2 - \frac{1}{m}$ and a matching upper bound of $\frac{5}{3}$ for $m = 3$. We believe that the upper bound is $2 - \frac{1}{m}$ for any m . Is there an LPT-based configuration s , such that there exists a coalition Γ for which $IR_{max}(s, \Gamma) > 2 - \frac{1}{m}$?
3. Is there a polynomial time approximation scheme for the minimum makespan problem that also provides a $(1 + \varepsilon)$ -SE?
4. In this paper we introduced three general measurements for the stability and performance of scheduling profiles in the context of deviations by coalitions. We believe that these measurements can be used to measure the stability of various algorithms to coalitional deviations and their performance in additional settings and games. We hope to see more work that makes use of these measurements within the framework of algorithmic game theory.

Acknowledgments. We thank Oded Schwartz for helpful discussions.

References

- [1] S. Albers. On the value of coordination in network design. In *SODA* 2008.
- [2] S. Albers, S. Elits, E. Even-Dar, Y. Mansour, and L. Roditty. On Nash Equilibria for a Network Creation Game. In *SODA* 2006.
- [3] N. Andelman, M. Feldman, and Y. Mansour. Strong Price of Anarchy. In *SODA* 2007.
- [4] E. Anshelevich, A. Dasgupta, J. M. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *FOCS*, pages 295–304, 2004.
- [5] E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler. Near-Optimal Network Design with Selfish Agents. In *STOC*, 2003.
- [6] R. Aumann. Acceptable Points in General Cooperative n-Person Games. In *Contributions to the Theory of Games*, volume 4, 1959.
- [7] B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in Worst-Case Equilibria. In *WAOA*, 2003.
- [8] D. B. Bernheim, B. Peleg, and M. D. Whinston. Coalition-proof nash equilibria: I concepts. *Journal of Economic Theory*, 42:1–12, 1987.
- [9] G. Christodoulou and E. Koutsoupias. On the Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games. In *ESA*, 2005.
- [10] G. Christodoulou, E. Koutsoupias, and A. Nanavati. Coordination mechanisms. *J. Daz, J. Karhumaki, A. Lepisto, and D. Sannella (Eds.), Automata, Languages and Programming, LNCS Vol. 3142: Springer*, (2):345–357, 2004.
- [11] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *SODA*, pages 413–420, 2002.
- [12] A. Epstein, M. Feldman, and Y. Mansour. Strong Equilibrium in Cost Sharing Connection Games. In *ACMEC*, 2007.
- [13] A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou, and S. Shenker. On a network creation game. In *PODC* 2003.
- [14] A. Fiat, H. Kaplan, M. Levi, and S. Olonetsky. Strong Price of Anarchy for Machine Load Balancing. In *ICALP*, 2007.
- [15] G. Finn and E. Horowitz. A linear time approximation algorithm for multiprocessor scheduling. *BIT Numerical Mathematics*, 19(3):312–320, 1979.

- [16] D. Fotakis, S. Kontogiannis, M. Mavronicolas, and P. Spiraklis. The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In *ICALP*, pages 510–519, 2002.
- [17] R. Graham. Bounds on multiprocessing timing anomalies. *SIAM J. Appl. Math.*, 17:263–269, 1969.
- [18] R. Holzman and N. Law-Yone. Strong equilibrium in congestion games. *Games and Economic Behavior*, 21:85–101, 1997.
- [19] R. Holzman and N. Law-Yone. Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences*, 46:193–205, 2003.
- [20] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413, 1999.
- [21] S. Leonardi and P. Sankowski. Network Formation Games with Local Coalitions. In *PODC*, 2007.
- [22] I. Milchtaich. Crowding games are sequentially solvable. *International Journal of Game Theory*, 27:501–509, 1998.
- [23] C. H. Papadimitriou. Algorithms, games, and the internet. In *proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 749–753, 2001.
- [24] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- [25] O. Rozenfeld and M. Tennenholtz. Strong and correlated strong equilibria in monotone congestion games. In *working paper. Technion, Israel*, 2006.
- [26] P. Schuurman and T. Vredeveld. Performance guarantees of local search for multiprocessor scheduling. *INFORMS Journal on Computing*, to appear.