Approximate Strong Equilibrium in Job Scheduling Games

Michal Feldman[1,⋆] and Tami Tamir[2]

[1] School of Business Administration and Center for the Study of Rationality,
Hebrew University of Jerusalem
`mfeldman@cs.huji.ac.il`
[2] School of Computer Science, The Interdisciplinary Center, Herzliya, Israel
`tami@idc.ac.il`

Abstract. A Nash Equilibriun (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.

[⋆] Research partially supported by a grant of the Israel Science Foundation, BSF, Lady Davis Fellowship, and an IBM faculty award.

B. Monien and U.-P. Schroeder (Eds.): SAGT 2008, LNCS 4997, pp. 58–69, 2008.

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 to[1]. These problems have been widely studied in recent years from a game theoretic perspective [21,3,10,11,15]. 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 [24].

Most game theoretic models applied to job scheduling problems, as well as other network games (e.g., [13,2,25,4]), 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.

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

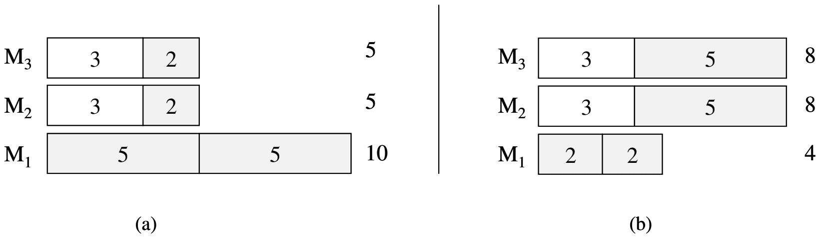

Fig. 1. An example of a configuration (a) that is a Nash equilibrium but is not resilient against coordinated deviations, since the jobs of load $\{5, 5, 2, 2\}$ all profit from the deviation demonstrated in (b)

---

[1] This 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.

also sustainable against joint deviations. Yet, as was already observed in [3], this may not be true[2]. 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 factor [3] . 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 approximation[4] and approximate stability.

We also consider a subclass of NE, produced by the *Longest Processing Time* (LPT) rule [18]. 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

---

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

[3] Throughout 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.

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

verify that every configuration produced by LPT is a NE [17]. Is it also a SE? Note that for the instance depicted in Figure 1, LPT would 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).

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

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

	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	$\frac{2}{\sqrt{34}-4}$	$\frac{4}{3} - \frac{1}{3m}$	$\frac{2}{\sqrt{34}-4}$	$\frac{5}{3}$ $(m=3)$	$2 - \frac{1}{m}$	$\frac{3}{2}$	$\frac{3}{2}$

$m = 2$ 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 [24,21] as the ratio between the worst-case NE and the optimum solution, and has been extensively studied in various settings, including job scheduling [21,10,11], network design [2,4,5,13], network routing [25,7,9], 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, [12,19,20,26] 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 [19,20,23,8] 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 preoperties. 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 [3,15,22].

2 Model and Preliminaries

In our job scheduling setting there is a set of m identical machines, $M = \{M_1, \ldots, M_m\}$, and n jobs, $N = \{1, \ldots, n\}$, 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 $s \in S : N \to M$ (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 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. All missing proofs (from this section as well as other sections) are given in the full version of this paper [14].

Observation 1 *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. □

Definition 2. *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 $i, j > 1$, $P_{i,j} = 0$.*

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

□

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 1. *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 2. *Any NE schedule on three machines is a $\frac{5}{4}$-SE.* □

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

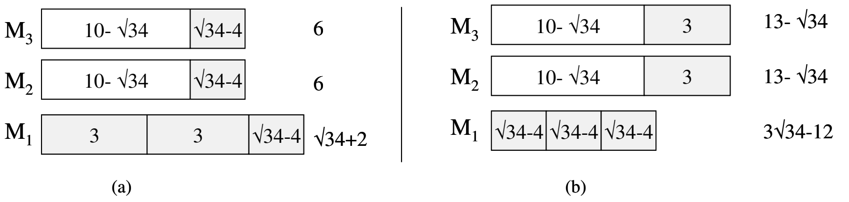

Fig. 2. An LPT-based coalition on 3 machines in which all migrating jobs improve by $\frac{2}{\sqrt{34}-4}$

The above analysis is tight as shown in Figure 2. Moreover, as for NE, this lower bound can be extended to any $m > 3$ by adding dummy jobs and machines. Thus,

Theorem 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.

Theorem 6. *Any schedule produced by LPT on m identical machines is a* $(\frac{4}{3} - \frac{1}{3m}) - SE$. ☐

Theorem 7. *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 8. *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 3(a). The coalition consists of $\{1, 1, 2r, 2r\}$. Their improved schedule is given in Figure 3(b). The improvement ratio of the jobs of load 1 is $2r/2 = r$. For $m > 3$, dummy machines and jobs can be added. ☐

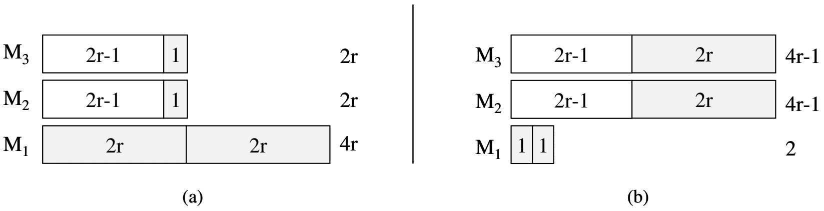

Fig. 3. An NE-based coalition in which the jobs of load 1 have improvement ratio r

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

Theorem 9. *For any LPT schedule on three machines, the maximum improvement ratio of any coalition is less than $\frac{5}{3}$.* ☐

The above analysis is tight, as demonstrated in Figure 4 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 10. *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 $\delta > 0$.*

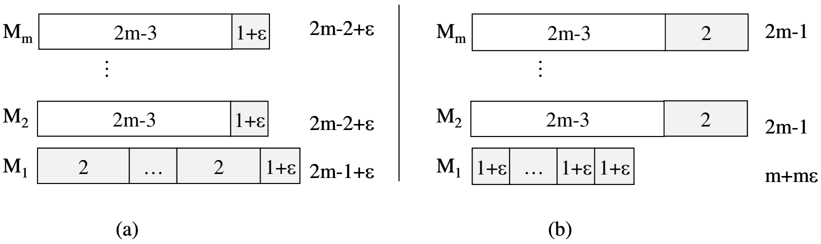

Fig. 4. An LPT-based coalition on m machines in which the job of load $1+\varepsilon$ assigned to M_1 has improvement ratio arbitrarily close to $2-\frac{1}{m}$

Note that the coalitional deviation in Figure 4 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 11. *For any m, the damage ratio caused by any NE-based coalition is less than* 2. □

The above analysis is tight as shown in Figure 3: The damage ratio of the jobs of load $2r-1$ is $(4r-1)/(2r)$, which can be arbitrarily close to 2. Formally,

Theorem 12. *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 13. *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. □

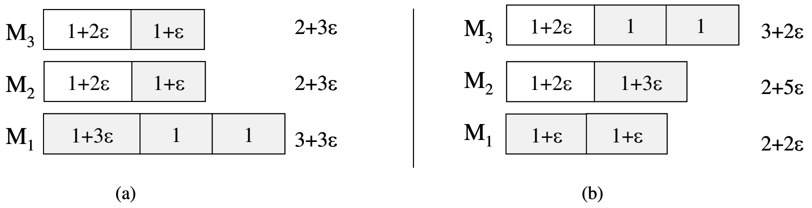

Fig. 5. An LPT-based coalition, in which the damage ratio of the job of load $1+2\varepsilon$ on M_3 is arbitrarily close to $\frac{3}{2}$

The above analysis is tight as shown in Figure 5. Moreover, by adding dummy machines and jobs it can be extended to any $m \geq 3$. Formally,

Theorem 14. *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.

Theorem 15. *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 $a_1, \ldots, a_n$ with total size $2B$, and the question whether there is a subset of total size B, construct the schedule in Figure 6(a). In this schedule on three machines there are $n+4$ jobs of loads $a_1, \ldots, a_n, B-2, B-2, B-1, B-1$. We assume w.l.o.g. that $min_i a_i \geq 3$, else the whole instance can be scaled. Thus, schedule 6(a) is a NE. For $m \geq 3$, add $m-3$ machines each with a single job of load $2B$.

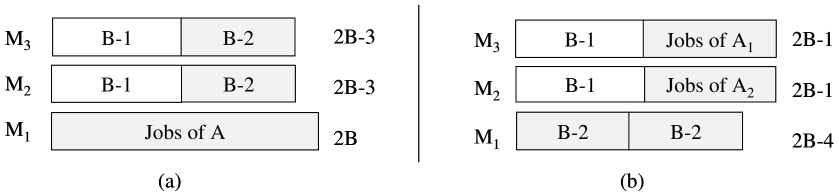

Fig. 6. Partition induces a coalition in a schedule on identical machines

Claim. The NE schedule in Figure 6(a) is a SE if and only if there is no partition. □

A direct corollary of the above proof is the following:

Corollary 2. *Given a NE schedule and a coalition, it is NP-hard to determine whether the coalition can deviate.*

Theorem 15 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 16. *Given a NE schedule on $m \geq 2$ unrelated machines, it is NP-hard to determine if it is a SE.* □

A direct corollary of the above proof is the following:

Corollary 3. *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.

Acknowledgments. We thank Oded Schwartz for helpful discussions.

References

1. Albers, S.: On the value of coordination in network design. In: SODA (2008)
2. Albers, S., Elits, S., Even-Dar, E., Mansour, Y., Roditty, L.: On Nash Equilibria for a Network Creation Game. In: SODA (2006)
3. Andelman, N., Feldman, M., Mansour, Y.: Strong Price of Anarchy. In: SODA (2007)
4. Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. In: FOCS, pp. 295–304 (2004)
5. Anshelevich, E., Dasgupta, A., Tardos, E., Wexler, T.: Near-Optimal Network Design with Selfish Agents. In: STOC (2003)

6. Aumann, R.: Acceptable Points in General Cooperative n-Person Games. In: Conti, R., Ruberti, A. (eds.) Optimization Techniques 1973. LNCS, vol. 4, p. 1959. Springer, Heidelberg (1973)
7. Azar, Y., Tsur, D., Richter, Y., Awerbuch, B.: Tradeoffs in Worst-Case Equilibria. In: Solis-Oba, R., Jansen, K. (eds.) WAOA 2003. LNCS, vol. 2909, pp. 41–52. Springer, Heidelberg (2004)
8. Bernheim, D.B., Peleg, B., Whinston, M.D.: Coalition-proof nash equilibria: I concepts. Journal of Economic Theory 42, 1–12 (1987)
9. Christodoulou, G., Koutsoupias, E.: On the Price of Anarchy and Stability of Correlated Equilibria of Linear Congestion Games. In: Brodal, G.S., Leonardi, S. (eds.) ESA 2005. LNCS, vol. 3669, pp. 59–70. Springer, Heidelberg (2005)
10. Christodoulou, G., Koutsoupias, E., Nanavati, A.: Coordination Mechanisms. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 345–357. Springer, Heidelberg (2004)
11. Czumaj, A., Vöcking, B.: Tight bounds for worst-case equilibria. In: SODA, pp. 413–420 (2002)
12. Epstein, A., Feldman, M., Mansour, Y.: Strong Equilibrium in Cost Sharing Connection Games. In: ACMEC (2007)
13. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C., Shenker, S.: On a network creation game. In: PODC (2003)
14. Feldman, M., Tamir, T.: Approximate Strong Equilibrium in Job Scheduling Games. `http://www.faculty.idc.ac.il/tami/Papers/approxSE.pdf`
15. Fiat, A., Kaplan, H., Levi, M., Olonetsky, S.: Strong Price of Anarchy for Machine Load Balancing. In: ICALP (2007)
16. Finn, G., Horowitz, E.: A linear time approximation algorithm for multiprocessor scheduling. BIT Numerical Mathematics 19(3), 312–320 (1979)
17. Fotakis, D., Kontogiannis, S., Mavronicolas, M., Spiraklis, P.: The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. In: Widmayer, P., Triguero, F., Morales, R., Hennessy, M., Eidenbenz, S., Conejo, R. (eds.) ICALP 2002. LNCS, vol. 2380, pp. 510–519. Springer, Heidelberg (2002)
18. Graham, R.: Bounds on multiprocessing timing anomalies. SIAM J. Appl. Math. 17, 263–269 (1969)
19. Holzman, R., Law-Yone, N.: Strong equilibrium in congestion games. Games and Economic Behavior 21, 85–101 (1997)
20. Holzman, R., Law-Yone, N.: Network structure and strong equilibrium in route selection games. Mathematical Social Sciences 46, 193–205 (2003)
21. Koutsoupias, E., Papadimitriou, C.H.: Worst-Case Equilibria. In: Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
22. Leonardi, S., Sankowski, P.: Network Formation Games with Local Coalitions. In: PODC (2007)
23. Milchtaich, I.: Crowding games are sequentially solvable. International Journal of Game Theory 27, 501–509 (1998)
24. Papadimitriou, C.H.: Algorithms, games, and the internet. In: proceedings of the 33rd Annual ACM Symposium on Theory of Computing, pp. 749–753 (2001)
25. Roughgarden, T., Tardos, E.: How bad is selfish routing? Journal of the ACM 49(2), 236–259 (2002)
26. Rozenfeld, O., Tennenholtz, M.: Strong and correlated strong equilibria in monotone congestion games. In: working paper, Technion, Israel (2006)
27. Schuurman, P., Vredeveld, T.: Performance guarantees of local search for multiprocessor scheduling. INFORMS Journal on Computing (to appear)