Scheduling with Machine-Dependent Priority Lists



Tami Tamir

Based on joint work with Vipin Ravindran Vijayalakshmi and Marc Schroder

# **Traditional Scheduling Algorithms**

- A centralized authority (a scheduler) determines the outcome.
- The centralized authority aims to maximize the system's utilization and the total users' welfare.
- All the users obey it.



# Job Scheduling Games

- The jobs are controlled by selfish agents who select the jobs assignment.
- No centralized authority.

Every job selects its machine, trying to maximize its own utility



### **Coordinated Mechanism**

Machines have a local scheduling policy. The jobs know this policy and select their machine accordingly.

Example: Assume that all the machines schedule the jobs in LPT (Longest first) order.







If the local policy is LPT, I'll better join  $M_1$ . This is my best-response.



If the local policy is SPT, my best-response is to join  $M_2$ .

# Coordinated Mechanism SPT Policy





Consider the 2<sup>nd</sup> job of length 7 in the resulting schedule



Hey, I can migrate and reduce my completion time from 21 to 14

7

Now, other jobs may have a beneficial migration...

# Best Response Dynamics (BRD)

- A local search method.
- Players proceed in turns, each performing a selfish improving step.
- An important question: Does BRD converge to a

pure Nash equilibrium.

A stable profile in which no player has an improving step.

# Our Work

We study coordinated mechanisms in which different machines may have different local policies.

For the associated game, we analyzed:

- Nash equilibrium existence and calculation
- BRD convergence
- Equilibrium inefficiency

Not less important: We studied the centralized version of this setting.

• A set J of n jobs



- Every job  $j \in J$  has processing time  $p_i$
- A set M of m parallel machines
  - Every machine  $i \in M$  has speed  $s_i$  and a priority list  $\pi_i: J \rightarrow \{1,...,n\}$ , defining its scheduling policy.





Example: J={ 
$$p_j = 2$$
  $4$   $2$   $2$   $1$  pr

processing times

m=2,  

$$s_1=1$$
  $\pi_1 = (e,d,c,b,a)$   
 $s_2=0.5$   $\pi_2 = (a,b,c,d,e)$ 



A profile of the game: A schedule  $\sigma: J \rightarrow M$ .  $C_j(\sigma)$  = the completion time of job j in profile  $\sigma$ 

Example: 
$$J=\{ \begin{array}{ccc} & & & & & \\ a & b & c & d & e \\ p_j = & 2 & 4 & 2 & 2 & 1 \\ \end{array} \right. product of the product o$$

processing times

m=2,  

$$s_1=1$$
  $\pi_1 = (e,d,c,b,a)$   
 $s_2=0.5$   $\pi_2 = (a,b,c,d,e)$ 



Does anyone have a beneficial migration?

Example: 
$$J=\{ \begin{array}{ccc} & 0 & 0 & 0 & 0 \\ a & b & c & d & e \\ p_j= & 2 & 4 & 2 & 2 & 1 \\ \end{array} \right)$$

processing times

m=2,  

$$s_1=1$$
  $\pi_1 = (e,d,c,b,a)$   
 $s_2=0.5$   $\pi_2 = (a,b,c,d,e)$ 



# **Game Theory Definitions**

A profile is a pure Nash equilibrium (NE) if no job can reduce its completion time by changing its strategy (migrating to a different machine)

A social optimum (SO) of a game is a profile that attains some optimality criteria. For example:

social optimum w.r.t total flow time (=sum of C<sub>j</sub>) social optimum w.r.t makespan (=maximal C<sub>i</sub>).

SO = Optimal solution for the centralized problem  $P|\pi|C_{max}$  or  $P|\pi|\sum_j C_j$ 



### **Interesting Questions**

- Calculating a NE for a given game instance
- What is the equilibrium inefficiency?

Price of Anarchy = worst NE / SO Price of Stability = best NE / SO

- Convergence of **Best-Response Dynamics** 











# Back to our example

$$s_1 = 1$$
  $\pi_1 = (e,d,c,b,a)$   
 $s_2 = 0.5$   $\pi_2 = (a,b,c,d,e)$ 



C<sub>max</sub>=8 A social optimum w.r.t Makespan, but not a NE.



A possible NE profile. C<sub>max</sub>=9

Price of anarchy  $\geq 9/8$ 

# **Related Work**

- Koutsoupias and Papadimitriou (1999)
- Czumaj and Vocking (2003)
- Christodoulou, Koutsoupias and Nanavati (2004)
- Cole, Correa, Gkatzelis, Mirrokni and Olver (2015)
- Immorlica, Li, Mirrokni and Schulz (2005)
- Farzad, Olver and Vetta (2008)
- Correa and Queyranne (2012)
- Cole, Correa, Gkatzelis, Mirrokni and Olver (2015)
- Hoeksma and Uetz (2019)
- Bosman, Frascaria, Olver, Sitters, Stougie (2019)

Selfish Scheduling / Coordinated mechanism /Prioritybased model of routing/ The centralized problem.

# **NE** Calculation

Given  $\langle J, M, \{s_i\}, \{\pi_i\} \rangle$ , calculate a NE profile





m=3, M={M<sub>1</sub>, M<sub>2</sub>, M<sub>3</sub>}  

$$s_1=1$$
  $\pi_1 = (a, b, c, d, e)$   
 $s_2=s_2=0.5$   $\pi_2 = \pi_3 = (e, d, b, c, a)$ 

One fast machine.

Two slow machines



Since a is the first on  $\pi_1$ , it is first on  $M_1$  in any NE schedule.

 $s_3=0.5, \pi_3=(e,d,b,c,a)$  $s_2=0.5, \pi_2=(e,d,b,c,a)$  $s_1=1, \pi_1=(a,b,c,d,e)$ 





Given that  $\boxed{a}$  is on  $M_1$ , Since  $\boxed{e}$  is the first on  $\pi_2$ , it is first on  $M_2$  (w.l.o.g) in any NE schedule

s<sub>3</sub>=0.5,  $\pi_3$  =(e,d,b,c,a) s<sub>2</sub>=0.5,  $\pi_2$  =(e,d,b,c,a) s<sub>1</sub>=1,  $\pi_1$  =(a,b,c,d,e)





So we know that if a NE exists, then a and e are first on  $M_1$ and  $M_2$ , respectively. Let's consider the possible locations of job d

s<sub>3</sub>=0.5,  $\pi_3$  =(e,d,b,c,a) s<sub>2</sub>=0.5,  $\pi_2$  =(e,d,b,c,a) s<sub>1</sub>=1,  $\pi_1$  =(a,b,c,d,e)







 $s_3=0.5, \pi_3=(e,d,b,c,a)$  $s_2=0.5, \pi_2=(e,d,b,c,a)$  $s_1=1, \pi_1=(a,b,c,d,e)$ 



Therefore, there is no NE in which d is on  $M_1$ 



 $s_3=0.5, \pi_3=(e,d,b,c,a)$  $s_2=0.5, \pi_2=(e,d,b,c,a)$  $s_1=1, \pi_1=(a,b,c,d,e)$ 



Therefore, there is no NE in which d is on  $M_2$ 





s<sub>3</sub>=0.5, 
$$π_3 = (e,d,b,c,a)$$
  
s<sub>2</sub>=0.5,  $π_2 = (e,d,b,c,a)$   
s<sub>1</sub>=1,  $π_1 = (a,b,c,d,e)$ 



Therefore, there is no NE in which  $\square$  is on  $M_3$ 



We conclude that there are games in which a NE does not exist



Can we characterize games that have a NE?

Unfortunately, No.

Theorem: Given an instance of a scheduling game, it is NP-complete to decide whether the game has a NE. Proof: Reduction from 3-bounded 3-dimensional matching



#### On the other hand:

We identified four classes of games for which a NE is guaranteed to exist.

 $G_1$ : Unit Jobs  $G_2$ : Two machines  $G_3$ : Identical machines  $G_4$ : Global priority list

Note: This characterization is tight. In our No-NE example, there are three machines, two of them are identical (same speed and same priority list)



### For each of the four classes, we present:

- A polynomial time algorithm for computing a NE
- A proof that BRD converges to a NE
- Tight analysis of the equilibrium inefficiency:

Objective Instance class	Makespan PoA and PoS	Sum of Completion Time PoA and PoS
$\mathcal{G}_1$ : Unit Jobs	1	1
$\mathcal{G}_2$ : Two machines	$\frac{\sqrt{5}+1}{2}$	$\Theta(n)$
$\mathcal{G}_3$ : Identical machines	$2-\frac{1}{m}$	$\Theta\left(\frac{n}{m}\right)$
$\mathcal{G}_4$ : Global priority list	$\Theta(m)$	$\Theta(n)$

#### Two machines

Theorem: If m = 2, then a NE exists and can be calculated efficiently.

**Proof:** Algorithm



#### Two machines

#### Algorithm:

- 1. Assign all the jobs on  $M_1$  according to  $\pi_1$ .
- 2. For k = 1,...,n, let job j for which  $\pi_2(j) = k$  perform

a best-response move.



### Two machines



Claim: The algorithm produces a NE. Proof:

Let  $\boldsymbol{\sigma}$  denote the schedule produced by the algorithm.

- **1**. Jobs on M<sub>1</sub> have no incentive to deviate (easy).
- **2**. Suppose a job j on  $M_2$  has an incentive to deviate.

Let  $\Delta$  be the set of jobs that have a higher priority on  $M_1$  than j and moved to  $M_2$  after j.





s<sub>1</sub>=1

 $s_2 = s \le 1$ 

Two machines

(ii)  $(P_B + p_i + P_\Delta)/s_2 < P_A + P_\Delta$ 

(i)  $P_A + p_i < (P_B + p_i)/s_2$ 

 $p_j + P_{\Delta}/s_2 < P_{\Delta}$ 

A Contradiction (since  $p_i \ge 0$  and  $s_2 \le 1$ )



Remark: A possible generalization of our setting considers unrelated machines (p<sub>ij</sub> is the processing time of job i if processed on machine j).

In this environment, a NE need not exist already with only two unrelated machines.

# Equilibrium inefficiency

The makespan of a profile  $\sigma$ , is  $C_{max}(\sigma) = max_{j \in J}C_j(\sigma)$ For a game G,

$$PoA(G) = \frac{max_{\sigma \in NE(G)} C_{max}(\sigma)}{min_{\sigma^*} C_{max}(\sigma^*)} =$$

makespan of the worst NE schedule min makespan (social optimum)

For a class of games G, define  $PoA(G) = sup_{G \in G} PoA(G)$ 

# Equilibrium inefficiency

The makespan of a profile  $\sigma$ , is  $C_{max}(\sigma) = max_{j \in J}C_j(\sigma)$ For a game G,

$$PoA(G) = \frac{max_{\sigma \in NE(G)} C_{max}(\sigma)}{min_{\sigma^*} C_{max}(\sigma^*)} =$$

makespan of the worst NE schedule min makespan (social optimum)

For a class of games G, define  $PoA(G) = sup_{G \in G} PoA(G)$ 

Instance class	Makespan PoA	
$\mathcal{G}_1$ : Unit Jobs	1	
$\mathcal{G}_2$ : Two machines	$\frac{\sqrt{5}+1}{2}$	
$\mathcal{G}_3$ : Identical machines	$2-\frac{1}{m}$	
$\mathcal{G}_4$ : Global priority list	$\Theta(m)$	

Theorem: Let G be a game played on two machines,  $s_1 = 1$ and  $s_2 \le 1$ , then  $PoA(G) \le min\left\{1 + s_2, 1 + \frac{1}{1 + s_2}\right\}$ 

Since 
$$1 + s = 1 + \frac{1}{1+s}$$
 for  $s = \frac{\sqrt{5}-1}{2}$ , the theorem implies that  $PoA(\mathcal{G}_2) \le \frac{\sqrt{5}+1}{2}$ .

Theorem: Let G be a game played on two machines,  $s_1 = 1$ and  $s_2 \le 1$ , then  $PoA(G) \le min\left\{1 + s_2, 1 + \frac{1}{1 + s_2}\right\}$ 

Proof: Let  $\sigma$  be a NE.

1. 
$$C_{max}(\sigma) \leq \sum_{j \in J} p_j$$
 (if all jobs on fast machine)  
2.  $C_{max}(\sigma^*) \geq \frac{\sum_{j \in J} p_j}{1 + s_2}$  (balanced)

Implying that  $C_{max}(\sigma) \leq (1 + s_2) \cdot C_{max}(\sigma^*).$ 

Theorem: Let G be a game played on two machines,  $s_1 = 1$ and  $s_2 \le 1$ , then  $PoA(G) \le min\left\{1 + s_2, 1 + \frac{1}{1 + s_2}\right\}$ 

Proof: Let a be the last job to complete in a NE  $\sigma$ .

$$C_{max}(\sigma) \leq \frac{p_a + \sum_{j \in J} p_j}{1 + s_2} \leq \left(1 + \frac{1}{1 + s_2}\right) \cdot C_{max}(\sigma^*).$$

Theorem: For every  $s \le 1$ , there exists a game with  $s_1 = 1$ ,  $s_2 = s$ , and  $PoS(G) = min\left\{1 + s, 1 + \frac{1}{1+s}\right\}$ .

 $PoS(G) = \frac{\text{makespan of the best NE schedule}}{\text{min makespan (social optimum)}}$ 

Theorem: For every  $s \le 1$ , there exists a game with  $s_1 = 1$ ,  $s_2 = s$ , and  $PoS(G) = min\left\{1 + s, 1 + \frac{1}{1+s}\right\}$ .

Proof: case 1: 
$$s \le \frac{\sqrt{5}+1}{2}$$
.  
Let J={x,y},  $p_x = 1$ ,  $p_y = \frac{1}{s}$   $\pi_1 = \pi_2 = (x, y)$ .

$$1 + s = 1 + \frac{1}{1+s}$$
  
for  $s = \frac{\sqrt{5}-1}{2}$ 



(\*) if 
$$s = \frac{\sqrt{5}-1}{2}$$
, take  $p_y = \frac{1}{s} - \epsilon$ )

$$PoS = 1+s$$

case 2: s > 
$$\frac{\sqrt{5}-1}{2}$$
.  
J={x,y,z}, p<sub>x</sub> = 1, p<sub>y</sub> = s<sup>2</sup>+s-1, p<sub>z</sub> = 1+s.  
 $\pi_1 = \pi_2 = (x, y, z)$ .

In all NE: (1) x is on the fast machine (2) y is on the slow machine since  $s^2 + s > (s^2+s-1)/s$ . (3) z is indifferent.  $p_x + p_z = (p_y + p_z)/s = 2+s$ .



# Equilibrium inefficiency, Identical machines

Theorem:

If  $s_i = 1$  for all  $i \in M$ , then  $PoA(G) \le 2 - \frac{1}{m}$ 

Proof:

We show that any NE is a possible outcome of Graham's List-scheduling algorithm

#### Theorem:

If  $s_i = 1$  for all  $i \in M$ , then it is NP-hard to approximate the best NE within a factor of  $2 - \frac{1}{m} - \epsilon$  for all  $\epsilon > 0$ . Proof: Reduction from 3D-matching.

### Back to centralized setting (not a game)

- A set J of n jobs, and a set M of m parallel machines
  - Every job  $j \in J$  has processing time  $p_i$
  - In case of unrelated machines,  $p_{ij}$  is the processing time of job j on machine i.
  - Every machine  $i \in M$  has a priority list  $\pi_i: J \rightarrow \{1,...,n\}$ , defining its scheduling policy.

The Goal: Find a schedule that minimizes  $\sum_{i} C_{i}$ 

Note: In the centralized setting, priority lists do not `upgrade' the problem of minimizing the Makespan

Without priority lists, both problems are solvable P $|\sum_{j} C_{j}$  - SPT is optimal [Smith 1956] R $|\sum_{j} C_{j}$  - can be represented as a bipartite weighted matching problem [Bruno, Coffman, Sethi 1974]

Theorem:  $P|\pi|\sum_j C_j$  is APX-hard  $(\mathbf{z})$ 

We therefore consider several restricted classes:

- Global priority list
- Fixed number of machines
- Fixed number of priority classes

#### Our results:

	$\pi_i$	$\pi_{global}$	$\pi_{LPT}$	$\pi_{i,c}$	$\pi_{global,c}$
Ρ	APX-hard	QPTAS	Р	APX-hard	Р
R	APX-hard	APX-hard	APX-hard	APX-hard	APX-hard

 $\pi_{i,c}$  and  $\pi_{global,c}$ : the jobs are partitioned into c job classes  $J_1, ..., J_c$ . For every  $1 < k \le c$ , every machine processes jobs from  $J_k$  after it processes jobs from  $\bigcup_{j < k} J_j$ . Note: in every optimal schedule, for every  $1 \le i \le m$  and  $1 \le k \le c$ , machine i processes jobs of  $J_k$  in SPT order.

#### Our results:

	$\pi_i$	$\pi_{global}$	$\pi_{LPT}$	$\pi_{i,c}$	$\pi_{global,c}$
Ρ	APX-hard	QPTAS	Р	APX-hard	Р
R	APX-hard	APX-hard	APX-hard	APX-hard	APX-hard
			1		
In P if m is a constant					

A Useful Observation: Let  $l_i$  denote the number of jobs on machine i.

The job with the k-th highest priority assigned to machine

*i* contributes exactly  $l_i + 1 - k$  times its processing time

to the sum of completion times.

the delay-coefficient of the job



 $p_{ij}$  is counted  $l_i + 1 - (l_i - 2) = 3$  times in  $\Sigma_j C_j$ 

#### $\pi_{\rm LPT}$ - Longest Processing Time First

 $\pi_{LPT}$  – every machine processes jobs in LPT order.

 $P | \pi_{LPT} | \sum_j C_j :$ A global priority list  $\pi = (1, 2, ..., n)$ , where  $p_1 \ge p_2 \ge ... \ge p_n$ .

 $R | \pi_{LPT} | \sum_{j} C_{j}$ : For machine i,  $\pi_{i} = (1_{i}, 2_{i}, ..., n_{i})$ , where  $p_{i,1^{i}} \ge p_{i,2^{i}} \ge ... \ge p_{i,n^{i}}$ .

# An optimal algorithm for $P|\pi_{LPT}|\sum_j C_j$

Claim: There exists an optimal schedule for  $P|\pi_{LPT}|\sum_j C_j$ in which for some  $l_1 \leq l_2 \leq \cdots \leq l_m$  such that  $\sum_i l_i = n$ , it holds that machine *i* processes the consequent subsequence of  $l_i$  jobs  $1 + \sum_{k \leq i} l_k, \dots, \sum_{k \leq i} l_k$ .

#### Illustration of the claim:



Assume m=3, then some optimal schedule looks like this:



# An optimal algorithm for $P|\pi_{LPT}|\sum_j C_j$

Claim: There exists an optimal schedule for  $P|\pi_{LPT}|\sum_j C_j$ in which for some  $l_1 \leq l_2 \leq \cdots \leq l_m$  such that  $\sum_i l_i = n$ , it holds that machine *i* processes the consequent subsequence of  $l_i$  jobs  $1 + \sum_{k \leq i} l_k, \dots, \sum_{k \leq i} l_k$ .

#### Illustration of the claim:



Assume m=3, then some optimal schedule looks like this:





# An optimal algorithm for $P|\pi_{LPT}|\sum_{j} C_{j}$

Proof: (for two machines) Assume that we know how many jobs are assigned to each of the machines. W.l.o.g., assume that  $l_1 \leq l_2$ . We show that in some optimal schedule,  $M_1$  processes the  $l_1$  longest jobs, and  $M_2$  processes the  $l_2$  shortest jobs.



# An optimal algorithm for $P|\pi_{LPT}|\sum_{j} C_{j}$

Consider the *i* -th job on machine 2. This job gets a coefficient of  $l_2 + 1 - i$ . The shortest possible job that can get this coefficient is job  $l_1 + i$ . Consider a job  $i \le l_1$ . The minimal coefficient job *i* can get is  $l_1 + 1 - i$  (for example, in every schedule, the longest job, has coefficient at least  $l_1$ ).



# An optimal algorithm for $P|\pi_{LPT}|\sum_{j} C_{j}$

When jobs  $j = 1, ..., l_1$  are on  $M_1$  and jobs  $j = l_1 + 1, ..., l_2$  are on  $M_{2}$  every coefficient (on  $M_2$ ) is matched with the shortest job that can get this coefficient, and every job (on  $M_1$ ) is matched with the minimal coefficient it can get.



# An optimal algorithm for $P|\pi_{LPT}|\sum_j C_j$

**Theorem:** P |  $\pi_{LPT}$  |  $\sum_{i} C_{i}$  is polynomial time solvable.

Proof: A dynamic programming based on the above claim



#### On the other hand:

With unrelated machines, the problem is hard and hard to approximate:

**Theorem:**  $\mathbb{R} | \pi_{LPT} | \sum_j C_j$  is APX-hard

# $\mathbb{R}|\pi_{LPT}|\sum_{j} C_{j}$ is APX-hard

Theorem:  $\mathbb{R}|\pi_{LPT}|\sum_{j} C_{j}$  is APX-hard Proof: (for now, NP-hardness only) Reduction from vertex-cover Given a graph G and an integer k, does G have a VC of size k?



# $\mathbb{R}[\pi_{LPT}] \sum_{i} C_{i}$ is APX-hard

Given G=(V,E) and k, construct an instance for R  $|\pi_{LPT}| \sum_i C_i$ : |V| machines, where M<sub>i</sub> corresponds to node  $i \in V$ . The set of jobs consists of two sets D and A. D includes |V|-k dummy jobs.  $\forall i, d, p_{i,d} = 1$ A includes |E| jobs, each corresponding to an edge  $e \in E$ .

 $p_{i,(u,v)} = \begin{cases} 0 & i = u \text{ or } i = v \\ 1 & otherwise \end{cases} \quad i \text{ is an endpoint of } (u,v)$ 

 $M=\{a,b,c,d,e\}$  $J = D \cup A$  $D = \{d_1, d_2, d_3\}$ A={(ab),(ad),(ae),...}



# $\mathbb{R}|\pi_{LPT}|\sum_{j} C_{j}$ is APX-hard

 $\pi_{LPT}$  implies that if a job (edge) is assigned on a machine corresponding to one of its endpoint then it is processed after any dummy job assigned to this machine.



Every dummy job goes to a different machine. All A-jobs have C<sub>j</sub>=0.

# $\mathbb{R}|\pi_{LPT}|\sum_{j} C_{j}$ is APX-hard

Hardness proof for APX-hardness a bit more technical. The reduction is from Max-k-VC of a bounded degree graph.

Given G, k, where max-degree(G) =  $\Delta$ , find U  $\subseteq$  V, |U|=k, such that the number of edges adjacent to vertices in U is maximal.

### Summary and open problems





- The introduction of machine-dependent priority lists opens a new world of optimization problems.
- Challenging analysis as a game as well as an optimization problem.
- General instances: no guaranteed NE, hard to approx.
- Some important classes behave nicely.

# Summary and open problems



#### To do list:

- Complexity status of  $P|\pi| \sum_{j} C_{j}$  (QPTAS but no hardness proof)
- Identify additional tractable/stable instances
- Approximation algorithms
- Priority-list can be viewed as a special case of machines-based precedence constraints (precedence constraints given by a chain). Study the general P|machine-based prec $|\sum_{i} C_{i}$
- Analyze instances with due-dates and lateness-related obj.

Assume a global priority list. What is the minimal number of machines required to complete all jobs on time?

#### **Questions?**



© marketoonist.com