A simple Graph-Theoretic Model for Selfish Restricted Scheduling

R. Elsässer[1], M. Gairing[1], T. Lücking[1], M. Mavronicolas[2], B. Monien[1]

[1] University of Paderborn, Germany.
[2] University of Cyprus, Cyprus.

Motivation and the Problem studied

Problem:

- m non-cooperative users
- n processing machines
- **task:** assign an unsplittable *unit* job to each user.
- **Objective:** stable assignment of users jobs
  $\Rightarrow$ modelled as a *Nash Equilibrium*
- Users & Machines Interaction exploits locality: Each user has access to only *two* machines.

Representation: *interaction Graph*

vertices $\longleftrightarrow$ machines
edges $\longleftrightarrow$ users
Any assignment of users corresponds to an *orientation* of the graph.

Framework

- Pure Nash Equilibria (pure NE): each user assigns its load exactly to *one* of its pair of machines.

- Mixed Nash Equilibrium (Mixed NE): Probability distribution on the pair of machines.
  In a mixed NE, the Social Cost (SC) = expected *makespan*= max of total load over all machines. $\Longrightarrow$

  best mixed NE = min makespan

  worst mixed NE = max makespan

Summary of Results

3-regular interaction Graphs:

- SC of a fully mixed NE of any d-regular graph is $d - f(d,n)$, where asymptotically tends to zero.
- *Standard fully mixed NE*: all probabilities assignments are $1/2$. $\Longleftrightarrow$ The best 3-regular interaction graph for this case is the *3-regular parallel links* graph.

Bound on the Coordination Ratio:

- For the more general case of restricted parallel links, a tight bound of $\Theta(\frac{\log n}{\log\log n})$ is known for pure NE [M. Gairing et all, STOC' 04]
  $\Longrightarrow O(\frac{\log n}{\log\log n})$ for our model.
  – We construct an $\Omega(\frac{\log n}{\log\log n})$ interaction graph with this ratio, thus the bound is *tight* for our model.

Summary of Results (Cont.)

Fully Mixed NE:

- There exists counterexample interaction graphs for which fully mixed Nash Equilibria many not exist.

Let *a fully mixed Nash dimension* = the dimension d of the smallest d-dimensional space that can contain all fully mixed NE.

- *Complete bipartite* graphs, we prove a dichotomy theorem characterizing unique existence.
  Hybercubes, we prove that fully mixed Nash dimension is the hybercube dimension for hybercubes of dimension 2 or 3.

Related Work

- Our model of interaction graphs is a special case of restricted parallel links introduced in [M. Gairing et al. MFCS04].
- [Awerbuch et al, WAOA04]: Coordination ratio for the model of restricted parallel links is $\Theta(\frac{\log n}{\log \log \log n})$ (tight), for all mixed NE. This implies the same bound for our model.
- The model of restricted parallel links is a generalization of the KP-model for selfish routing of [Koutsoupias, Papadimitriou, STACS'99].

Definitions

Let $[k] = \{1, \cdots, \}$, $k \geq 1$.

- interaction Graphs: $G(V, E)$. edges$\longleftrightarrow$users, vertices$\longleftrightarrow$machines. Assume m users, n machines.
  $\Rightarrow$ An edge connects two vertices if and only if the user can place his job onto the two machines.

- Strategies and Assignments: *Pure Assignment*: each users plays only one strategy.
  Pure assignment $L = \langle l_1, \cdots, l_m \rangle$.
  Mixed strategy: probability distribution over strategies.
  Mixed assignment $P = (p_{ij})_{i \in [n], j \in [m]}$.
  Fully mixed assignment F: all probabilities are strictly positive.
  Standard Fully mixed assignment $\widetilde{F}$: all probabilities are $1/2$.
  Fully mixed Nash dimension of a graph G = the dimension d of the smallest d-dimensional space that can contain all fully mixed NE of G.

Definitions (Cont.): Cost measures

- In a pure assignment L,
  load of a machine j, λ_j is the number of users assigned to j.
  Individual cost of user i is $\lambda_i = |k :_k = l_i|$, the load of the machine it chooses.

- Mixed assignment L, the *expected load* of a machine j, is the expected number of users assigned to j.
  Expected individual cost of user i on machine j is $\lambda_{ij} = 1 + \Sigma_{k\in[m],k\neq i}\, p_{kj}$.
  The *Expected Individual Cost* for user $i \in [m]$, is $\lambda_i\, \Sigma_{j\in[n]p_{ij}\lambda_{ij}}$.

- *Social Cost* in a mixed assignment P, $SC(G, P)$, is the maximum load over all machines of G.
  The *optimum* $OPT(G)$ is the least possible social cost over all pure assignments.

- *Coordination Ratio*, CR_G is the maximum over all NE P of the ratio $\frac{SC(G,P)}{OPT(G)}$. CR is the maximum CR_G over all graphs G.

Definitions (Cont.): *Graph Orientations*

- C_r: a cycle of r vertices, $K_{r,s}$: bipartite graph, H_r: hybercube of dimension r, *necklace* is a graph conststing of 2 vertices and 3 paralel edges, $G_{||}(n)$ are the *parallel links* graph, i.e. the graph consisting of $n/2$ necklaces.
- An *orientation* of G: directions of its edges.
  The *makespan* of a vertex in an orientation α (makespan of an orientation) is the (maximum) in-degree of it (of all vertices) in α.
  d-orientation is an orientation with makespan d in the graph G.

3-Regular Graphs: Rough Estimation

Consider a standard fully mixed NE, $\widetilde{F}$. Let $q_d(G)$ the probability such a random orientation has makespan at most $d-1$.

Lemma 1. Let I an independent set of G. Then $q_d(G) \leq (1 - \frac{1}{2^d})^{|I|}$.

Theorem 1. For a d-regular graph G with n vertices,
$SC(\widetilde{F}, G) = d - f(d,n), \quad f(d,n) \to 0$ as $n \to \infty$.

Proof. Every maximal ind. set $I = \lceil \frac{n}{d+1} \rceil$. By Lemma 1,
$\Rightarrow q_d(G) \leq (1 - \frac{1}{2^d})^{\frac{n}{d+1}}$. Thus, $SC(\widetilde{F}, G) \geq q_d(G) + d(1 - q_d(G)) = d - f(n,d)$, where $f(n,d)$ asymptotically tends to zero.

□

3-Regular Graphs: Catroids and the Two-Sisters Lemma

- **Definition 1.(Cactoids)** A cactoid is a pair $\widehat{G} = \langle V, \widehat{E} \rangle$, V is the vertices, $\widehat{E}$ consists of undirected edges between vertices and pointers to vertices, i.e. loose edges incident to one single vertex.
- Let an arbitrary orientation of G, σ, called *standard*.
  $x_\alpha(e) \rightarrow \{0, 1\}$ for each $e \in \widehat{E}$ in any possible orientation α: is 1 (or 0) if e and the same orientation in α as in σ (otherwise).
  Assume two vertices u, v, called *two-sisters*, with incident pointers π_u, π_v, pointing away of u, v in σ.
  Let $P_{\widehat{G}}(i, j)$ the probability that an α with $x_\alpha(u) = i$ and $x_\alpha(v) = j$, $i, j \in \{0, 1\}$ is a 2-*orientation*.
- Clearly, $P_{\widehat{G}}(1, 1) \geq P_{\widehat{G}}(0, 0), P_{\widehat{G}}(0, 1), P_{\widehat{G}}(1, 0)$.
- We prove that $P_{\widehat{G}}(1, 1)$ is upper bounded by their sum..

3-Regular Graphs: The Two-Sisters Lemma

Lemma 2. (Two-sisters) For any 3-regular cactoid $\widehat{G} = \langle V, \widehat{E} \rangle$ and any two sisters $u, v \in V$,
it holds that, $P_{\widehat{G}}(0,0) + P_{\widehat{G}}(0,1) + P_{\widehat{G}}(1,0) \geq P_{\widehat{G}}(1,1)$.

Proof. Let b_1, b_2 and b_3, b_4, the other edges incident to the sisters u, v, respectively.
Let $\widehat{G'}$ obtained by $\widehat{G}$ by deleting u, v and their pointers π_u, π_v.
Let $P_{\widehat{G'}}(x_1, x_2, x_3, x_4)$ the probability a random orientation α of $\widehat{G'}$ with $x_\alpha(b_i) = x_i$, $1 \leq i \leq 4$ is a 2-orientation.

1. We express $P_{\widehat{G}}(i,j)$ $i, j \in \{0,1\}$ as functions of $P_{\widehat{G'}}(x_1, x_2, x_3, x_4)$.
2. By, induction on the number of vertices of $\widehat{G}$, we prove that, the statement holds for $\widehat{G'}$.
3. Using 1. , we return to $\widehat{G}$ and get the same statement.

$\square$

3-Regular Graphs: Orientations and Social Costs

Theorem 2. For every 3-regular graph G, with n vertices it holds that $|3\text{–}\mathit{or}(G)| \geq |3\text{–}\mathit{or}(G_{||}(n))|$, where $\mathit{or}(H)$ is the number of orientations of a graph H.

Proof.

- We start from the graph $G_0 = G = (V, E_0)$ and iteratively define $G_i = (V, E_i)$, $1 \leq i \leq r$, $r \leq n$ s.t.
  G_r equals $G_{||}(n)$ and $|3\text{–}\mathit{or}(G_i)| \geq |3\text{–}\mathit{or}(G_{||}(G_{i+1}))|$.
- Note: Each connected component of any regular graph, is either isomorphic to a necklace or
  it contains a path of length 3 connecting four different vertices, such that only the middle edge of this path can be a parallel edge.
- If in G_i all connected components are necklaces, then $G_i = G_{||}(n)$.

Proof of Theorem 2. (Cont. 2/5)

- Otherwise, some component of G_i contains a path c, a, b, d with 4 different vertices a, b, c, d.
  Construct a new graph $G_{i+1} = (V, E_{i+1})$ by deleting edges $\{a, c\}$, $\{b, d\}$ from E_i and adding edges $\{a, b\}$, $\{c, d\}$ to the graph as follows:

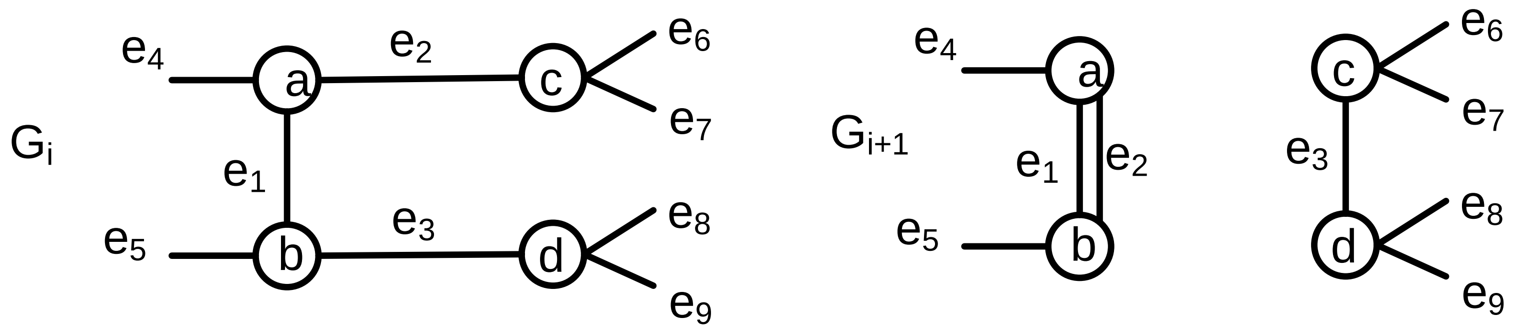

Figure 1: Construction of graph G_{i+1} from graph G_i.

- In the figure, all edges are different. This is *not always* the case.
- At each iteration, the number of single edges is decreased by at least one. Thus, # of iterations is at most n.

Proof of Theorem 2. (Cont. 3/5)

We prove the statement when,
Case 1: All edges $e_1, \cdots e_9$ are different.
Case 2: Some of the edges are equal.

Here we present only the **Case 1**:

- Consider the graphs G_1, G_2. There exists an one-to-one correspondence between their edges. Thus, an orientation of $G_1 \Leftrightarrow$ an orientation of G_2.
- We define an injective mapping $F : 3\text{–}or(G_2) \to 3\text{–}or(G_1)$
  Set $C_2 = \{\alpha; \alpha \in 3\text{–}or(G_2),\ \alpha; \notin 3\text{–}or(G_1)\}$ and
  $C_1 = \{\alpha; \alpha \in 3\text{–}or(G_1),\ \alpha; \notin 3\text{–}or(G_2)\}$.
  Define $F(\alpha) = \alpha$ for $\alpha \in 3\text{–}or(G_2)\backslash G_2$ and $F : C_2 \to C_1$ is injective.
  Thus, the mapping F is *injective*.
- We will show that F always exists in **Case 1**..

Proof of Theorem 2: Case 1 (Cont. 4/5)

- Let α an arbitrary orientation. All $u \notin \{a, b, c, d\}$ have the makespan in G_1 and G_2 with respect to α.
- We can show that vertices a, b, c, d have all makespan 3 in G_1.
- Using above info, we construct C_2:

$$C_2 = \{\alpha \notin 3\text{–}(G_1); x_1 = x_2 = x_3 = 0 \wedge x_5 = 1 \wedge x_6 \cdot x_7 = x_8 \cdot x_9 = 0\}$$
$$\cup\{\alpha \notin 3\text{–}(G_1); x_2 = x_3 = x_6 = x_7 = 1 \wedge x_1 \cdot x_4 = 0 \wedge (x_1 = 1 \vee x_5 = 0)\}$$

- Similarly, we construct C_1:

$$C_1 = \{\alpha \notin 3\text{–}(G_2); x_1 = 0 \wedge x_2 = x_3 = x_5 = 1 \wedge x_6 \cdot x_7 = 0\}$$
$$\cup\{\alpha \notin 3\text{–}(G_2); x_2 = x_3 = 0 \wedge x_6 = x_7 = 1 \wedge x_8 \cdot x_9 = 0 \wedge ((x_1 = 1 \vee x_5 = 0)\}$$

Proof of Theorem 2: Case 1 (Cont. 5/5)

We define F by considering four cases about orientations $\alpha \in C_2$:

1. Consider $\alpha \in C_2$ with $x_2 = x_3 = x_6 = x_7 = 1 \wedge x_1 \cdot x_4 = 0 \wedge x_8 \cdot x_9 = 0 \wedge (x_1 = 1 \vee x_5 = 0)$

   Set $F(x_1, 1, 1, x_4, x_5, 1, 1, x_8, x_9, \ldots) = (x_1, 0, 0, x_4, x_5, 1, 1, x_8, x_9, \ldots))$

   Note: vertices $\{a, b, c, d\}$ have the same connections to vertices outside $\{a, b, c, d\}$; therefore $\alpha \notin 3\text{–}or(G_1)$, thus $F(\alpha) \notin 3\text{–}or(G_2)$.
   Thus, $F(\alpha) \in C_1$.

   2-4. More complicated... prove the same result.

$\square$

Theorem 2 consequences

Corollary 1. For an 3-regular graph G with n vertices, $SC(G, \widetilde{F}) \geq SC(G_{||}(n), \widetilde{F}) = 3 - (3/4)^{n/2}$.

- Equality does not hold in Corollary 1: there exist a 3-regular graph for which the SC of its fully mixed NE is larger than for the corresponding parallel links graph.

Coordination Ratio

Theorem 3. Restricted to pure NE, $CR = \Theta\left(\frac{\log n}{\log \log n}\right)$.

Proof. *Upper bound*: Our model is a special case of the restricted parallel links. $\Rightarrow$ The upper bound $O(\frac{\log n}{\log \log n})$ of [M. Gairing et all, MFCS04] also holds for our model.

Lower bound: Let G a complete tree with height k, where each vertex in layer l of the tree has $k - l$ children.
Let $k^{\underline{l}} = k(k-1) \ddots \ldots \cdot^{\cdot^{\cdot}} (k-1)$ the l-th *falling factorial* of k. Then $n = \Sigma_{0 \leq l \leq k} k^{\underline{l}} < (k+1)! = \Gamma(k+2)$. This implies $k > \Gamma^{-1}(n) - 2$.

1. Denote L_1 the pure assignment in which all users are assigned toward the root.
   Then the individual cost of user in layer l is $k - l$. Also, the user can not improve by moving its vertex in layer $(l+1)$.
   Thus, L_1 is a pure NE with Social Cost k.

Theorem 3 proof. (Cont.)

2. Denote L_2 the pure assignment in which all users are assigned toward the leaves.
Then the individual cost of all users is 1.
Thus, the Social Cost of L_2 is 1.

$$\Rightarrow \max_{G,L} \frac{SC(G,L)}{OPT(G)} \geq \frac{SC(G,L_1)}{SC(G,L_2)} = k > \Gamma^1(n) - 2 = \Omega(\frac{\log n}{\log \log n}).$$

$\square$

The fully Mixed Nash Equilibrium

Consider a fully Mixed NE, P. For each edge $jk \in E$, let jk the user corresponding to the edge jk.
Denote $\widehat{p}_{jk}$ and $\widehat{p}_{kj}$ the probabilities according to P that user jk chooses machines j and k, resp.
For each machine $j \in V$, the expected load of machine j excluding a set of edges $\widetilde{E}$, denoted by $\pi_P \backslash \widetilde{E} = \Sigma_{kj \in E \backslash \widetilde{E}} \widehat{p}_{kj}$.

Lemma 3. (The 4-Cycle Lemma) Take any 4-cycle C_4 in a graph G and any two vertices $u, v \in C_4$ that are non-adjacent in C_4. Consider a NE P for G. Then, $\pi_P(u) \backslash C_4 = \pi_P(v) \backslash C_4$.

Counterexample 1. There is no fully mixed NE for trees and meshes.

Counterexample 2. For each graph in Figure 1, there is no fully mixed NE.

Fully mixed NE: Uniqueness and Dimensional Results

Theorem 4. Consider the complete bipartite graph $K_{r,s}$, where $s \geq r \geq 2$ and $s \geq 3$. Then the fully mixed NE F for $K_{r,s}$ exists uniquely if and only if $r > 2$. Moreover, in case $r = 2$, the fully mixed Nash dimension of $K_{r,s}$ is $s - 1$.

Observation 2. Consider a hybercube H_r, for any $r \geq 2$. Then, the fully mixed Nash dimension of H_r is at least r.

Theorem 4. Consider the hybercube H_r, where $r \in \{2, 3\}$ Then the fully mixed Nash dimension is r.

Worst-Case NE

Counterexample 3. There is an interaction graph for which no fully mixed NE has worst Social Cost.

Counterexample 4. There is an interaction graph for which there exists a fully mixed NE with worst Social Cost.