

A Simple Graph-Theoretic Model for Selfish Restricted Scheduling [★]

Robert Elsässer¹, Martin Gairing¹, Thomas Lücking¹, Marios Mavronicolas², and Burkhard Monien¹

¹ Faculty of Computer Science, Electrical Engineering and Mathematics, University of Paderborn, Fürstenallee 11, 33102 Paderborn, Germany. Email: {elsa, gairing, luck, bm}@uni-paderborn.de

² Department of Computer Science, University of Cyprus, P. O. Box 20537, Nicosia CY-1678, Cyprus. Email: mavronic@ucy.ac.cy

Abstract. In this work, we introduce and study a simple, graph-theoretic model for selfish *scheduling* among m non-cooperative *users* over a collection of n *machines*; however, each user is restricted to assign its unsplitable *load* to one from a pair of machines that are allowed for the user. We model these bounded interactions using an *interaction graph*, whose vertices and edges are the machines and the users, respectively. We study the impact of our modeling assumptions on the properties of Nash equilibria in this new model. The main findings of our study are outlined as follows:

- We prove, as our main result, that the *parallel links* graph is the *best-case* interaction graph – the one that minimizes expected *makespan* of the *standard fully mixed Nash equilibrium* – among all *3-regular* interaction graphs. The proof employs a graph-theoretic lemma about *orientations* in 3-regular graphs, which may be of independent interest.
- We prove a lower bound on *Coordination Ratio* [16] – a measure of the cost incurred to the system due to the selfish behavior of the users. In particular, we prove that there is an interaction graph incurring Coordination Ratio $\Omega\left(\frac{\log n}{\log \log n}\right)$. This bound is shown for pure Nash equilibria.
- We present counterexample interaction graphs to prove that a *fully mixed Nash equilibrium* may sometimes not exist at all. Moreover, we prove properties of the fully mixed Nash equilibrium for *complete bipartite* graphs and *hypercube* graphs.

1 Introduction

Motivation and Framework. Consider a group of m non-cooperative *users*, each wishing to assign its unsplitable unit *job* onto a collection of n processing (identical) *machines*. The users seek to arrive at a stable assignment of their jobs for their joint interaction. As usual, such stable assignments are modeled as *Nash equilibria* [21], where no user can unilaterally improve its objective by switching to a different strategy.

[★] This work has been partially supported by the *DFG-Sonderforschungsbereich 376 Massive Parallelität: Algorithmen, Entwurfsmethoden, Anwendungen*, by the VEGA grant No. 2/3164/23, by the *IST* Program of the European Union under projects *FLAGS* (contract number IST-2001-33116) and *DELIS* (contract number 001907), and by research funds at University of Cyprus.

We use a structured and sparse representation of the relation between the users and the machines that exploits the locality of their interaction; such locality almost always exists in complex scheduling systems. More specifically, we assume that each user has access (that is, finite cost) to only *two* machines; its cost on other machines is infinitely large, giving it no incentive to switch there. The (expected) cost of a user is the (expected) load of the machine it chooses. Interaction with just a few neighbors is a basic design principle to guarantee efficient use of resources in a distributed system. Restricting the number of interacting neighbors to just two is then a natural starting point for the theoretical study of the impact of selfish behavior in a distributed system with local interactions.

Our representation is based on the *interaction graph*, whose vertices and (undirected) edges represent the machines and the users, respectively. Multiple edges are allowed; however, for simplicity, our interaction multigraphs will be called *interaction graphs*. The model of interaction graphs is interesting because it is the simplest, non-trivial model for selfish scheduling on restricted parallel links. In this model, any assignment of users to machines naturally corresponds to an *orientation* of the interaction graph. (Each edge is directed to the machine where the user is assigned.)

We will consider *pure Nash equilibria*, where each user assigns its load to exactly one of its two allowed machines with probability one; we will also consider *mixed Nash equilibria*, where each user employs a probability distribution to choose between its two allowed machines. Of particular interest to us is the *fully mixed Nash equilibrium* [20] where every user has strictly positive probability to choose each of its two machines. In the *standard fully mixed Nash equilibrium*, all probabilities are equal to $\frac{1}{2}$. It is easy to see that the standard fully mixed Nash equilibrium exists if and only if the (multi)graph is regular.

With each (mixed) Nash equilibrium, we associate a *Social Cost* [16] which is the expected *makespan* - the expectation of the maximum, over all machines, total load on the machine. *Best-case* and *worst-case* Nash equilibria minimize and maximize Social Cost, respectively. For a given type of Nash equilibrium such as the standard fully mixed Nash equilibrium, best-case and worst-case graphs among a graph class minimize and maximize Social Cost of Nash equilibria of the given type, respectively. The assignment of users to machines that minimizes Social Cost might not necessarily be a Nash equilibrium; call *Optimum* this least possible Social Cost. We will investigate *Coordination Ratio* [16] - the worst-case ratio over all Nash equilibria, of Social Cost over Optimum. We are interested in understanding the interplay between the topology of the underlying interaction graph and the various existence, algorithmic, combinatorial, structural and optimality properties of Nash equilibria in this new model of selfish restricted scheduling with bounded interaction.

Contribution and Significance. We partition our results into three major groups: 3-regular interaction graphs (Section 3). It is easy to prove that the Social Cost of the standard fully mixed Nash equilibrium for any d -regular graph is $d - f(d, n)$, where $f(d, n)$ is a function that goes to 0 as n goes to infinity. This gives a general but rather rough estimation of Social Cost for d -regular graphs; moreover, it does not say how the specific structure of each particular 3-regular graph affects the Social Cost of the standard fully mixed Nash equilibrium. We continue to prove much sharper estimations

for the special class of 3-regular graphs. Restricting our model of interaction graphs to 3-regular graphs led us to discover some nice structural properties of orientations in 3-regular graphs, which were motivated by Nash equilibria. However, we have so far been unable to generalize these properties to regular graphs of degree higher than 3.

We pursue a thorough study of 3-regular interaction graphs; these graphs further restrict the bounded interaction by insisting that each machine is accessible to just three users. Specifically, we focus on the standard fully mixed Nash equilibrium where all probabilities of assigning users to machines are $\frac{1}{2}$. We ask which the best 3-regular interaction graph is in this case. This question brings into context the problem of comparing against each other the expected number of 2-orientations and 3-orientations - those with makespan 2 and 3, respectively. The manner in which these numbers outweigh each other brings Social Cost closer to either 2 or 3. We develop some deep graph-theoretic lemmas about 2- and 3-orientations in 3-regular graphs to prove, as our main result, that the simplest 3-regular parallel links graph is the best-case 3-regular graph in this setting. The proof decomposes any 3-regular graph down to the parallel links graph in a way that Social Cost of the standard fully mixed Nash equilibrium does not increase. The graph theoretic lemmas about 2- and 3-orientations are proved using both counting and mapping techniques; both the lemmas and their proof techniques are, we believe, of more general interest and applicability.

Bound on Coordination Ratio (Section 4). For the more general model of restricted parallel links, a tight bound of $\Theta(\frac{\log n}{\log \log n})$ on Coordination ratio restricted to pure Nash equilibria was shown in [9, Theorem 5.2] and independently in [1, Theorem 1]. This implies an upper bound of $O(\frac{\log n}{\log \log n})$ on the Coordination Ratio for pure Nash equilibria in our model as well. We construct an interaction graph incurring Coordination Ratio $\Omega(\frac{\log n}{\log \log n})$ to prove that this bound is tight for the model of interaction graphs as well. The construction extends an approach followed in [9, Lemma 5.1] that proved the same lower bound for the more general model of restricted parallel links.

The Fully Mixed Nash Equilibrium (Section 5). We pursue a thorough study of fully mixed Nash equilibria across interaction graphs. Our findings are outlined as follows:

- There exist counterexample interaction graphs for which fully mixed Nash equilibria may not exist. Among them are all trees and meshes. These counterexamples provide some insight about a possible graph-theoretic characterization of interaction graphs admitting a fully mixed Nash equilibrium. 4-cycles and 1-connectivity are factors expected to play a role in this characterization.
- We next consider the case where infinitely many fully mixed Nash equilibria may exist. In this case, the fully mixed Nash dimension is defined to be the dimension d of the smallest d -dimensional space that can contain all fully mixed Nash equilibria. For complete bipartite graphs, we prove a dichotomy theorem that characterizes unique existence. The proof employs arguments from Linear Algebra. For hypercubes, we have only been able to prove that the fully mixed Nash dimension is the hypercube dimension for hypercubes of dimension 2 or 3. We conjecture that this is true for all hypercubes, but we have only been able to observe that the hypercube dimension is a lower bound on the fully mixed Nash dimension (for all hypercubes).

- We are finally interested in understanding whether (or when) the fully mixed Nash equilibrium is the worst-case one in this setting. We present counterexample interaction graphs to show that the fully mixed Nash equilibrium is sometimes the worst-case Nash equilibrium, but sometimes not. For the hypercube, there is a pure Nash equilibrium that is worse (with respect to Social Cost) than the fully mixed one. On the other hand, for the 3-cycle the fully mixed Nash equilibrium has worst Social Cost.

Related Work and Comparison. Our model of interaction graphs is the special case of the model of restricted parallel links introduced and studied in [9], where each user is now further restricted to have access to only two machines. The work in [9] focused on the problem of computing pure Nash equilibria for that more general model. Awerbuch et al. [1] also considered the model of restricted parallel links, and proved a tight upper bound of $\Theta(\frac{\log n}{\log \log \log n})$ on Coordination Ratio for all (mixed) Nash equilibria. This implies a corresponding upper bound for our model of interaction graphs. It is an open problem whether this bound of $O(\frac{\log n}{\log \log \log n})$ is tight for the model of interaction graphs, or whether a better upper bound on Coordination Ratio for all (mixed) Nash equilibria can be proved.

The model of restricted parallel links is, in turn, a generalization of the so called KP-model for selfish routing [16], which has been extensively studied in the last five years; see e.g. [3–5, 7–11, 19, 20]. Social Cost and Coordination Ratio were originally introduced in [16]. Bounds on Coordination Ratio are proved in [3, 8–10, 20]. The fully mixed Nash equilibrium was introduced and studied in [20], where its unique existence was proved for the original KP-model. The Fully Mixed Nash Equilibrium Conjecture, stating that the fully mixed Nash equilibrium maximizes Social Cost, was first explicitly stated in [11]. It was proved to hold for special cases of the KP-model [11, 19] and for variants of this model [9, 10]. Recently the Fully Mixed Nash Equilibrium Conjecture was disproved for the original KP-model and the case that job sizes are *non-identical* [6]. This stands in sharp contrast to the model considered in this paper where job sizes are *identical*.

The model of interaction graphs is an alternative to *graphical games* [14] studied in the Artificial Intelligence community. The essential difference is that in graphical games, users and resources are modeled as vertices and edges, respectively. The problem of computing Nash equilibria for graphical games has been studied in [13, 14, 18]. Other studied variants of graphical games include the *network games* studied in [12], *multi-agent influence diagrams* [15] and *game networks* [17].

2 Framework and Preliminaries

For all integers $k \geq 1$, denote $[k] = \{1, \dots, k\}$.

Interaction Graphs. We consider a graph $G = (V, E)$ where edges and vertices correspond to users and machines, respectively. Assume there are m users and n machines, respectively, where $m > 1$ and $n > 1$. Each user has a unit job. From here on, we shall refer to users and edges (respectively, machines and vertices) interchangeably. So, an edge connects two vertices if and only if the user can place his job onto the two machines.

Strategies and Assignments. A *pure strategy* for a user is one of the two machines it connects; so, a pure strategy represents an assignment of the user's job to a machine. A *mixed strategy* for a user is a probability distribution over its pure strategies. A *pure assignment* $\mathbf{L} = \langle \ell_1, \dots, \ell_m \rangle$ is a collection of pure strategies, one for each user. A pure assignment induces an orientation of the graph G in the natural way. A *mixed assignment* $\mathbf{P} = (p_{ij})_{i \in [n], j \in [m]}$ is a collection of mixed strategies, one for each user. A mixed assignment \mathbf{F} is *fully mixed* [20, Section 2.2] if all probabilities are strictly positive. The *standard fully mixed assignment* $\tilde{\mathbf{F}}$ is the fully mixed one where all probabilities are equal to $\frac{1}{2}$. The *fully mixed dimension* of a graph G is the dimension d of the smallest d -dimensional space that contains all fully mixed Nash equilibria for this graph.

Cost Measures. For a pure assignment \mathbf{L} , the *load* of a machine $j \in [n]$ is the number of users assigned to j . The *Individual Cost* for user $i \in [m]$ is $\lambda_i = |\{k : \ell_k = i\}|$, the load of the machine it chooses. For a mixed assignment $\mathbf{P} = (p_{ij})_{i \in [m], j \in [n]}$, the *expected load* of a machine $j \in [n]$ is the expected number of users assigned to j . The *Expected Individual Cost* for user $i \in [m]$ on machine $j \in [n]$ is the expectation, according to \mathbf{P} , of the Individual Cost for user i on machine j , then, $\lambda_{ij} = 1 + \sum_{k \in [m], k \neq i} p_{kj}$. The *Expected Individual Cost* for user $i \in [m]$ is $\lambda_i = \sum_{j \in [n]} p_{ij} \lambda_{ij}$.

Associated with a mixed assignment \mathbf{P} is the *Social Cost* defined as the expectation, according to \mathbf{P} , of makespan (that is, maximum load), that is, $\text{SC}(G, \mathbf{P}) = \mathcal{E}_{\mathbf{P}}(\max_{v \in [n]} |\{k : \ell_k = v\}|)$, that is, Social Cost is the expectation, according to \mathbf{P} , of makespan (that is, maximum load). The *Optimum* $\text{OPT}(G)$ is defined as the least possible, over all pure assignments $\mathbf{L} = \langle \ell_1, \dots, \ell_n \rangle \in [n]^m$, makespan; that is, $\text{OPT}(G) = \min_{\mathbf{L} \in [n]^m} \max_{v \in [n]} |\{k : \ell_k = v\}|$.

Nash Equilibria and Coordination Ratio. We are interested in a special class of (pure or) mixed assignments called *Nash equilibria* [21] that we describe here. The mixed assignment \mathbf{P} is a Nash equilibrium [9, 16] if for each user $i \in [m]$, it minimizes $\lambda_i(\mathbf{P})$ over all mixed assignments that differ from \mathbf{P} only with respect to the mixed strategy of user i . Thus, in a Nash equilibrium, there is no incentive for a user to unilaterally deviate from its own mixed strategy in order to decrease its Expected Individual Cost. Clearly, this implies that $\lambda_{ij} = \lambda_i$ if $p_{ij} > 0$ whereas $\lambda_{ij} \geq \lambda_i$ otherwise. We refer to these conditions as *Nash equations* and *Nash inequalities*, respectively.

The *Coordination Ratio* CR_G for a graph G is the maximum, over all Nash equilibria \mathbf{P} , of the ratio $\frac{\text{SC}(G, \mathbf{P})}{\text{OPT}(G)}$; thus, $\text{CR}_G = \max_{\mathbf{P}} \frac{\text{SC}(G, \mathbf{P})}{\text{OPT}(G)}$. The *Coordination Ratio* CR is the maximum, over all graphs G and Nash equilibria \mathbf{P} , of the ratio $\frac{\text{SC}(G, \mathbf{P})}{\text{OPT}(G)}$; thus, $\text{CR} = \max_{G, \mathbf{P}} \frac{\text{SC}(G, \mathbf{P})}{\text{OPT}(G)}$. Our definitions for CR_G and CR extend the original definition of Coordination Ratio by Koutsoupias and Papadimitriou [16] to encompass interaction graphs.

Graphs and Orientations. Some special classes of graphs we shall consider include the cycle C_r on r vertices; the complete bipartite graph (or biclique) $K_{r,s}$ which is a simple bipartite graph with partite sets of size r and s respectively, such that two vertices are adjacent if and only if they are in different partite sets; the hypercube H_r of dimension r whose vertices are binary words of length r connected if and only if their Hamming distance is 1. For a graph G , denote Δ_G the maximum degree of G . A graph

is d -regular if all vertices have the same degree d . The graph consisting of 2 vertices and 3 parallel edges will be called *necklace*. Also, for even n , $G_{\parallel}(n)$ will denote the *parallel links* graph, i.e., the graph consisting of $\frac{n}{2}$ necklaces.

An *orientation* of an undirected graph G results when assigning directions to its edges. The *makespan* of a vertex in an orientation α is the in-degree it has in α . The *makespan* of an orientation is the maximum vertex makespan. For any integer d , a d -orientation is an orientation with makespan d in a graph G ; denote $d\text{-or}(G)$ the set of d -orientations of G .

3 3-Regular Graphs

In this section, we consider the problem of determining the best-case d -regular graph among the class of all d -regular graphs with a given number of vertices (and, therefore, with the same number of edges), with respect to the Social Cost of the standard fully mixed Nash equilibrium, where all probabilities are equal to $1/2$.

A Rough Estimation. We start with a rough estimation of the Social Cost of any d -regular graph G , where $d \geq 2$. We first prove a technical lemma about the probability that such a random orientation has makespan at most $d - 1$. Denote this probability $q_d(G)$.

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

We are now ready to prove:

Theorem 1. *For a d -regular graph G with n vertices, $\text{SC}(\tilde{\mathbf{F}}, G) = d - f(n, d)$, where $f(n, d) \rightarrow 0$ as $n \rightarrow \infty$.*

Cactoids and the Two-Sisters Lemma. The rest of our analysis will deal with 3-regular graphs. We will be able to significantly strengthen and improve Theorem 1 for the special case of 3-regular graphs. We define a structure that we will use in our proofs.

Definition 1 (Cactoids). *A cactoid is a pair $\hat{G} = \langle V, \hat{E} \rangle$, where V is a set of vertices and \hat{E} is a set consisting of undirected edges between vertices, and pointers to vertices, i.e., loose edges incident to one single vertex.*

A cactoid is called 3-regular if each vertex is incident to three elements from \hat{E} . A cactoid may be considered as a standard multigraph if we add a special vertex and we replace each pointer by an edge which connects the special vertex with the vertex the pointer is incident to.

Consider now any arbitrary but fixed orientation σ of \hat{G} . Call it *standard orientation*. We will now define variables $x_\alpha(e)$ for each $e \in \hat{E}$, which take values from $\{0, 1\}$ in each possible orientation α of \hat{G} . The values are defined with reference to the standard orientation σ . So, take any arbitrary orientation α of \hat{G} . For each $e \in \hat{E}$, $x_\alpha(e) = 1$ if e has the same direction in α and σ , and 0 otherwise. Note that $x_\sigma(e) = 1$ for all $e \in \hat{E}$.

We now continue with a lemma that estimates the probability that a random orientation is a 2-orientation in a 3-regular cactoid \hat{G} . Consider two vertices u and v called the two sisters with incident pointers π_u and π_v . Assume that in the standard orientation

σ , π_u and π_v point away from u and v , respectively. Denote $P_{\widehat{G}}(i, j)$ the probability that a random orientation α with $x_\alpha(u) = i$ and $x_\alpha(v) = j$, where $i, j \in \{0, 1\}$, is a 2-orientation. Clearly, by our assumption on the standard orientation σ , $P_{\widehat{G}}(1, 1)$ is not smaller than each of $P_{\widehat{G}}(0, 0)$, $P_{\widehat{G}}(0, 1)$ and $P_{\widehat{G}}(1, 0)$. However, we prove that $P_{\widehat{G}}(1, 1)$ is upper bounded by their sum.

Lemma 2 (Two Sisters Lemma). *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. Denote b_1, b_2 and b_3, b_4 the other edges or pointers incident to the two sisters u and v , respectively. Define the standard orientation σ so that these edges or pointers point towards u or v , respectively. Denote \widehat{G}' the cactoid obtained from \widehat{G} by deleting the two sisters u and v and their pointers π_u and π_v . Define $P_{\widehat{G}'}(x_1, x_2, x_3, x_4)$ the probability that a random orientation α of the cactoid \widehat{G}' with $x_\alpha(b_i) = x_i$ for $1 \leq i \leq 4$ is a 2-orientation. Then,

$$\begin{aligned} P_{\widehat{G}}(1, 1) &= \frac{1}{16} \sum_{x_1, x_2, x_3, x_4 \in \{0, 1\}} P_{\widehat{G}'}(x_1, x_2, x_3, x_4), \\ P_{\widehat{G}}(0, 0) &= \frac{1}{16} \sum_{x_1 \cdot x_2 = 0, x_3 \cdot x_4 = 0} P_{\widehat{G}'}(x_1, x_2, x_3, x_4), \\ P_{\widehat{G}}(0, 1) &= \frac{1}{16} \sum_{x_1, x_2 \in \{0, 1\}, x_3 \cdot x_4 = 0} P_{\widehat{G}'}(x_1, x_2, x_3, x_4), \\ P_{\widehat{G}}(1, 0) &= \frac{1}{16} \sum_{x_1 \cdot x_2 = 0, x_3, x_4 \in \{0, 1\}} P_{\widehat{G}'}(x_1, x_2, x_3, x_4). \end{aligned}$$

Set now $D = 16 \cdot (P_{\widehat{G}}(0, 0) + P_{\widehat{G}}(0, 1) + P_{\widehat{G}}(1, 0) - P_{\widehat{G}}(1, 1))$. It suffices to prove that $D \geq 0$. Clearly,

$$D = 2 \sum_{x_1 \cdot x_2 = 0, x_3 \cdot x_4 = 0} P_{\widehat{G}'}(x_1, x_2, x_3, x_4) - P(1, 1, 1, 1).$$

Use now the cactoid \widehat{G}' to define the probabilities $Q(i, j)$ and $R(i, j)$ where $i, j \in \{0, 1\}$ as follows: $Q(i, j)$ is the probability that a random orientation α of the cactoid \widehat{G}' with $x_\alpha(b_1) = i$ and $x_\alpha(b_2) = j$ is a 2-orientation; $R(i, j)$ is the probability that a random orientation α of the cactoid \widehat{G}' with $x_\alpha(b_3) = i$ and $x_\alpha(b_4) = j$ is a 2-orientation. Clearly,

$$\begin{aligned} Q_{\widehat{G}'}(i, j) &= \sum_{x_3, x_4 \in \{0, 1\}} P_{\widehat{G}'}(i, j, x_3, x_4), \\ R_{\widehat{G}'}(i, j) &= \sum_{x_1, x_2 \in \{0, 1\}} P_{\widehat{G}'}(x_1, x_2, i, j). \end{aligned}$$

We proceed by induction on the number of vertices of \widehat{G} . So, it suffices to assume the claim for the cactoid \widehat{G}' and prove the claim for the cactoid \widehat{G} . Assume inductively

that $Q_{\widehat{G}'}(0,0) + Q_{\widehat{G}'}(0,1) + Q_{\widehat{G}'}(1,0) \geq Q_{\widehat{G}'}(1,1)$ and $R_{\widehat{G}'}(0,0) + R_{\widehat{G}'}(0,1) + R_{\widehat{G}'}(1,0) \geq R_{\widehat{G}'}(1,1)$. These inductive assumptions and the definitions of $Q_{\widehat{G}'}$ and $R_{\widehat{G}'}$ imply that

$$\begin{aligned} \sum_{\substack{x_3, x_4 \in \{0,1\} \\ x_1 \cdot x_2 = 0}} P_{\widehat{G}'}(x_1, x_2, x_3, x_4) &\geq \sum_{x_3, x_4 \in \{0,1\}} P_{\widehat{G}'}(1, 1, x_3, x_4), \\ \sum_{\substack{x_1, x_2 \in \{0,1\} \\ x_3 \cdot x_4 = 0}} P_{\widehat{G}'}(x_1, x_2, x_3, x_4) &\geq \sum_{x_1, x_2 \in \{0,1\}} P_{\widehat{G}'}(x_1, x_2, 1, 1). \end{aligned}$$

From the first inequality we obtain,

$$\sum_{\substack{x_3 \cdot x_4 = 0 \\ x_1 \cdot x_2 = 0}} P_{\widehat{G}'}(x_1, x_2, x_3, x_4) \geq \sum_{x_3, x_4 \in \{0,1\}} P_{\widehat{G}'}(1, 1, x_3, x_4) - \sum_{x_1 \cdot x_2 = 0} P_{\widehat{G}'}(x_1, x_2, 1, 1).$$

From the second inequality we get,

$$\sum_{\substack{x_1 \cdot x_2 = 0 \\ x_3 \cdot x_4 = 0}} P_{\widehat{G}'}(x_1, x_2, x_3, x_4) \geq \sum_{x_1, x_2 \in \{0,1\}} P_{\widehat{G}'}(x_1, x_2, 1, 1) - \sum_{x_3 \cdot x_4 = 0} P_{\widehat{G}'}(1, 1, x_3, x_4).$$

This yields $2 \sum_{\substack{x_1 \cdot x_2 = 0 \\ x_3 \cdot x_4 = 0}} P_{\widehat{G}'}(x_1, x_2, x_3, x_4) \geq 2P_{\widehat{G}'}(1, 1, 1, 1)$ by adding up the last two inequalities, which implies $D \geq 0$, and the claim follows. \square

Orientations and Social Cost. In this section, we prove a graph-theoretic result, namely that the regular parallel links graph minimizes the number of 3-orientations among all 3-regular graphs with the same number of vertices.

Theorem 2. *For every 3-regular graph G with n vertices it holds that $|3\text{-or}(G)| \geq |3\text{-or}(G_{\parallel}(n))|$.*

Proof. In order to prove the claim, we start from the graph $G_0 = G = (V, E_0)$ and iteratively define graphs $G_i = (V, E_i)$, $1 \leq i \leq r$, for some $r \leq n$, in a way that G_r equals $G_{\parallel}(n)$ and $|3\text{-or}(G_i)| \geq |3\text{-or}(G_{i+1})|$ holds for all $1 \leq i < r$.

Note that in each 3-regular graph, each connected component 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 is equal to $G_{\parallel}(n)$, otherwise some connected component of G_i contains a path c, a, b, d with 4 different vertices a, b, c, d . In the latter case, construct a new graph $G_{i+1} = (V, E_{i+1})$ by deleting the edges $\{a, c\}, \{b, d\}$ from E_i and adding the edges $\{a, b\}, \{c, d\}$ to the graph as described in the following paragraph.

As illustrated in Figure 1, the edges incident to vertices a, b, c, d are numbered by some j , where $1 \leq j \leq 9$. In this figure, all the edges are different. This does not necessarily have to be the case. It may happen that $e_4 = e_5$ resulting in two parallel edges between a and b in G_i and three parallel edges between a and b in G_{i+1} . It may also happen that e_6 or e_7 is equal to e_8 or e_9 . It is not possible that e_6 or e_7 is equal to e_2 (or that e_8 or e_9 is equal to e_3) since we assumed that in the path c, a, b, d only the middle edge may be a parallel edge. It may be also possible that e_4 is equal to e_8 or e_9 , and that e_5 is equal to e_6 or e_7 . Note also that in each iteration step, the number of single edges is decreased by at least 1. So the number of iteration steps is at most n .



Fig. 1. Constructing the graph G_{i+1} from G_i .

First, we will show that $|3\text{-or}(G_i)| \geq |3\text{-or}(G_{i+1})|$ holds if all the edges e_1, \dots, e_9 are different. We will consider the more general case in which some of the e_j 's are equal at the end of the proof. To make the notation simpler, we set $i = 1$, i.e., we consider the graphs G_1 and G_2 . Note that there is a one-to-one correspondence between edges in G_1 and edges in G_2 . This implies that any arbitrary orientation in G_1 can be interpreted as an orientation in G_2 and vice versa. Take the standard orientation of G_1 to be the one consistent with the arrows in Figure 2. The interpretation of this orientation for G_1 yields the standard orientation for G_2 (also shown in Figure 2).



Fig. 2. The standard orientations in G_1 and G_2 .

We will prove our claim by defining an injective mapping $F : 3\text{-or}(G_2) \rightarrow 3\text{-or}(G_1)$. We want to use the identity mapping as far as possible. We set $C_2 = \{\alpha ; \alpha \in 3\text{-or}(G_2), \alpha \notin 3\text{-or}(G_1)\}$ and $C_1 = \{\alpha ; \alpha \in 3\text{-or}(G_1), \alpha \notin 3\text{-or}(G_2)\}$, and we will define F such that $F(\alpha) = \alpha$ for $\alpha \in 3\text{-or}(G_2) \setminus C_2$ and that $F : C_2 \rightarrow C_1$ is injective. Note that a mapping $F : 3\text{-or}(G_2) \rightarrow 3\text{-or}(G_1)$ defined this way is injective, since if $\beta \in C_1$, then $\beta \notin 3\text{-or}(G_2)$ and therefore β is not an image when using the identity function.

Let α be an arbitrary orientation. Note that all vertices $u \notin \{a, b, c, d\}$ have the same makespan in G_1 and in G_2 with respect to α . We identify first the class C_2 and consider the vertices a, b, c, d . We observe:

- | | |
|---|---|
| a has makespan 3 in G_2 | d has makespan 3 in G_2 |
| $\Rightarrow x_1 = x_2 = x_4 = 1$ | $\Rightarrow x_3 = 0, x_8 = x_9 = 1$ |
| $\Rightarrow a$ has makespan 3 in G_1 | $\Rightarrow d$ has makespan 3 in G_1 |

$$\begin{array}{ll}
b \text{ has makespan 3 in } G_2 & c \text{ has makespan 3 in } G_2 \\
\Rightarrow x_1 = x_2 = 0, x_5 = 1 & \Rightarrow x_3 = x_6 = x_7 = 1 \\
x_3 = 1 \Rightarrow b \text{ has makespan 3 in } G_1 & x_2 = 0 \Rightarrow c \text{ has makespan 3 in } G_1 \\
x_3 = 0, x_8 = x_9 = 1 & x_1 = 0 \wedge x_5 = 1 \\
\Rightarrow d \text{ has makespan 3 in } G_1 & \Rightarrow b \text{ has makespan 3 in } G_1 \\
x_6 = x_7 = 1 & x_2 = 1 \wedge x_1 = x_4 = 1 \\
\Rightarrow c \text{ has makespan 3 in } G_1 & \Rightarrow a \text{ has makespan 3 in } G_1
\end{array}$$

Collecting this characterization, we construct the class C_2 as

$$\begin{aligned}
C_2 = \{ & \alpha \notin 3\text{-or}(G_1); x_1 = x_2 = x_3 = 0 \wedge x_5 = 1 \wedge x_6 \cdot x_7 = x_8 \cdot x_9 = 0 \} \\
& \cup \{ \alpha \notin 3\text{-or}(G_1); x_2 = x_3 = x_6 = x_7 = 1 \\
& \quad \wedge x_1 \cdot x_4 = 0 \wedge (x_1 = 1 \vee x_5 = 0) \}.
\end{aligned}$$

In a similar way, we construct the class C_1 as

$$\begin{aligned}
C_1 = \{ & \alpha \notin 3\text{-or}(G_2); x_1 = 0 \wedge x_2 = x_3 = x_5 = 1 \wedge x_6 \cdot x_7 = 0 \} \\
& \cup \{ \alpha \notin 3\text{-or}(G_2); x_2 = x_3 = 0 \wedge x_6 = x_7 = 1 \\
& \quad \wedge x_8 \cdot x_9 = 0 \wedge (x_1 = 1 \vee x_5 = 0) \}.
\end{aligned}$$

Now, to define F , we consider 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 $F(x_1, 1, 1, x_4, x_5, 1, 1, x_8, x_9, \dots) = (x_1, 0, 0, x_4, x_5, 1, 1, x_8, x_9, \dots)$.




Fig. 3. The mapping F .

Note that vertices from $\{a, b, c, d\}$ have the same connections to vertices outside $\{a, b, c, d\}$; therefore, $\alpha \notin 3\text{-or}(G_1)$ implies that $F(\alpha) \notin 3\text{-or}(G_2)$. This implies that $F(\alpha) \in C_1$.

- (2) Consider $\alpha \in C_2$ with $x_1 = x_2 = x_3 = 0 \wedge x_5 = 1 \wedge x_6 \cdot x_7 = 0 \wedge x_8 \cdot x_9 = 0$.
Set $F(0, 0, 0, x_4, 1, x_6, x_7, x_8, x_9, \dots) = (0, 1, 1, x_4, 1, x_6, x_7, x_8, x_9, \dots)$.
In a way similar to case (1), we conclude that $F(\alpha) \in C_1$.

After these two cases, any orientation $\alpha \in H_2$ with

$$\begin{aligned}
H_2 = \{ & \alpha \in C_2 \mid x_2 = x_3 = x_6 = x_7 = 1 \\
& \quad \wedge x_1 \cdot x_4 = 0 \wedge x_8 = x_9 = 1 \wedge (x_1 = 1 \vee x_5 = 0) \}
\end{aligned}$$

has not been mapped by F , and orientations $\beta \in H_1$ with

$$H_1 = \{\beta \in C_1 \mid x_2 = x_3 = 0 \wedge x_6 = x_7 = 1 \wedge x_1 = x_4 = 1 \wedge x_8 \cdot x_9 = 0\} \\ \cup \{\beta \in C_1 \mid x_1 = 0, x_2 = x_3 = x_5 = 1 \wedge x_6 \cdot x_7 = 0 \wedge x_8 = x_9 = 1\}$$

are not images under F . We continue with these orientations.

(3) Set

$$H_{21} = \{\alpha \in C_2 ; x_2 = x_3 = x_6 = x_7 = x_8 = x_9 = 1 \wedge x_1 = 1 \wedge x_4 = 0\} \\ H_{11} = \{\beta \in C_1 ; x_2 = x_3 = 0 \wedge x_1 = x_4 = x_6 = x_7 = 1 \wedge x_8 \cdot x_9 = 0\}$$

We will show that $|H_{21}| \leq |H_{11}|$ holds.

Consider the cactoids T_{21} and T_{11} obtained by omitting the vertices a, b, c, d from H_{21} and H_{11} , respectively. T_{21} and T_{11} consist of edges and 6 pointers $e_j, 4 \leq j \leq 9$. Fixing the directions of the pointers in the same way as in the definitions of H_{21} and H_{11} , respectively, the number of 2-orientations of T_{21} is equal to $|H_{21}|$ and the number of 2-orientations of T_{11} is equal to $|H_{11}|$. See Figure 4 for an illustration. The pointers e_6 and e_7 have the same directions in T_{21} and T_{11} and e_5 has no specified direction in both cases. Edge e_4 has different directions in T_{21} and T_{11} . Directing edge e_4 in T_{21} towards vertex a would lead to an increased number of 2-orientations since the other vertex incident to e_4 has in this case makespan 2 with




Fig. 4. Orientations from the sets H_{21} and H_{11} .

a larger probability. Let \tilde{T}_{21} be the cactoid obtained from T_{21} by directing edge e_4 towards a . Then \tilde{T}_{21} and T_{11} differ only in the directions given to edges e_8 and e_9 . Let $P(i, j)$ be the probability of a 2-orientation in G_2 if $x_8 = i$ and $x_9 = j$. Set $m = \frac{3}{2}n$. Then, $\frac{|H_{11}|}{2^{m-3}} = P(0, 0) + P(0, 1) + P(1, 0) \geq P(1, 1) \geq \frac{|H_{21}|}{2^{m-3}}$, because of Lemma 2. It follows that $|H_{21}| \leq |H_{11}|$.

(4) To finish the first part of the proof, set

$$H_{22} = \{\beta \in C_2 ; x_2 = x_3 = x_6 = x_7 = x_8 = x_9 = 1 \wedge x_1 = x_5 = 0\} \\ H_{12} = \{\beta \in C_1 ; x_1 = 0 \wedge x_2 = x_3 = x_5 = x_8 = x_9 = 1 \wedge x_6 \cdot x_7 = 0\}$$

See Figure 5 below for an illustration. In the same way as in case (3), we show that $|H_{22}| \leq |H_{12}|$.




Fig. 5. Orientations from the sets H_{22} and H_{12} .

Since $H_2 = H_{21} \cup H_{22}$ and $H_1 = H_{11} \cup H_{12}$, there exists an injective mapping $F : 3\text{-or}(G_2) \rightarrow 3\text{-or}(G_1)$ in the case that all edges e_4, \dots, e_9 are different.

Now we consider the case that some of these edges are equal. If $e_i = e_j$ then in each orientation α the variables x_i and x_j get opposite values. Recall that the construction and proof of injectivity of the mapping F , which we described above, was done in 3 steps:

- (i) We defined $F(\alpha) = \alpha$ for all $\alpha \in 3\text{-or}(G_2) \setminus C_2$
- (ii) In cases (1) and (2) for some well defined $\alpha = (x_1, \dots, x_9, \dots)$, the value $F(\alpha)$ is obtained by negating x_2 and x_3 and leaving the other directions unchanged.
- (iii) $|H_2| \leq |H_1|$ is shown for the remaining cases.

Steps (i) and (ii) are not influenced by setting $x_i = \bar{x}_j$ for some $i, j \in \{4, \dots, 9\}, i \neq j$. So it remains to consider step (iii). If $e_i = e_j$ for $i \in \{6, 7\}, j \in \{8, 9\}$, then $x_i = \bar{x}_j$ holds and this implies that $H_2 = \emptyset$, since for all $\alpha \in H_2$ it holds $x_6 = x_7 = x_8 = x_9 = 1$. Clearly, this implies $|H_2| \leq |H_1|$.

So we can assume now that $e_i \neq e_j$ for $i \in \{6, 7\}, j \in \{8, 9\}$ and we consider the case $e_4 = e_5$. We will show first that $|H_{21}| \leq |H_{11}|$ holds also in this case. Consider the cactoids T_{21} and T_{11} obtained by deleting the vertices a, b, c, d from H_{21} and H_{11} . Since edge $e_4 = e_5$ connects vertices a and b , it is also deleted when the cactoids are formed. Each of the cactoids T_{11} and T_{21} has now only the 4 pointers $e_j, 6 \leq j \leq 9$. A simple inspection of the proof given above shows that $|H_{21}| \leq |H_{11}|$ holds also in this case. Furthermore, $|H_{22}| \leq |H_{21}|$ can be shown in the same way. The cases $e_4 = e_8$ and $e_5 = e_6$ can be handled in a very similar way. This completes the proof of the claim. \square

Our main result follows now as an immediate consequence of Theorem 2.

Corollary 1. For a 3-regular graph G with n vertices, $\text{SC}(G, \tilde{\mathbf{F}}) \geq \text{SC}(G_{\parallel}(n), \tilde{\mathbf{F}}) = 3 - \left(\frac{3}{4}\right)^{n/2}$.

We can also show that equality does *not* hold in Corollary 1.

Example 1. There is a 3-regular graph for which the Social Cost of the standard fully mixed Nash equilibrium is larger than for the corresponding parallel links graph.

4 Coordination Ratio

In this section, we present a bound on the Coordination Ratio for pure Nash equilibria.

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

Observation 1. *Restricted to pure Nash equilibria, for any interaction graph G , $\text{CR}_G \leq \Delta_G$, and this bound is tight.*

5 The Fully Mixed Nash Equilibrium

In this section, we study the fully mixed Nash equilibrium. For a graph $G = (V, E)$, for each edge $jk \in E$, denote jk the user corresponding to the edge jk . Denote \hat{p}_{jk} and \hat{p}_{kj} the probabilities (according to \mathbf{P}) that user jk chooses machines j and k , respectively. For each machine $j \in V$, the expected load of machine j excluding a set of edges \tilde{E} , denoted $\pi_{\mathbf{P}}(j) \setminus \tilde{E}$, is the sum $\sum_{k: jk \in E \setminus \tilde{E}} \hat{p}_{kj}$. As a useful combinatorial tool for the analysis of our counterexamples, we prove:

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 Nash equilibrium \mathbf{P} for G . Then, $\pi_{\mathbf{P}}(u) \setminus C_4 = \pi_{\mathbf{P}}(v) \setminus C_4$.*

Non-Existence Results. We first observe:

Counterexample 1. *There is no fully mixed Nash equilibrium for trees and meshes.*

We remark that the crucial property of trees that was used in the proof of Counterexample 1 is that each tree contains at least one leaf. Thus, Counterexample 1 actually applies to the more general class of graphs with no vertex of degree 1. We continue to prove:

Counterexample 2. *For each graph in Figure 6, there is no fully mixed Nash equilibrium.*

Our six counterexample graphs suggest that the existence of 4-cycles across the “boundary” of a graph or 1-connectivity may be crucial factors that disallow the existence of fully mixed Nash equilibria. Of course, this remains yet to be proved.

Uniqueness and Dimension Results. For *Complete Bipartite Graphs*, we prove:

Theorem 4. *Consider the complete bipartite graph $K_{r,s}$, where $s \geq r \geq 2$ and $s \geq 3$. Then, the fully mixed Nash equilibrium \mathbf{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 $s - 1$.*

Hypercube Graphs. Observe first that, in general, any point in $(0, 1)^r$ is mapped to a fully mixed Nash equilibrium with equal Nash probabilities on all edges of the same dimension (and “pointing” to the same direction). This implies:

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




Fig. 6. Six counterexample graphs

To show that r is also an upper bound, we need to prove that no other fully mixed Nash equilibria exist. We manage to do this only for $r \in \{2, 3\}$.

Theorem 5. *Consider the hypercube H_r , for $r \in \{2, 3\}$. Then, the fully mixed Nash dimension is r .*

Worst-Case Equilibria. We present two counterexamples to show that a fully mixed Nash equilibrium is not necessarily the worst-case Nash equilibrium, but it can be.

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

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

Acknowledgments: We thank Paul Spirakis and Karsten Tiemann for helpful discussions on the topic of our work.

References

1. B. Awerbuch, Y. Azar, Y. Richter and D. Tsur, "Tradeoffs in Worst-Case Equilibria," *Proceedings of the First International Workshop on Approximation and Online Algorithms*, pp. 41–52, LNCS 2909, 2003.
2. G. Christodoulou, E. Koutsoupias and A. Nanavati, "Coordination Mechanisms," *Proceedings of the 31st International Colloquium on Automata, Languages and Programming*, pp. 345–357, LNCS 3142, 2004.
3. A. Czumaj and B. Vöcking, "Tight Bounds for Worst-Case Equilibria," *Proceedings of the 13th Annual ACM Symposium on Discrete Algorithms*, pp. 413–420, 2002.
4. E. Even-Dar, A. Kesselman and Y. Mansour, "Convergence Time to Nash Equilibria," *Proceedings of the 30th International Colloquium on Automata, Languages and Programming*, pp. 502–513, LNCS 2719, 2003.
5. R. Feldmann, M. Gairing, T. Lücking, B. Monien and M. Rode, "Nashification and the Coordination Ratio for a Selfish Routing Game," *Proceedings of the 30th International Colloquium on Automata, Languages and Programming*, pp. 514–526, LNCS 2719, 2003.

6. S. Fischer and B. Vöcking, "A Counterexample to the Fully Mixed Nash Equilibrium Conjecture," *Technical Report, RWTH Aachen*, <http://aib.informatik.rwth-aachen.de/2005/2005-11.ps.gz>, 2005.
7. D. Fotakis, S. Koutsoyannis, E. Koutsoupias, M. Mavronicolas and P. Spirakis, "The Structure and Complexity of Nash Equilibria for a Selfish Routing Game," *Proceedings of the 29th International Colloquium on Automata, Languages and Programming*, pp. 123–134, LNCS 2380, 2002.
8. M. Gairing, T. Lücking, M. Mavronicolas and B. Monien, "Computing Nash Equilibria for Scheduling on Restricted Parallel Links," *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pp. 613–622, 2004.
9. M. Gairing, T. Lücking, M. Mavronicolas and B. Monien, "The Price of Anarchy for Polynomial Social Cost," *Proceedings of the 29th International Symposium on Mathematical Foundations of Computer Science*, pp. 574–585, LNCS 3153, 2004.
10. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and M. Rode, "Nash Equilibria in Discrete Routing Games with Convex Latency Functions," *Proceedings of the 31st International Colloquium on Automata, Languages and Programming*, pp. 645–657, LNCS 3142, 2004.
11. M. Gairing, T. Lücking, M. Mavronicolas, B. Monien and P. Spirakis, "Extreme Nash Equilibria," *Proceedings of the 8th Italian Conference on Theoretical Computer Science*, pp. 1–20, LNCS 2841, 2003. Full version: "The Structure and Complexity of Extreme Nash Equilibria," accepted to *Theoretical Computer Science*.
12. G. Gottlob, G. Greco and F. Scarcello, "Pure Nash Equilibria: Hard and Easy Games," *Proceedings of the 9th Conference on Theoretical Aspects of Rationality and Knowledge*, pp. 215–230, 2003.
13. S. Kakade, M. Kearns, J. Langford and L. Ortiz, "Correlated Equilibria in Graphical Games," *Proceedings of the 4th ACM Conference on Electronic Commerce*, pp. 42–47, 2003.
14. M. Kearns, M. Littman and S. Singh, "Graphical Models for Game Theory," *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*, pp. 253–260, 2001.
15. D. Koller and B. Milch, "Multi-Agent Influence Diagrams for Representing and Solving Games," *Games and Economic Behavior*, Vol. 45, No. 1, pp. 181–221, 2003.
16. E. Koutsoupias and C. H. Papadimitriou, "Worst-Case Equilibria," *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pp. 404–413, LNCS 1563, 1999.
17. P. La Mura, "Game Networks," *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, pp. 335–342, 2000.
18. M. Littman, M. Kearns and S. Singh, "An Efficient Exact Algorithm for Solving Tree-Structured Graphical Games," *Proceedings of the 15th Conference on Neural Information Processing Systems – Natural and Synthetic*, pp. 817–823, 2002.
19. T. Lücking, M. Mavronicolas, B. Monien, M. Rode, P. Spirakis and I. Vrto, "Which is the Worst-Case Nash equilibrium?," *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science*, pp. 551–561, LNCS 2747, 2003.
20. M. Mavronicolas and P. Spirakis, "The Price of Selfish Routing," *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pp. 510–519, 2001.
21. J. F. Nash, "Equilibrium Points in N -Person Games," *Proceedings of the National Academy of Sciences*, Vol. 36, pp. 48–49, 1950.