

Selfish Load Balancing Under Partial Knowledge*

Elias Koutsoupias¹, Panagiota N. Panagopoulou^{2,3}, and Paul G. Spirakis^{2,3}

¹ Department of Informatics and Telecommunications, University of Athens, Greece
elias@di.uoa.gr

² Computer Engineering and Informatics Department, Patras University, Greece

³ Research Academic Computer Technology Institute, Greece
{panagopp,spirakis}@cti.gr

Abstract. We consider n selfish agents or players, each having a load, who want to place their loads to one of two bins. The agents have an incomplete picture of the world: They know some loads exactly and only a probability distribution for the rest. We study Nash equilibria for this model, we compute the Price of Anarchy for some cases and show that sometimes extra information adversely affects the Divergence Ratio (a kind of subjective Price of Anarchy).

1 Introduction

We consider a simple version of load balancing in which n agents or players, each having a load, want to place their loads to one of two bins. We assume that the players are selfish and each one wants to minimize their own expected load of their own bin. This is a typical problem in the study of the Price of Anarchy. What distinguishes our approach here is that we aim at studying how the information regime affects the Price of Anarchy. We consider local strategies for the players in which a player has *exact* knowledge about the loads only of some players and *statistical* knowledge about the rest. Such a situation can practically arise in distributed situations, especially when there is not sufficient time for the n agents to communicate and coordinate.

There are n selfish players (agents) and $m = 2$ bins. Each player i has a load w_i and has to select a bin in $\{0, 1\}$ to place her load. The cost c_i of player i is the total load of the bin she selects, i.e., $c_i = \sum_{k: b_k = b_i} w_k$. The cost is apparently influenced by the decision of the other players which suggests that game theory is a proper framework to study this situation. However, before we are able to analyze the situation, we need to specify what kind of information is available to the players. What does player i know about the load of some other player j ? Among the many possibilities, we will concentrate on two extreme cases: Player

* Partially supported by the EU within the 6th Framework Programme under contracts 001907 “Dynamically Evolving, Large Scale Information Systems” (DELIS) and 015964 “Algorithmic Principles for Building Efficient Overlay Computers” (AEOLUS), and by the General Secretariat for Research and Technology of the Greek Ministry of Development within the Programme PENED 2003.




Fig. 1. Some information regime graph for 3 players

i knows either the exact value w_j (a real number in $[0, 1]$) or a probability distribution on w_j (the uniform distribution in $[0, 1]$). We can represent the information regimes by directed graphs which we will call *information regime graphs* (example in Fig. 1). An edge (i, j) denotes that player i knows the value w_j , while its absence denotes that the player i knows only that the value w_j is a random variable uniformly distributed in $[0, 1]$.

At the one extreme, when every player knows all w_j 's (which corresponds to a complete graph), the problem has been studied before as selfish task allocation [8]. At the other extreme, when each player knows only the distributions of the other players but not their actual values, the situation is very similar to the model studied by Papadimitriou and Yannakakis [11] for the purpose of quantifying the value of information. There are two main differences between our consideration and the model of [11]: First, we take the game-theoretic approach and assume that the players are selfish and care only about their own objective, while in the model of [11] the players want to optimize the global objective. Second, the cost in [11] was the probability of overflow, but in this work, the cost of a player is the total load of the bin she selects.

The main motivation of our work is to study the effects of information hiding on selfish systems. What happens to the Price of Anarchy in the more general case when players know only some values w_i ?

In this work, we consider only pure Nash equilibria. A trivial observation is that when we allow mixed Nash equilibria, there is always the fully-mixed equilibrium in which every player selects uniformly at random a bin. This Nash equilibrium has been studied extensively. For example, it was shown in [8] that the Price of Anarchy is $3/2$ in this case even when there are only 2 players.

To be able to compute the Price of Anarchy, we need to decide about the social cost and the optimal cost, that is the numerator and denominator of the Price of Anarchy. We consider the social cost to be either the sum (equivalent to the average) of the cost of all players or the maximum cost among players (which is equal to the makespan). The definition of the optimal cost is clear for the complete information regime, but in the case of incomplete information there are at least two natural choices. Either the optimal algorithm is a centralized one which knows all w_i 's, or the optimal algorithm is a distributed algorithm which knows exactly as much as the players of the game. In the latter case, each node knows only the w_i 's indicated by the information regime graph. The objective of each node, unlike the selfish objective that contributes to the numerator of the Price of Anarchy, is to minimize the global objective, the social cost. In this work, we will compute the Price of Anarchy using a distributed optimal algorithm.

Related Work. The game of n non-cooperating agents and m parallel links (or bins) was defined in [8] and has an extended literature e.g., [2,3,4,7,9,10]. In particular, the work of Monien et al [6] is closer in spirit to this work since they consider Bayesian games. Their model though is very different than ours.

The load balancing problem studied in this work was originally introduced by Papadimitriou and Yannakakis [11] in an effort to understand the crucial economic value of information [1] as a computational resource in a distributed system (e.g. in the context of Computational Complexity [12]). That work considered only $n = 3$ agents. In order to understand how the optimum solution achieved by the agents varies as a function of the amount of information available to them, Papadimitriou and Yannakakis [11] considered each possible communication pattern and discovered the corresponding optimal decision protocol to be unexpectedly sophisticated. For the special case where no communication is allowed, i.e. when each agent i is aware of only her own load w_i , it was conjectured that the simple decision rule: “if $w_i \leq 1 - \frac{1}{\sqrt{7}}$ then put w_i on bin 0 else put w_i on bin 1” is optimal; the conjecture is still open. Georgiades et al. [7] studied the extension of the load balancing problem to the case of n agents. Their work was focused on the special cases of oblivious decision rules, for which agents do not “look at” their inputs, and non-oblivious decision rules, for which they do. In either case, optimality conditions were derived in the form of combinatorial polynomial equations.

Our Results. We study the Nash equilibria of these games and give some results about the Divergence Ratio and the Price of Anarchy. In particular since the players have incomplete picture of the world, the cost that they compute may differ from the actual cost. To capture this situation, we define the Divergence Ratio, a kind of subjective Price of Anarchy. We show that for the regime of total lack of information, the Divergence Ratio is $(n+1)/n$ for even n and 1 for odd n . For the regime of complete information the Divergence Ratio is $4/3$. In contrast, the Divergence Ratio is $\Theta(n)$ for some intermediate information regimes. We also estimate the Price of Anarchy for the total lack of information regime.

2 The Model

Consider a set $N = \{1, 2, \dots, n\}$ of n selfish agents. Each agent i has a *load* $w_i \in [0, 1]$. Let $\mathbf{w} = (w_i)_{i \in N}$ denote a vector of input loads, one per agent. For any pair of agents $(i, j) \in N \times N$, agent i knows either (a) the exact value of w_j or (b) that w_j is uniformly distributed on $[0, 1]$. For each $i \in N$ let $I_i = \{j \in N : \text{agent } i \text{ knows the exact value of } w_j\}$, and thus for each $j \notin I_i$, agent i knows that w_j is uniformly distributed on $[0, 1]$. Denote by $\mathbf{I} = (I_i)_{i \in N}$ the collection of I_i for all $i \in N$. Let's also denote the cardinality of I_i by γ_i .

Each agent $i \in N$ has to select one of two available bins (bin 0 and bin 1) to put her load. The bin is selected based on the values of w_i 's in I_i . Thus, a *strategy* for agent i is a function s_i from $[0, 1]^{\gamma_i}$ to $\{0, 1\}$: $s_i((w_j)_{j \in I_i})$.

Of particular interest are the *single-threshold* strategies where each agent i places her load to the first bin iff w_i is below a threshold value t_i . This threshold value depends on the known loads of the agents, i.e. t_i is a function of $(w_j)_{j \in I_i}$.

Here we study threshold strategies, but *most of our results can be extended easily to all strategies*. The reason is simple: What matters in most of our considerations is the expected load that a player places in one bin. In particular, for a threshold t_i , the expected load that player i places in bin 0 is $\int_0^{t_i} w_i dw_i = t_i^2/2$ and in bin 1 it is $\int_{t_i}^1 w_i dw_i = 1/2 - t_i^2/2$. From the point of view of the other players that do not know the value w_i , only the expected value $t_i^2/2$ matters. But this, in general, can be achieved with many strategies (for example, by the inverse threshold strategy that places the load w_i in bin 0 when its value is greater than $\bar{t}_i$ where $\bar{t}_i^2/2 = 1/2 - t_i^2/2$).

A *strategy profile* $\mathbf{s} = (s_1, \dots, s_n)$ is a collection of strategies, one for each agent. Denote by $(s'_i, \mathbf{s}_{-i})$ the strategy profile that is identical to $\mathbf{s}$ except for agent i , who chooses strategy s'_i instead of s_i . Since we will mainly consider threshold strategies, we will denote such a strategy profile by $\mathbf{t} = (t_1, \dots, t_n) \in [0, 1]^n$.

For weights $\mathbf{w}$ and strategies $\mathbf{s}$, we define two costs of each player i . The *subjective cost* of a player is the expected cost that player computes based on her knowledge of the weights in I_i and assuming that the other weights are uniformly (and independently) distributed in $[0, 1]$. The *(objective) cost* is the actual cost which can be computed from the full knowledge of $\mathbf{w}$ and $\mathbf{s}$. The *social cost* is the cost of the system for these weights and strategies and it is either the sum of the cost of all agents or the maximum cost among agents; the latter corresponds to the makespan. In analogy to the costs of the players, we have the subjective social cost and the (objective) social cost. We will denote the subjective cost by $\text{cost}_i(\mathbf{s}; \mathbf{w}; I_i)$.

Definition 1. The strategy profile $\mathbf{s} = (s_1, \dots, s_n)$ is a Nash equilibrium if and only if, for all $\mathbf{w}$, $i \in N$, and s'_i , $\text{cost}_i(\mathbf{s}; \mathbf{w}; I_i) \leq \text{cost}_i((s'_i, \mathbf{s}_{-i}); \mathbf{w}; I_i)$.

3 The Structure of Nash Equilibria

Lemma 1. For any threshold strategy profile $\mathbf{t} \in [0, 1]^n$ and for all $i \in N$,

$$i \notin I_i \Rightarrow \text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) = t_i \left(\sum_{j \in I_i: w_j \leq t_j} w_j - \sum_{j \in I_i: w_j > t_j} w_j + \sum_{j \notin I_i, j \neq i} t_j^2 - \frac{n - \gamma_i - 1}{2} \right) + \frac{1}{2} + \sum_{j \in I_i: w_j > t_j} w_j + \frac{n - \gamma_i - 1}{2} - \frac{1}{2} \sum_{j \notin I_i, j \neq i} t_j^2 \quad (1)$$

$$i \in I_i \Rightarrow \text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) = \begin{cases} \sum_{j \in I_i: w_j \leq t_j} w_j + \frac{1}{2} \sum_{j \notin I_i} t_j^2 & \text{if } w_i \leq t_i \\ \sum_{j \in I_i: w_j > t_j} w_j + \frac{n - \gamma_i}{2} - \frac{1}{2} \sum_{j \notin I_i} t_j^2 & \text{if } w_i > t_i \end{cases} \quad (2)$$

Proof. Fix a strategy profile $\mathbf{t} \in [0, 1]^n$ and an agent $i \in N$. Assume that $i \notin I_i$. Then i does not know the exact value of her own load, so she expects that her load will be put in bin 0 with probability t_i and in bin 1 with probability $1 - t_i$.

Therefore (i) with probability t_i , the cost for agent i is her own expected load, plus the sum of the loads of all $j \in I_i$ such that $w_j \leq t_j$, plus the expected load that every other agent $j \notin I_i$ puts on bin 0 (that is, $t_j^2/2$), and (ii) with

probability $1 - t_i$, the cost for agent i is her own expected load, plus the sum of the loads of all $j \in I_i$ such that $w_j > t_j$, plus the expected load that every other agent $j \notin I_i$ puts on bin 1 (that is, $1/2 - t_j^2/2$).

Agent's i expected load is $1/2$ (i.e. $\int_0^1 w_i dw_i$). Therefore, if $i \notin I_i$, then the cost of agent $i \in N$ in the strategy profile $\mathbf{t}$ is

$$\begin{aligned} \text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) &= \frac{1}{2} + t_i \left(\sum_{j \in I_i: w_j \leq t_j} w_j + \sum_{j \notin I_i, j \neq i} \frac{t_j^2}{2} \right) \\ &\quad + (1 - t_i) \left(\sum_{j \in I_i: w_j > t_j} w_j + \sum_{j \notin I_i, j \neq i} \frac{1 - t_j^2}{2} \right) \\ &= t_i \left(\sum_{j \in I_i: w_j \leq t_j} w_j - \sum_{j \in I_i: w_j > t_j} w_j + \sum_{j \notin I_i, j \neq i} t_j^2 - \frac{n - \gamma_i - 1}{2} \right) \\ &\quad + \frac{1}{2} + \sum_{j \in I_i: w_j > t_j} w_j + \frac{n - \gamma_i - 1}{2} - \frac{1}{2} \sum_{j \notin I_i, j \neq i} t_j^2. \end{aligned}$$

Now, if $i \in I_i$, i.e. i knows the exact value of her load, then

$$\begin{aligned} \text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) &= \begin{cases} \sum_{j \in I_i: w_j \leq t_j} w_j + \sum_{j \notin I_i} t_j^2/2 & \text{if } w_i \leq t_i \\ \sum_{j \in I_i: w_j > t_j} w_j + \sum_{j \notin I_i} (1/2 - t_j^2/2) & \text{if } w_i > t_i \end{cases} \\ &= \begin{cases} \sum_{j \in I_i: w_j \leq t_j} w_j + \frac{1}{2} \sum_{j \notin I_i} t_j^2 & \text{if } w_i \leq t_i \\ \sum_{j \in I_i: w_j > t_j} w_j + \frac{n - \gamma_i}{2} - \frac{1}{2} \sum_{j \notin I_i} t_j^2 & \text{if } w_i > t_i \end{cases}. \quad \square \end{aligned}$$

3.1 The Total Lack of Information Case, $I_i = \emptyset$

Assume $I_i = \emptyset$ for all $i \in N$. From Equation 1 of Lemma 1, the cost of any agent $i \in N$ for a strategy profile $\mathbf{t} = (t_1, \dots, t_n) \in [0, 1]^n$ is

$$\text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) = t_i \left(\sum_{j \neq i} t_j^2 - \frac{n-1}{2} \right) + \frac{n}{2} - \frac{1}{2} \sum_{j \neq i} t_j^2. \quad (3)$$

Proposition 1. *Consider the case where $I_i = \emptyset$ for all $i \in N$. Then the strategy profile $\mathbf{t} \in [0, 1]^n$ is a Nash equilibrium if and only if, for all $i \in N$,*

$$\begin{aligned} t_i = 0 &\Rightarrow \sum_{j \neq i} t_j^2 \geq \frac{n-1}{2} \\ t_i = 1 &\Rightarrow \sum_{j \neq i} t_j^2 \leq \frac{n-1}{2} \\ t_i \in (0, 1) &\Rightarrow \sum_{j \neq i} t_j^2 = \frac{n-1}{2}. \end{aligned}$$

Proof. Fix a strategy profile $\mathbf{t} \in [0, 1]^n$. By Definition 1, $\mathbf{t}$ is a Nash equilibrium if and only if, for all $i \in N$, $\text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) = \min_{t'_i \in [0, 1]} \text{cost}_i((t'_i, \mathbf{t}_{-i}); \mathbf{w}; I_i)$. Now fix an agent $i \in N$. From Equation 3, $\text{cost}_i(\mathbf{t}; \mathbf{w}; I_i)$ is an affine function of t_i . Therefore, it is minimized at $t_i = 0$ only if it is non-decreasing, i.e. only if $\sum_{j \neq i} t_j^2 \geq \frac{n-1}{2}$. It is minimized at $t_i = 1$ only if it is non-increasing, i.e. only if $\sum_{j \neq i} t_j^2 \leq \frac{n-1}{2}$. Finally, it is minimized at some $t_i \in (0, 1)$ only if it is a constant function, i.e. only if $\sum_{j \neq i} t_j^2 = \frac{n-1}{2}$. $\square$

Observe that, in a Nash equilibrium, all $i \in N$ such that $t_i \in (0, 1)$ must have equal t_i 's. This is because, for all $i \in N$ such that $t_i \in (0, 1)$, $\sum_{j \neq i} t_j^2 = \frac{n-1}{2}$ which implies that $t_i^2 = \sum_{j \in N} t_j^2 - \frac{n-1}{2}$.

With this, we can now characterize all the Nash equilibria of the total lack of information case.

Theorem 1. *Consider the case where $I_i = \emptyset$ for all $i \in N$. Then the strategy profile $\mathbf{t} \in [0, 1]^n$ is a Nash equilibrium if and only if κ agents choose threshold 1, λ agents choose threshold $t_A \in (0, 1)$, $n - \kappa - \lambda$ agents choose threshold 0 and*

- (1) $\frac{n-1}{2} - \lambda \leq \kappa \leq \frac{n-1}{2}$, $\lambda > 1$, $t_A^2 = \frac{n-1}{2(\lambda-1)} - \frac{\kappa}{\lambda-1}$ or
- (2) n is even, $\kappa = \frac{n}{2}$, $\lambda = 0$ or
- (3) n is odd, $\kappa = \frac{n+1}{2}$, $\lambda = 0$ or
- (4) n is odd, $\kappa = \frac{n-1}{2}$, $\lambda = 0$ or
- (5) n is odd, $\kappa = \frac{n-1}{2}$, $\lambda = 1$.

Moreover, the maximum, over all Nash equilibria, Social Cost is $\frac{n+1}{4}$.

Proof. In order to find all Nash equilibria we have to find all the possible partitions of the set of agents into three sets A , B and C so that

- For all $i \in A$, $t_i = t_A$ for some $t_A \in (0, 1)$ and $\sum_{j \neq i} t_j^2 = \frac{n-1}{2}$, or equivalently

$$\begin{aligned} (|A| - 1) \cdot t_A^2 + |B| \cdot 0 + |C| \cdot 1 &= \frac{n-1}{2} \\ (|A| - 1)t_A^2 + |C| &= \frac{n-1}{2}. \end{aligned}$$

- For all $i \in B$, $t_i = 0$ and $\sum_{j \neq i} t_j^2 \geq \frac{n-1}{2}$, or equivalently

$$|A|t_A^2 + |C| \geq \frac{n-1}{2}.$$

- For all $i \in C$, $t_i = 1$ and $\sum_{j \neq i} t_j^2 \leq \frac{n-1}{2}$, or equivalently

$$|A|t_A^2 + |C| - 1 \leq \frac{n-1}{2}.$$

We consider the following cases.

1. $|A| = 0$. Then there is a Nash equilibrium if and only if $\frac{n-1}{2} \leq |C| \leq \frac{n+1}{2}$. Since $|B| = n - |C|$, it must also hold that $\frac{n-1}{2} \leq |B| \leq \frac{n+1}{2}$. This is possible if and only if (1) n is even and $|B| = |C| = n/2$, in which case the cost for any agent is $n/4$, or (2) if n is odd, $|B| = \frac{n+1}{2}$ and $|C| = \frac{n-1}{2}$, in which case the cost for any agent that chooses threshold 0 is $\frac{n+1}{4}$ and the cost for any agent that chooses threshold 1 is $\frac{n-1}{4}$ or if (3) n is odd, $|B| = \frac{n-1}{2}$ and $|C| = \frac{n+1}{2}$, in which case the cost for any agent that chooses threshold 0 is $\frac{n-1}{4}$ and the cost for any agent that chooses threshold 1 is $\frac{n+1}{4}$.
2. $|A| = 1$. Then there is a Nash equilibrium if and only if $|C| = \frac{n-1}{2}$ and $0 < t_A < 1$. Then $|B| = n - 1 - \frac{n-1}{2} = \frac{n-1}{2}$. So in this case we have a Nash equilibrium if and only if n is odd, $|B| = |C| = \frac{n-1}{2}$ and $0 < t_A < 1$. Moreover, the cost for any agent that chooses threshold 0 is $\frac{n+1}{4} - \frac{t_A^2}{2}$, the cost for any agent that chooses threshold 1 is $\frac{n-1}{4} + \frac{t_A^2}{2}$, and the cost for the agent that chooses $t_A \in (0, 1)$ is $\frac{n+1}{4}$.
3. $|A| > 1$. Then there is a Nash equilibrium if and only if

$$\begin{aligned} (|A| - 1)t_A^2 + |C| &= \frac{n-1}{2} \quad \text{and} \\ |A|t_A^2 + |C| &\geq \frac{n-1}{2} \quad \text{and} \\ |A|t_A^2 + |C| - 1 &\leq \frac{n-1}{2}. \end{aligned}$$

Simple calculations yield that the above are equivalent to $|C| \leq \frac{n-1}{2}$ and $|C| \geq \frac{n+1}{2} - |A|$ and $t_A^2 = \frac{n-1}{2(|A|-1)} - \frac{|C|}{|A|-1}$. Furthermore, the cost for any agent $i \in N$ such that $t_i = 0$ is

$$\text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) = \frac{n}{2} - \frac{1}{2}(\lambda t_A^2 + \kappa) < \frac{n+1}{4},$$

the cost for any agent $i \in N$ such that $t_i = 1$ is

$$\text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) = \kappa - 1 + \lambda t_A^2 - \frac{n-1}{2} + \frac{n}{2} - \frac{1}{2}(\kappa - 1 + \lambda t_A^2) < \frac{n+1}{4},$$

and the cost for any agent $i \in N$ such that $t_i = t_A$ is

$$\text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) = \frac{n}{2} - \frac{1}{2}((\lambda - 1)t_A^2 + \kappa) = \frac{n+1}{4}. \quad \square$$

4 The Divergence Ratio

In order to study the impact that the information regime plays on the performance of the system, we first consider the *Divergence Ratio*. This is essentially the subjective Price of Anarchy and it is the worst-case ratio of the system cost at a Nash equilibrium over the minimum system cost.

We first define the *social cost* $\text{SC}(\mathbf{t}, \mathbf{I})$ to be the maximum subjective selfish cost over all agents, i.e. $\text{SC}(\mathbf{t}, \mathbf{I}) = \max_{i \in N} \text{cost}_i(\mathbf{t}; \mathbf{w}; I_i)$. Notice that here consider the social cost as the maximum among the costs of all players. We could

define it also as the sum of the costs of all players. All results in this section extend easily to this social cost as well, although we omit them from this abstract for lack of space.

The *Players' Optimum* $\text{PO}(\mathbf{I})$ is the minimum, over all possible strategy profiles $\mathbf{t} \in [0, 1]^n$, Social Cost: $\text{PO}(\mathbf{I}) = \min_{\mathbf{t} \in [0, 1]^n} \text{SC}(\mathbf{t}, \mathbf{I})$. That is, the Players' Optimum corresponds to a strategy profile that minimizes the maximum cost seen by the agents.

The *Divergence Ratio* $\text{DR}(\mathbf{I})$ is the worst-case, over all weights $\mathbf{w}$ and Nash equilibria $\mathbf{t}$, of the ratio $\frac{\text{SC}(\mathbf{t}, \mathbf{I})}{\text{PO}(\mathbf{I})}$:

$$\text{DR}(\mathbf{I}) = \max_{\mathbf{w}} \max_{\mathbf{t}: \mathbf{t} \text{ N.E.}} \frac{\text{SC}(\mathbf{t}, \mathbf{I})}{\text{PO}(\mathbf{I})}.$$

4.1 The Case of $I_i = \emptyset$

We will now show that the Divergence Ratio for the total lack of information regime is small. Recall from Theorem 1 that the Social Cost for this information regime is $(n + 1)/4$. We now need to compute the Players' Optimum for this regime.

Lemma 2. *Consider the case where $I_i = \emptyset$ for all $i \in N$. Then $\text{PO}(\mathbf{I}) = \frac{n}{4}$ if n is even and $\text{PO}(\mathbf{I}) = \frac{n+1}{4}$ if n is odd.*

Proof (Sketch). Let $\mathbf{t}^*$ denote the strategy profile that corresponds to the Players' Optimum. We consider the following cases:

- Case 1: $t_i^* \in \{0, 1\}$ for all $i \in N$. In this case, the set of agents N is partitioned into 2 subsets N_0 and N_1 such that $t_i^* = 0$ for all $i \in N_0$ and $t_i^* = 1$ for all $i \in N_1$, $|N_0| + |N_1| = n$ and the Social Cost is minimized, i.e. $\text{PO}(\mathbf{I}) = \min_{N_0 \subseteq N} \max \left\{ \frac{|N_0|}{2}, \frac{n - |N_0|}{2} \right\}$. If n is even, then $|N_0| = n/2$ and $\text{PO}(\mathbf{I}) = \frac{n}{4}$. If n is odd, then $|N_0| = (n - 1)/2$ or $|N_0| = (n + 1)/2$ and $\text{PO}(\mathbf{I}) = \frac{n+1}{4}$.
- Case 2: There exists some $i \in N$ such that $t_i^* \in (0, 1)$. By contradiction, it can be shown that in this case there must be some agent j such that $\text{cost}_j(\mathbf{t}^*; \mathbf{w}; \mathbf{I}) \geq \frac{n+1}{4}$ if n is odd or some agent j such that $\text{cost}_j(\mathbf{t}^*; \mathbf{w}; \mathbf{I}) \geq \frac{n}{4}$ if n is even. $\square$

Theorem 1 and Lemma 2 immediately yield:

Theorem 2. *Consider the case where $I_i = \emptyset$ for all $i \in N$. Then $\text{DR}(\mathbf{I}) = 1 + \frac{1}{n}$ if n is even and $\text{DR}(\mathbf{I}) = 1$ if n is odd.*

4.2 The Case of $I_i = N$

Assume that $I_i = N$ for all $i \in N$. From Equation 2 of Lemma 1, the cost of any agent $i \in N$ for a strategy profile $\mathbf{t} = (t_1, \dots, t_n) \in [0, 1]^n$ is

$$\text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) = \begin{cases} \sum_{j \in N: w_j \leq t_j} w_j & \text{if } w_i \leq t_i \\ \sum_{j \in N: w_j > t_j} w_j & \text{if } w_i > t_i \end{cases}.$$

An important observation in this case is that, since each agent knows the exact values of the loads of all other agents, it suffices to consider single-threshold strategies of the form $t_i = 0$ or $t_i = 1$, for all $i \in N$. For example, consider a Nash equilibrium $\mathbf{t}$ such that $0 < t_k < 1$ for some $k \in N$. Assume that $w_k \leq t_k$. All agents know that $w_k \leq t_k$, so all agents know that w_k goes on bin 0. Therefore, all strategy profiles $(t'_k, \mathbf{t}_{-k})$ such that $w_k \leq t'_k \leq 1$ are Nash equilibria, which are equivalent (as regards the selfish costs of the agents) to the Nash equilibrium $\mathbf{t}$. Similar arguments apply for the case $w_k > t_k$. Therefore the Nash equilibria in this case correspond to the *pure* Nash equilibria of the KP model [8] with n agents and 2 links. So in this case the Divergence Ratio is identical to the *pure* Price of Anarchy in the KP model with n agents and 2 links. In [5] (Theorem 7.1), it is shown that the pure Price of Anarchy in this setting is at most $\frac{4}{3}$ and this bound is tight. We give an alternative proof of this bound:

Theorem 3. *Consider the case where $I_i = N$ for all $i \in N$. Then $\text{DR}(\mathbf{I}) = \frac{4}{3}$.*

Proof. Consider an arbitrary Nash equilibrium $\mathbf{t}^*$. Then there is an equivalent, with respect to the costs and the Divergence Ratio, Nash equilibrium $\mathbf{t}$ such that for all $i \in N$, $t_i = 0$ or $t_i = 1$. The total load on bin 0 is $B_0(\mathbf{t}) = \sum_{i:t_i=1} w_i$ and the total load on bin 1 is $B_1(\mathbf{t}) = \sum_{i:t_i=1} w_i$. Therefore the cost for agent $i \in N$ is

$$\text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) = \begin{cases} B_0(\mathbf{t}) & \text{if } t_i = 1 \\ B_1(\mathbf{t}) & \text{if } t_i = 0 \end{cases}.$$

Without loss of generality, assume that $B_0(\mathbf{t}) \geq B_1(\mathbf{t})$. Thus $\text{SC}(\mathbf{t}, \mathbf{I}) = B_0(\mathbf{t})$. Moreover, $\text{PO}(\mathbf{I}) \geq \frac{\sum_{i \in N} w_i}{2} = \frac{B_0(\mathbf{t}) + B_1(\mathbf{t})}{2}$. Now if only one agent, say agent i , places her load on bin 0 (i.e. $t_i = 1$ and $t_j = 0$ for all $j \neq i$) then $\text{PO}(\mathbf{I}) = B_0(\mathbf{t})$ and $\text{DR}(\mathbf{I}) = 1$.

Otherwise, there are at least two loads on bin 0. Thus there exists an agent i that chooses bin 0 (i.e. with $t_i = 1$) such that $w_i \leq \frac{B_0(\mathbf{t})}{2}$. Since $\mathbf{t}$ is a Nash equilibrium, it holds that $B_0(\mathbf{t}) \leq B_1(\mathbf{t}) + w_i$, implying that $B_0(\mathbf{t}) \leq B_1(\mathbf{t}) + \frac{B_0(\mathbf{t})}{2}$ and that $B_1(\mathbf{t}) \geq \frac{B_0(\mathbf{t})}{2}$. Therefore, $\text{PO}(\mathbf{I}) \geq \frac{B_0(\mathbf{t}) + B_1(\mathbf{t})}{2} \geq \frac{3B_0(\mathbf{t})}{4}$ and $\text{DR}(\mathbf{I}) = \max_{\mathbf{t}: \mathbf{t} \text{ N.E.}} \frac{\text{SC}(\mathbf{t}, \mathbf{I})}{\text{PO}(\mathbf{I})} \leq \frac{4}{3}$.

Now consider the case where n is even and $n > 2$, $w_1 = w_2 = (n-2)\alpha$ and $w_i = \alpha$ for all $i \neq 1, 2$, for some $\alpha \in \left(0, \frac{1}{n-2}\right]$. Then the strategy profile $\mathbf{t}$ where $t_1 = t_2 = 1$ and $t_i = 0$ for all $i \neq 1, 2$ is a Nash equilibrium which gives a Social Cost equal to $2(n-2)\alpha$. In this case however $\text{PO}(\mathbf{I}) = \frac{3}{2}(n-2)\alpha$ and thus $\text{DR}(\mathbf{I}) \geq \frac{\frac{2(n-2)\alpha}{2}}{\frac{3}{2}(n-2)\alpha} = \frac{4}{3}$. $\square$

4.3 The Case of Arbitrary I_i

We have shown that, in the case where the agents have no information about the exact value of their weights, then the Divergence Ratio is very close to 1.

The same holds when the agents have complete information. In contrast, we will next show that if $i \in I_i$ and the cardinality of I_i is sufficiently small, then the Divergence Ratio grows significantly.

Theorem 4. *If $\gamma_i = \gamma \leq \frac{n-2}{3}$ and $i \in I_i$ for all $i \in N$, then $\text{DR}(\mathbf{I}) \geq \frac{n+\gamma+2}{4\gamma+4}$.*

Proof. For the proof, we focus on the instance where $w_i = 1$ for all $i \in N$. Our goal is to find (a) a Nash equilibrium $\mathbf{t}$ of low Social Cost, so as to upper bound the Players' Optimum, and (b) a Nash equilibrium $\mathbf{t}'$ of high Social Cost, so as to lower bound the worst possible Social Cost:

- (a) Consider the strategy profile $\mathbf{t}$ such that $t_i = 1 - \frac{1}{n-\gamma}$ for all $i \in N$. Then the cost for any agent $i \in N$ is

$$\text{cost}_i(\mathbf{t}; \mathbf{w}; I_i) = \gamma + \frac{n-\gamma}{2} - \frac{n-\gamma}{2} \left(1 - \frac{1}{n-\gamma}\right)^2 = \gamma + 1 - \frac{1}{2(n-\gamma)}.$$

The profile $\mathbf{t}$ is a Nash equilibrium, since the cost for i if she chose bin 0 would be

$$1 + \frac{n-\gamma}{2} \left(1 - \frac{1}{n-\gamma}\right)^2 \geq \gamma + 1 + \frac{1}{2(n-\gamma)} > \text{cost}_i(\mathbf{t}; \mathbf{w}; I_i).$$

The Social Cost of $\mathbf{t}$ is $\text{SC}(\mathbf{t}, \mathbf{I}) = \max_{i \in N} \text{cost}_i(\mathbf{t}; \mathbf{w}; \mathbf{I}) \leq \gamma + 1$.

- (b) Now consider the profile $\mathbf{t}'$ where $t'_i = \sqrt{\frac{1}{2} + \frac{\gamma-1}{n-\gamma}}$ for all $i \in N$ (since $\gamma \leq \frac{n-2}{3}$, $t'_i \in (0, 1)$). Then the cost for any agent $i \in N$ is

$$\text{cost}_i(\mathbf{t}'; \mathbf{w}; I_i) = \gamma + \frac{n-\gamma}{2} - \frac{n-\gamma}{2} (t'_i)^2 = \frac{n+\gamma+2}{4}.$$

The profile $\mathbf{t}'$ is also a Nash equilibrium, since the cost for i if she chose bin 0 would be

$$1 + \frac{n-\gamma}{2} (t'_i)^2 = \frac{n+\gamma+2}{4} = \text{cost}_i(\mathbf{t}'; \mathbf{w}; I_i).$$

The Social Cost of $\mathbf{t}'$ is $\text{SC}(\mathbf{t}', \mathbf{I}) = \max_{i \in N} \text{cost}_i(\mathbf{t}'; \mathbf{w}; \mathbf{I}) = \frac{n+\gamma+2}{4}$.

Thus the Divergence Ratio is

$$\text{DR}(\mathbf{I}) = \max_{\hat{\mathbf{t}}: \hat{\mathbf{t}} \text{ N.E.}} \frac{\text{SC}(\hat{\mathbf{t}}, \mathbf{I})}{\text{PO}(\mathbf{I})} \geq \frac{\text{SC}(\mathbf{t}', \mathbf{I})}{\text{SC}(\mathbf{t}, \mathbf{I})} \geq \frac{\frac{n+\gamma+2}{4}}{\gamma+1} = \frac{n+\gamma+2}{4\gamma+4}. \quad \square$$

Corollary 1. *If $\gamma_i = \gamma = o(n)$ and $i \in I_i$ for all $i \in N$, then $\lim_{n \rightarrow \infty} \text{DR}(\mathbf{I}) = \infty$.*

Interestingly, the above observations hold also for the case of the social cost which is the sum of the costs of all players and it is an easy consequence of the symmetry of the Nash equilibria in this section.

5 The Price of Anarchy

We now consider the objective cost and the Price of Anarchy for the total lack of information regime. In particular, for a given strategy profile $\mathbf{t}$ (or more general strategy profile $\mathbf{s}$), we define the objective cost of player i to be the expected load of the bin selected by i . The expectation is over uniformly distributed $\mathbf{w} \in [0, 1]^n$. Accordingly, we define the social cost as the sum of the costs of all players. Notice that the social cost here is the sum of the cost of all players. The case of max social cost is more complicated and we leave it as an interesting open problem.

The Price of Anarchy is the worst-case ratio of the social cost at a Nash equilibrium over the expected optimum. To compute the optimum, we assume that the optimal algorithm is distributed and has the same information as the agents. In the total lack of information regime, this is easy to define when we consider only pure strategies: $\lceil n/2 \rceil$ of the agents select bin 0 and the rest select bin 1. The expected optimal cost of each agent is either $\frac{1}{2}\lceil n/2 \rceil$ or $\frac{1}{2}\lfloor n/2 \rfloor$. The sum of the costs of all agents is

$$\text{OPT} = \begin{cases} \frac{n^2}{4} & \text{for even } n \\ \frac{n^2+1}{4} & \text{for odd } n \end{cases}.$$

Theorem 5. *The Price of Anarchy for the total lack of information regime is*

$$\text{PA} = \begin{cases} \frac{n+1}{n} & \text{for even } n \\ \frac{n(n+1)}{n^2+1} & \text{for odd } n \end{cases}.$$

Proof. From the characterization of the Nash equilibria in Theorem 1, we can compute the cost for each agent. For the non-deterministic agents, i.e., the agents that have a threshold in $(0, 1)$, the cost is exactly $(n+1)/4$. This follows either by computing the cost explicitly: $1/2$, which is the expected cost of her own load, plus the expected cost of bin 0 due to other agents $\kappa/2 + t_A^2/(\lambda-1) = (n+1)/4$ (the notation is from Theorem 1). A more direct way follows from the definition of the Nash equilibrium: The expected load on the bins due to other agents should be the same which happens when the load is $(n-1)/4$. Adding the expected load $1/2$ of her own, the cost of agent i is $(n+1)/4$.

By similar considerations, for the deterministic agents the cost is at most $(n+1)/4$. The worst case happens when all agents are non-deterministic. The total cost of all players is $n(n+1)/4$. The Price of Anarchy follows by dividing by the optimal OPT. $\square$

References

1. Arrow, K.: The Economics of Information. Harvard University Press (1984)
2. 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)

3. Fotakis, D., Kontogiannis, S., Koutsoupias, E., Mavronicolas, M., Spirakis, 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. 123–134. Springer, Heidelberg (2002)
4. Gairing, M., Lücking, T., Mavronicolas, M., Monien, B., Rode, M.: Nash Equilibria in Discrete Routing Games with Convex Latency Functions. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 645–657. Springer, Heidelberg (2004)
5. Gairing, M., Lücking, T., Mavronicolas, M., Monien, B., Spirakis, P.: Structure and Complexity of Extreme Nash Equilibria. *Theoretical Computer Science* 343(1-2), 133–157 (2005)
6. Gairing, M., Monien, B., Tiemann, K.: Selfish routing with incomplete information. *SPAA 2005*, 203–212 (2005)
7. Georgiades, S., Mavronicolas, M., Spirakis, P.: Optimal, Distributed Decision-Making: The Case of No Communication. In: *Proceedings of the 12th International Symposium on Fundamentals of Computation Theory*, pp. 293–303 (1999)
8. Koutsoupias, E., Papadimitriou, C.H.: Worst-Case Equilibria. In: Meinel, C., Tiison, S. (eds.) *STACS 99*. LNCS, vol. 1563, pp. 404–413. Springer, Heidelberg (1999)
9. Lücking, T., Mavronicolas, M., Monien, B., Rode, M., Spirakis, P., Vrtó, I.: Which is the Worst-case Nash Equilibrium? In: Rován, B., Vojtáš, P. (eds.) *MFCS 2003*. LNCS, vol. 2747, pp. 551–561. Springer, Heidelberg (2003)
10. Mavronicolas, M., Spirakis, P.: The Price of Selfish Routing. In: Mavronicolas, M., Spirakis, P. (eds.) *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, July 2001, pp. 510–519 (2001), Also, accepted to *Algorithmica*
11. Papadimitriou, C.H., Yannakakis, M.: On the Value of Information in Distributed Decision-Making. In: *Proceedings of the 10th Annual ACM Symposium on Principles of Distributed Computing*, pp. 61–64. ACM Press, New York (1991)
12. Yao, A.C.: Some Complexity Questions Related to Distributive Computing. In: *Proceedings of the 11th ACM Symposium on Theory of Computing (STOC 1979)*, pp. 209–213. ACM Press, New York (1979)