Selfish Load Balancing under Partial Knowledge

Panagiota N. Panagopoulou

Research Academic Computer Technology Institute
University of Patras, Greece
panagopp@cti.gr

joint work with Elias Koutsoupias and Paul Spirakis

AEOLUS Workshop on Scheduling
Nice, March 8–9, 2007

Outline

1. The Model
   - Agents and strategies
   - Selfish Costs and Nash equilibria
   - The Divergence Ratio
2. Zero Knowledge
   - All Nash equilibria
   - The Divergence Ratio
3. Arbitrary Knowledge
   - Bounding the Players' Optimum
   - Bounding the worst Social Cost
   - A lower bound on the Divergence Ratio
4. Full Knowledge
   - The Divergence Ratio

Outline

1 The Model
- Agents and strategies
- Selfish Costs and Nash equilibria
- The Divergence Ratio

2 Zero Knowledge
- All Nash equilibria
- The Divergence Ratio

3 Arbitrary Knowledge
- Bounding the Players' Optimum
- Bounding the worst Social Cost
- A lower bound on the Divergence Ratio

4 Full Knowledge
- The Divergence Ratio

Agents, loads and information

- A set $N = \{1, 2, \ldots, n\}$ of $n > 1$ selfish *agents*
- Each $i \in N$ has a *load* $w_i \in [0, 1]$
- Two bins (bin 0 and bin 1) of unbounded capacity
- Each agent has to select one of the two available bins to put her load
- For any $(i, j) \in N \times N$, agent i knows either
  - (a) the exact value of w_j or
  - (b) that w_j is uniformly distributed on $[0, 1]$
- Let $I_i = \{j \in N : \text{agent } i \text{ knows the exact value of } w_j\}$ and denote $\mathbf{I} = (I_i)_{i \in N}$

Single-threshold strategies

- A *strategy* for agent i is a function $s_i : [0,1] \rightarrow \{0,1\}$ such that $s_i(w_i)$ is the bin that agent i selects when her load is w_i.
- We only consider *single-threshold* strategies, i.e. a strategy for agent i is some $t_i \in [0,1]$ so that

$$s_i(w_i) = \begin{cases} 0 & w_i \leq t_i \\ 1 & w_i > t_i \end{cases}$$

- A *strategy profile* $\mathbf{t} = (t_1, \ldots, t_n) \in [0,1]^n$ is a combination of strategies, one for each agent.
- Denote by $(t'_i, \mathbf{t}_{-i})$ the strategy profile that is identical to $\mathbf{t}$ except for agent i, who chooses strategy t'_i instead of t_i.

Selfish Costs and Nash equilibria

The *Selfish Cost* $\text{Cost}_i(\mathbf{t}; I_i)$ of agent i is the expected load of the bin she selects, based on

1. her information about the exact loads of all $j \in I_i$ and
2. her knowledge that the loads of all $j \notin I_i$ are uniformly distributed on $[0, 1]$.

Selfish Costs and Nash equilibria

The *Selfish Cost* $\text{Cost}_i(\mathbf{t}; I_i)$ of agent i is the expected load of the bin she selects, based on

1. her information about the exact loads of all $j \in I_i$ and
2. her knowledge that the loads of all $j \notin I_i$ are uniformly distributed on $[0, 1]$.

In a *Nash equilibrium*, no agent can decrease her Selfish Cost by deviating:

Definition

The strategy profile $\mathbf{t} = (t_1, \ldots, t_n) \in [0, 1]^n$ is a Nash equilibrium if and only if, for all $i \in N$,

$$\text{Cost}_i(\mathbf{t}; I_i) \leq \text{Cost}_i\left((t_i', \mathbf{t}_{-i}); I_i\right) \quad \forall t_i' \in [0, 1].$$

The Divergence Ratio

- Associated with a strategy profile $\mathbf{t} \in [0,1]^n$ is the *Social Cost*:

$$\mathsf{SC}(\mathbf{t}, \mathbf{I}) = \max_{i \in N} \mathsf{Cost}_i(\mathbf{t}; I_i) \ .$$

- The *Players' Optimum* $\mathsf{PO}(\mathbf{I})$ is the minimum, over all possible strategy profiles $\mathbf{t} \in [0,1]^n$, Social Cost:

$$\mathsf{PO}(\mathbf{I}) = \min_{\mathbf{t} \in [0,1]^n} \mathsf{SC}(\mathbf{t}, \mathbf{I}) \ .$$

- The *Divergence Ratio* $\mathsf{DR}(\mathbf{I})$ is the maximum, over all Nash equilibria $\mathbf{t}$, of the ratio $\frac{\mathsf{SC}(\mathbf{t}, \mathbf{I})}{\mathsf{PO}(\mathbf{I})}$:

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

Outline

1. The Model
   - Agents and strategies
   - Selfish Costs and Nash equilibria
   - The Divergence Ratio
2. Zero Knowledge
   - All Nash equilibria
   - The Divergence Ratio
3. Arbitrary Knowledge
   - Bounding the Players' Optimum
   - Bounding the worst Social Cost
   - A lower bound on the Divergence Ratio
4. Full Knowledge
   - The Divergence Ratio

Zero Knowledge

Assume $I_i = \emptyset$ for all $i \in N$. Then

$$\mathrm{Cost}_i(\mathbf{t}; I_i) \quad =$$

Zero Knowledge

Assume $I_i = \emptyset$ for all $i \in N$. Then

$$\mathrm{Cost}_i(\mathbf{t}; I_i) = t_i \left(\frac{t_i}{2} + \sum_{j \neq i} t_j \frac{t_j}{2} \right)$$

Zero Knowledge

Assume $I_i = \emptyset$ for all $i \in N$. Then

$$\begin{aligned} \text{Cost}_i(\mathbf{t}; I_i) &= t_i \left(\frac{t_i}{2} + \sum_{j \neq i} t_j \frac{t_j}{2} \right) \\ &\quad + (1 - t_i) \left(\frac{t_i + 1}{2} + \sum_{j \neq i} (1 - t_j) \frac{t_j + 1}{2} \right) \end{aligned}$$

Zero Knowledge

Assume $I_i = \emptyset$ for all $i \in N$. Then

$$\begin{aligned} \mathrm{Cost}_i(\mathbf{t}; I_i) &= t_i \left(\frac{t_i}{2} + \sum_{j \neq i} t_j \frac{t_j}{2} \right) \\ &+ (1 - t_i) \left(\frac{t_i + 1}{2} + \sum_{j \neq i} (1 - t_j) \frac{t_j + 1}{2} \right) \end{aligned}$$

Proposition

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

Characterization of Nash equilibria

Since $\mathrm{Cost}_i(\mathbf{t}; I_i) = t_i \left(\sum_{j \neq i} t_j^2 - \frac{n-1}{2} \right) + \frac{n}{2} - \frac{1}{2} \sum_{j \neq i} t_j^2$:

Proposition

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 $i \in N$,

$$t_i = 0 \quad \Rightarrow \quad \sum_{j \neq i} t_j^2 \geq \frac{n-1}{2}$$

$$t_i = 1 \quad \Rightarrow \quad \sum_{j \neq i} t_j^2 \leq \frac{n-1}{2}$$

$$t_i \in (0,1) \quad \Rightarrow \quad \sum_{j \neq i} t_j^2 = \frac{n-1}{2}$$

All Nash equilibria

Theorem

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,* $\kappa = \frac{n-1}{2}$, $\lambda = 1$.

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

All Nash equilibria

Sketch of 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 $(|A|-1)t_A^2 + |C| = \frac{n-1}{2}$.
- For all $i \in B$, $t_i = 0$ and $|A|t_A^2 + |C| \geq \frac{n-1}{2}$.
- For all $i \in C$, $t_i = 1$ and $|A|t_A^2 + |C| - 1 \leq \frac{n-1}{2}$.

We consider the cases $|A| = 0, |A| = 1$ and $|A| > 1$ so as to find all Nash equilibria and calculate their Social Cost. □

The Divergence Ratio

Lemma

$\mathrm{PO}(\mathbf{I}) = \frac{n}{4}$ *if n is even and* $\mathrm{PO}(\mathbf{I}) = \frac{n+1}{4}$ *if n is odd.*

Therefore

Theorem

Consider the case where $I_i = \emptyset$ for all $i \in N$. Then

- $\mathrm{DR}(\mathbf{I}) = 1 + \frac{1}{n}$ *if n is even and*
- $\mathrm{DR}(\mathbf{I}) = 1$ *if n is odd.*

Outline

1. The Model
   - Agents and strategies
   - Selfish Costs and Nash equilibria
   - The Divergence Ratio
2. Zero Knowledge
   - All Nash equilibria
   - The Divergence Ratio
3. Arbitrary Knowledge
   - Bounding the Players' Optimum
   - Bounding the worst Social Cost
   - A lower bound on the Divergence Ratio
4. Full Knowledge
   - The Divergence Ratio

Arbitrary Knowledge

- Assume arbitrary I_i's for all $i \in N$.
- We will show that, if $i \in I_i$ and the cardinality of I_i is sufficiently small for all $i \in N$, then the divergence ratio can be as bad as n.

Arbitrary Knowledge

- Assume arbitrary I_i's for all $i \in N$.
- We will show that, if $i \in I_i$ and the cardinality of I_i is sufficiently small for all $i \in N$, then the divergence ratio can be as bad as n.

Sketch of Proof:

- Assume that $i \in I_i$ and $|I_i| \leq \frac{n-2}{3}$ for all $i \in N$.
- Consider the instance where $w_i = 1$ for all $i \in N$.
- Our goal is to find
  1. a Nash equilibrium $\mathbf{t}$ of low Social Cost, so as to upper bound the Players' Optimum, and
  2. a Nash equilibrium $\mathbf{t}'$ of high Social Cost, so as to lower bound the worst possible Social Cost.

Bounding the Players' Optimum

Consider the strategy profile $\mathbf{t}$ such that $t_i = 1 - \frac{1}{n-|I_i|}$ for all $i \in N$. Then

$$\mathrm{Cost}_i(\mathbf{t}; I_i) = |I_i| + 1 - \frac{1}{2(n-|I_i|)}$$

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

$$\begin{aligned} 1 + \frac{n-|I_i|}{2}\left(1 - \frac{1}{n-|I_i|}\right)^2 &\geq |I_i| + 1 + \frac{1}{2(n-|I_i|)} \\ &> \mathrm{Cost}_i(\mathbf{t}; I_i) \end{aligned}$$

The Social Cost of the Nash equilibrium $\mathbf{t}$ is

$$\mathrm{SC}(\mathbf{t}, \mathbf{I}) = \max_{i \in N} \mathrm{Cost}_i(\mathbf{t}; \mathbf{I}) \leq \max_{i \in N} |I_i| + 1$$

Bounding the worst Social Cost

Now consider the profile $\mathbf{t}'$ where $t'_i = \frac{\sqrt{8+16\frac{|I_i|-1}{n-|I_i|}}}{4}$ for all $i \in N$. Then

$$\text{Cost}_i(\mathbf{t}'; I_i) = \frac{n + |I_i| + 2}{4}$$

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

$$\begin{aligned} 1 + \frac{n - |I_i|}{2}(t'_i)^2 &= \frac{n + |I_i| + 2}{4} \\ &= \text{Cost}_i(\mathbf{t}'; I_i) \end{aligned}$$

The Social Cost the Nash equilibrium $\mathbf{t}'$ is

$$\text{SC}(\mathbf{t}', \mathbf{I}) = \max_{i \in N} \text{Cost}_i(\mathbf{t}'; \mathbf{I}) = \frac{n + \max_{i \in N} |I_i| + 2}{4}$$

A lower bound on the Divergence Ratio

Thus the Divergence Ratio is

$$\mathrm{DR}(\mathbf{l}) = \max_{\hat{\mathbf{t}}:\hat{\mathbf{t}} \text{ N.E.}} \frac{\mathrm{SC}(\hat{\mathbf{t}}, \mathbf{l})}{\mathrm{PO}(\mathbf{l})} \geq \frac{\mathrm{SC}(\mathbf{t}', \mathbf{l})}{\mathrm{SC}(\mathbf{t}, \mathbf{l})} \geq \frac{n + \max_{i \in N} |l_i| + 2}{4 \max_{i \in N} |l_i| + 4} .$$

A lower bound on the Divergence Ratio

Thus the Divergence Ratio is

$$\mathrm{DR}(\mathbf{I}) = \max_{\hat{\mathbf{t}}:\hat{\mathbf{t}}\ \mathrm{N.E.}} \frac{\mathrm{SC}(\hat{\mathbf{t}}, \mathbf{I})}{\mathrm{PO}(\mathbf{I})} \geq \frac{\mathrm{SC}(\mathbf{t}', \mathbf{I})}{\mathrm{SC}(\mathbf{t}, \mathbf{I})} \geq \frac{n + \max_{i \in N} |I_i| + 2}{4 \max_{i \in N} |I_i| + 4} .$$

Theorem

If $|I_i| \leq \frac{n-2}{3}$ *and* $i \in I_i$ *for all* $i \in N$*, then* $\mathrm{DR}(\mathbf{I}) \geq \frac{n + \max_i |I_i| + 2}{4 \max_i |I_i| + 4}$.

A lower bound on the Divergence Ratio

Thus the Divergence Ratio is

$$\mathrm{DR}(\mathbf{I}) = \max_{\hat{\mathbf{t}}:\hat{\mathbf{t}} \text{ N.E.}} \frac{\mathrm{SC}(\hat{\mathbf{t}}, \mathbf{I})}{\mathrm{PO}(\mathbf{I})} \geq \frac{\mathrm{SC}(\mathbf{t}', \mathbf{I})}{\mathrm{SC}(\mathbf{t}, \mathbf{I})} \geq \frac{n + \max_{i \in N} |I_i| + 2}{4 \max_{i \in N} |I_i| + 4} \;.$$

Theorem

If $|I_i| \leq \frac{n-2}{3}$ *and* $i \in I_i$ *for all* $i \in N$, *then* $\mathrm{DR}(\mathbf{I}) \geq \frac{n+\max_i |I_i|+2}{4 \max_i |I_i|+4}$.

Corollary

If $|I_i| = o(n)$ *and* $i \in I_i$ *for all* $i \in N$, *then* $\lim_{n \to \infty} \mathrm{DR}(\mathbf{I}) = \infty$.

Outline

1. The Model
   - Agents and strategies
   - Selfish Costs and Nash equilibria
   - The Divergence Ratio
2. Zero Knowledge
   - All Nash equilibria
   - The Divergence Ratio
3. Arbitrary Knowledge
   - Bounding the Players' Optimum
   - Bounding the worst Social Cost
   - A lower bound on the Divergence Ratio
4. Full Knowledge
   - The Divergence Ratio

Full Knowledge

- Assume that $I_i = N$ for all $i \in N$.
- The cost of $i \in N$ for a strategy profile $\mathbf{t} = (t_1, \ldots, t_n) \in [0,1]^n$ is

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

- It suffices to consider single-threshold strategies of the form $t_i = 0$ or $t_i = 1$, for all $i \in N$.

Theorem

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

The Divergence Ratio

Proof of the upper bound.

- Consider a Nash equilibrium $\mathbf{t}$. The total loads on the bins are $B_0(\mathbf{t}) = \sum_{i:t_i=1} w_i$ and $B_1(\mathbf{t}) = \sum_{i:t_i=0} w_i$. Assume that $B_0(\mathbf{t}) \geq B_1(\mathbf{t})$. Thus $\mathrm{SC}(\mathbf{t}, \mathbf{I}) = B_0(\mathbf{t})$.
- Moreover, $\mathrm{PO}(\mathbf{I}) \geq \frac{\sum_{i \in N} w_i}{2} = \frac{B_0(\mathbf{t}) + B_1(\mathbf{t})}{2}$.
- If only one agent places her load on bin 0 then $\mathrm{DR}(\mathbf{I}) = 1$.
- Otherwise, there exists an agent i who chooses bin 0 such that $w_i \leq \frac{B_0(\mathbf{t})}{2}$ implying that $B_0(\mathbf{t}) \leq B_1(\mathbf{t}) + \frac{B_0(\mathbf{t})}{2}$. Therefore,

$$\mathrm{DR}(\mathbf{I}) = \max_{\mathbf{t}:\mathbf{t}\ \text{N.E.}} \frac{\mathrm{SC}(\mathbf{t}, \mathbf{I})}{\mathrm{PO}(\mathbf{I})} \leq \frac{B_0(\mathbf{t})}{\frac{B_0(\mathbf{t}) + B_1(\mathbf{t})}{2}} \leq \frac{4}{3} \ .$$

□

The Divergence Ratio

Proof of Tightness.

- 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, $\mathrm{PO}(\mathbf{I}) = \frac{3}{2}(n-2)\alpha$ and thus

$$\mathrm{DR}(\mathbf{I}) \geq \frac{2(n-2)\alpha}{\frac{3}{2}(n-2)\alpha} = \frac{4}{3} \quad .$$

□

Summary

- If $I_i = \emptyset$ for all $i \in N$, then the Divergence Ratio is almost 1.
- If $|I_i|$ is constant for all $i \in N$, then the Divergence Ratio is lower bounded by n.
- If $|I_i| = o(n)$ for all $i \in N$, then the Divergence Ratio tends to infinity with n.
- If $I_i = N$ for all $i \in N$, then the Divergence Ratio is $\frac{4}{3}$.

Thank you