A Resource-Sharing Game with Relative Priorities

J. Doncel (jdoncel@laas.fr)
U. Ayesta, O. Brun and B. J. Prabhu

LAAS-CNRS
Toulouse, France

IFIP Performance
Turin, Italy
October 2014

Motivation

File hosting $\Rightarrow$ premium packages

- Higher payment, higher bandwidth

Motivation

File hosting ⇒ premium packages

- Higher payment, higher bandwidth

Cloud computing ⇒ pay for service

- Amazon EC2: Instances

Motivation

File hosting ⇒ premium packages

- Higher payment, higher bandwidth

Cloud computing ⇒ pay for service

- Amazon EC2: Instances
- Microsoft Azure, IBM, Google...

Motivation

File hosting $\Rightarrow$ premium packages

- Higher payment, higher bandwidth

Cloud computing $\Rightarrow$ pay for service

- Amazon EC2: Instances
- Microsoft Azure, IBM, Google...

$\Rightarrow$ Better service increasing payment

Outline

Outline

Model: system parameters

$\mathcal{C} = \{1, 2, \ldots, R\}$ set of players (classes of users) paying for service

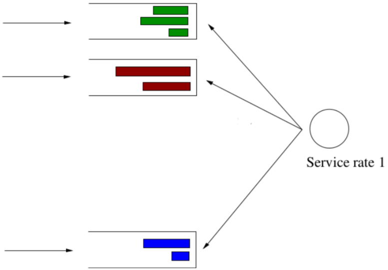

$\Rightarrow$ Quality of service of classes: function of processing time

Model: system parameters

$\mathcal{C} = \{1, 2, \ldots, R\}$ set of players (classes of users) paying for service

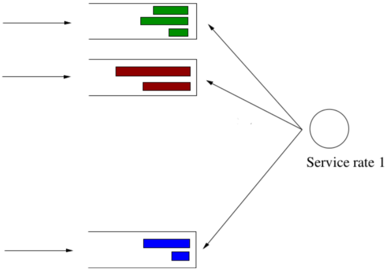

$\Rightarrow$ Quality of service of classes: function of processing time

- λ_i: arrival rate
- B_i: service requirement r. v.
- $\rho_i = \lambda_i \mathbb{E}(B_i)$: class-i load
- $T_i(\boldsymbol{g})$: response time of tasks of class i
- $\mathbb{E}(T_i(\boldsymbol{g})) = \overline{T}_i(\boldsymbol{g})$

Model: resource sharing

- *Strategy* of players: amount that they pay $g_i \in [\epsilon, \infty)$.
- Capacity is shared according to the discriminatory processor sharing (DPS).

Model: resource sharing

- *Strategy* of players: amount that they pay $g_i \in [\epsilon, \infty)$.
- Capacity is shared according to the discriminatory processor sharing (DPS).
- r_i: rate allocated to class i,

$$r_i(N_1, \dots, N_R) = \frac{g_i N_i}{\sum_{j=1}^{R} g_j N_j},$$

when there are N_j tasks of class j in the system.

Model: resource sharing

- *Strategy* of players: amount that they pay $g_i \in [\epsilon, \infty)$.
- Capacity is shared according to the discriminatory processor sharing (DPS).
- r_i: rate allocated to class i,

$$r_i(N_1, \ldots, N_R) = \frac{g_i N_i}{\sum_{j=1}^{R} g_j N_j},$$

  when there are N_j tasks of class j in the system.
- Higher payment $\Rightarrow$ higher rate (better service)

Model: resource sharing

- *Strategy* of players: amount that they pay $g_i \in [\epsilon, \infty)$.
- Capacity is shared according to the discriminatory processor sharing (DPS).
- r_i: rate allocated to class i,

$$r_i(N_1, \ldots, N_R) = \frac{g_i N_i}{\sum_{j=1}^{R} g_j N_j},$$

  when there are N_j tasks of class j in the system.
- Higher payment $\Rightarrow$ higher rate (better service)

Objective:

Minimize payment ensuring the QoS requirement

Known Results of DPS

Difficult model

- $\overline{T}_k(\boldsymbol{g})$ in close-form only for $R = 2$ and exp serv times.

Known Results of DPS

Difficult model

- $\overline{T}_k(\boldsymbol{g})$ in close-form only for $R = 2$ and exp serv times.
- for exponentially distributed required service times and $R \geq 2$:

$$\overline{T}_k(\boldsymbol{g}) \left(1 - \sum_{j=1}^{R} \frac{\lambda_j g_j}{\mu_j g_j + \mu_k g_k} \right) - \sum_{j=1}^{R} \frac{\lambda_j g_j \overline{T}_j(\boldsymbol{g})}{\mu_j g_j + \mu_k g_k} = \frac{1}{\mu_k}, \; k \in \mathcal{C}.$$

Known Results of DPS

Difficult model

- $\overline{T}_k(\boldsymbol{g})$ in close-form only for $R = 2$ and exp serv times.
- for exponentially distributed required service times and $R \geq 2$:

$$\overline{T}_k(\boldsymbol{g}) \left(1 - \sum_{j=1}^{R} \frac{\lambda_j g_j}{\mu_j g_j + \mu_k g_k} \right) - \sum_{j=1}^{R} \frac{\lambda_j g_j \overline{T}_j(\boldsymbol{g})}{\mu_j g_j + \mu_k g_k} = \frac{1}{\mu_k}, \; k \in \mathcal{C}.$$

- For non-exponentially distributed service times, $\overline{\mathbf{T}}(\boldsymbol{g})$ is a solution to a set of integro-differential equations (Fayolle et al 80).

Known Results of DPS

Difficult model

- $\overline{T}_k(\boldsymbol{g})$ in close-form only for $R = 2$ and exp serv times.
- for exponentially distributed required service times and $R \geq 2$:

$$\overline{T}_k(\boldsymbol{g}) \left(1 - \sum_{j=1}^{R} \frac{\lambda_j g_j}{\mu_j g_j + \mu_k g_k} \right) - \sum_{j=1}^{R} \frac{\lambda_j g_j \overline{T}_j(\boldsymbol{g})}{\mu_j g_j + \mu_k g_k} = \frac{1}{\mu_k}, \ k \in \mathcal{C}.$$

- For non-exponentially distributed service times, $\overline{\mathbf{T}}(\boldsymbol{g})$ is a solution to a set of integro-differential equations (Fayolle et al 80).
- when $\rho \to 1$,

$$(1-\rho)\ T_i(\boldsymbol{g}) \xrightarrow{d} T_i(\boldsymbol{g};1) = X \cdot \frac{\mathbb{E}(B_i)}{g_i}, \quad i \in \mathcal{C}, \tag{1}$$

where X is an exponentially distributed random variable

Game (OPT-M)

Each player i

$$\min_{g_i \geq \epsilon} \quad \rho_i g_i \qquad \text{(OPT-M)}$$
$$\text{subject to} \quad \overline{T}_i(\boldsymbol{g}) \leq c_i.$$

Game (OPT-M)

Each player i

$$\min_{g_i \geq \epsilon} \quad \rho_i g_i \qquad \text{(OPT-M)}$$
$$\text{subject to} \quad \overline{T}_i(\boldsymbol{g}) \leq c_i.$$

Question:

What price should a player pay?

Summary of Results

	Game (OPT-M)		Game (OPT-HT)	
	N. Classes	Serv. Times	N. Classes	Serv. Times
Feasibility	Arbitrary	Exponential	Arbitrary	General
Existence of NE	Arbitrary	General	Arbitrary	General
Uniqueness of NE	2	General	Arbitrary	General
NE Characterization	2	Exponential	Arbitrary	General
Price of Anarchy	2	General	Arbitrary	General
BR Convergence (feasible point)	Arbitrary	General	Arbitrary	General
BR Convergence (any point)	2	Exponential	2	General

Summary of main results

Outline

Existence of Equilibrium

Definition (Feasibility)

The game (OPT-M) is feasible if and only if it exists a performance vector such that $\overline{T}_i(\boldsymbol{g}) \leq c_i$, $i \in \mathcal{C}$.

Existence of Equilibrium

Definition (Feasibility)

The game (OPT-M) is feasible if and only if it exists a performance vector such that $\overline{T}_i(\boldsymbol{g}) \leq c_i$, $i \in \mathcal{C}$.

Proposition

With general service time distributions, if the game is feasible, then

- there exists a Nash Equilibrium
- the dynamics of best-response converge to a Nash Equilibrium if the starting point $\boldsymbol{g}$ satisfies that $\overline{T}_i(\boldsymbol{g}) \leq c_i$, $i \in \mathcal{C}$.

Existence of Equilibrium

Definition (Feasibility)

The game (OPT-M) is feasible if and only if it exists a performance vector such that $\overline{T}_i(\boldsymbol{g}) \leq c_i$, $i \in \mathcal{C}$.

Proposition

With general service time distributions, if the game is feasible, then

- there exists a Nash Equilibrium
- the dynamics of best-response converge to a Nash Equilibrium if the starting point $\boldsymbol{g}$ satisfies that $\overline{T}_i(\boldsymbol{g}) \leq c_i$, $i \in \mathcal{C}$.

Since:

- $\overline{T}_i(\boldsymbol{g})$ decreases with g_i and increases with g_j, $\forall j \neq i$
- $\overline{T}_i(c\boldsymbol{g}) = \overline{T}_i(\boldsymbol{g})$

Feasibility

For exponential service times

Proposition

The game (OPT-M) *is feasible if and only if*

$$\sum_{i\in r} \rho_i c_i \geq W_r, \ \forall r \text{ subset of } \mathcal{C}.$$

where $\bar{\rho}_r = \sum_{i\in r} \rho_i$ *and* $W_r = \frac{1}{1-\bar{\rho}_r} \sum_{i\in r} \frac{\rho_i}{\mu_i}$.

Characterization of the NE

Particular case

If $\exists$ $\boldsymbol{g}$ such that $\overline{T}_i(\boldsymbol{g}) = c_i \Rightarrow$ infinite equilibria

$$\boldsymbol{g}^{NE} = c\,\boldsymbol{g},\ \forall c.$$

Characterization of the NE

Particular case

If $\exists$ $\boldsymbol{g}$ such that $\overline{T}_i(\boldsymbol{g}) = c_i \Rightarrow$ infinite equilibria

$$\boldsymbol{g}^{NE} = c\,\boldsymbol{g},\ \forall c.$$

Two players and exponential service times

Let $\mathbb{E}(B_i) = 1/\mu_i$. Assume $c_1\mu_1 \leq c_2\mu_2$.
If the game is feasible, then the unique equilibrium is

- Let $\boldsymbol{g}^{PS} = (\epsilon, \epsilon)$. If $\overline{T}_i(\boldsymbol{g}^{PS}) \leq c_i$, then $\mathbf{g}^{\mathbf{NE}} = \boldsymbol{g}^{PS}$,
- otherwise, $\mathbf{g}^{\mathbf{NE}} = (g_1^{NE}, \epsilon)$, where $g_1^{NE} = \epsilon\,\frac{-\mu_1\rho_2+\mu_2(1-\rho_2)[\mu_1 c_1(1-\rho)-1]}{-\mu_1\rho_2-\mu_1(1-\rho_1)[\mu_1 c_1(1-\rho)-1]}$.

Example:

2 classes and exp serv times

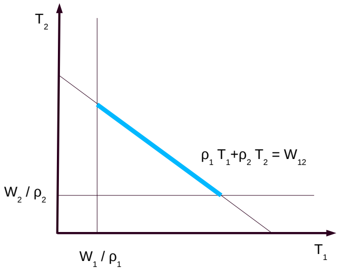

Figure: Set of performance vectors in a DPS queue

2 classes and exp serv times

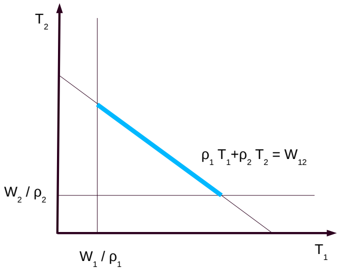

Figure: Set of performance vectors in a DPS queue

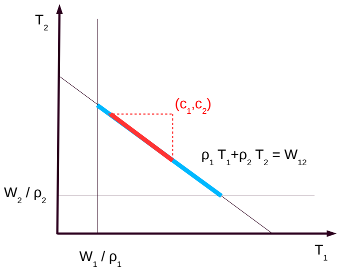

Figure: Set of performance vectors such that $\overline{T}_i(\boldsymbol{g}) \leq c_i$ (feasibility)

$$\textit{Feasibility} \iff \begin{array}{l} \rho_i c_i \geq W_i, i = 1, 2 \\ \rho_1 c_1 + \rho_2 c_2 \geq W_{12} \end{array}$$

Example: 2 classes and exp serv times

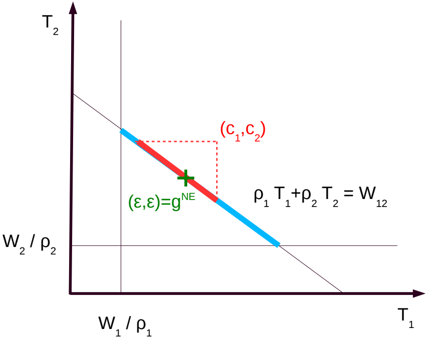

Figure: The case $\overline{T}_i(\boldsymbol{g}^{PS}) \leq c_i$.

Example: 2 classes and exp serv times

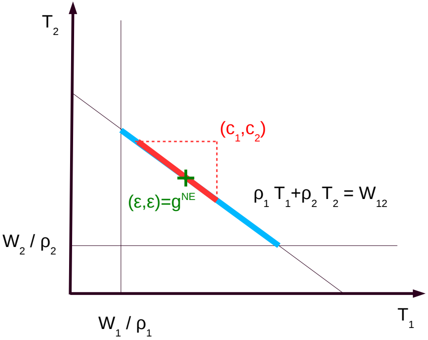

Figure: The case $\overline{T}_i(\boldsymbol{g}^{PS}) \leq c_i$.

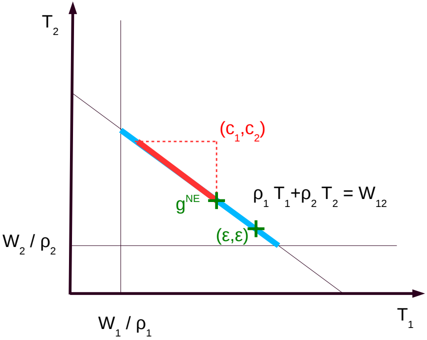

Figure: The case $\overline{T}_1(\boldsymbol{g}^{PS}) > c_1$.

Outline

The HT Game

- *Heavy traffic* : when the system load is close to 1.

The HT Game

- *Heavy traffic* : when the system load is close to 1.

Proposition (Verloop et al 2011)

When scaled with $1-\rho$, the response time of class-i jobs has a proper distribution as $\rho \to 1$.

$$(1-\rho)\ T_i(\boldsymbol{g}) \xrightarrow{d} T_i(\boldsymbol{g};1) = X \cdot \frac{\mathbb{E}(B_i)}{g_i}, \quad i \in \mathcal{C}, \tag{2}$$

where $\xrightarrow{d}$ denotes convergence in distribution and X is an exponentially distributed random variable with mean

$$\mathbb{E}(X) = \frac{\sum_k \lambda_k \mathbb{E}\left(B_k^2\right)}{\sum_k \lambda_k \mathbb{E}\left(B_k^2\right) \frac{1}{g_k}}. \tag{3}$$

The HT Game

- *Heavy traffic* : when the system load is close to 1.

Proposition (Verloop et al 2011)

When scaled with $1 - \rho$*, the response time of class-i jobs has a proper distribution as* $\rho \to 1$*.*

$$(1 - \rho)\ T_i(\boldsymbol{g}) \xrightarrow{d} T_i(\boldsymbol{g}; 1) = X \cdot \frac{\mathbb{E}(B_i)}{g_i}, \quad i \in \mathcal{C}, \tag{2}$$

where $\xrightarrow{d}$ *denotes convergence in distribution and* X *is an exponentially distributed random variable with mean*

$$\mathbb{E}(X) = \frac{\sum_k \lambda_k \mathbb{E}\left(B_k^2\right)}{\sum_k \lambda_k \mathbb{E}\left(B_k^2\right) \frac{1}{g_k}}. \tag{3}$$

Assume convergence in mean:

$$(1 - \rho)\overline{T}_i(\boldsymbol{g}; \rho) = \overline{T}_i(\boldsymbol{g}; 1) = \frac{\mathbb{E}(B_i)}{g_i} \frac{\sum_k \lambda_k \mathbb{E}\left(B_k^2\right)}{\sum_k \lambda_k \mathbb{E}\left(B_k^2\right) \frac{1}{g_k}}$$

Characterization of NE

Theorem

Assume $\mathbb{E}(B_i)/\tilde{c}_i$ decreasing with i, where $\tilde{c}_i = c_i\,(1-\rho)$.
If the game is feasible, the unique Nash equilibrium is

$$
\begin{aligned}
g_i^{NE} &= \epsilon \frac{\tilde{t}_m/\mathbb{E}(B_m)}{\tilde{c}_i/\mathbb{E}(B_i)}, \text{ for all } i < m, \\
g_i^{NE} &= \epsilon, \text{ for all } i \geq m,
\end{aligned}
$$

where m is the minimum value such that there exists $\tilde{t}_m \leq \tilde{c}_m$ verifying

$$
\frac{\tilde{t}_m}{\mathbb{E}(B_m)} = \frac{\sum_{k=1}^{R} \lambda_k \mathbb{E}\left(B_k^2\right) - \sum_{k=1}^{m-1} \lambda_k \frac{\mathbb{E}\left(B_k^2\right)}{\mathbb{E}(B_k)} \tilde{c}_k}{\sum_{k=m}^{R} \lambda_k \mathbb{E}\left(B_k^2\right)}.
$$

Arbitrary number of classes and general service times distribution.

Approximating (OPT-M)

Let $\mathbb{E}(B_i)/c_i \geq \mathbb{E}(B_j)/c_j$, if $i < j$.

Using $\overline{T}_i(\boldsymbol{g}) = \frac{\overline{T}_i(\boldsymbol{g};1)}{1-\rho} \Rightarrow$ Approximated NE:

Corollary

$$g_i^{NE} = \epsilon \frac{t_m/\mathbb{E}(B_m)}{c_i/\mathbb{E}(B_i)}, \text{ for all } i < m,$$

$$g_i^{NE} = \epsilon, \text{ for all } i \geq m,$$

where $m = 1, \ldots, R$ is the minimum value such that there exists a value $t_m \leq c_m$ verifying

$$\frac{t_m}{\mathbb{E}(B_m)} = \frac{\sum_{k=1}^{R} \frac{\lambda_k \mathbb{E}\left(B_k^2\right)}{(1-\rho)} - \sum_{k=1}^{m-1} \lambda_k \frac{\mathbb{E}\left(B_k^2\right)}{\mathbb{E}(B_k)} c_k}{\sum_{k=m}^{R} \lambda_k \mathbb{E}\left(B_k^2\right)}. \tag{4}$$

$$\tilde{c}_i = c_i\,(1-\rho),$$

$$\tilde{t}_m = t_m\,(1-\rho)$$

Outline

4 classes and exp serv times: homogeneous players

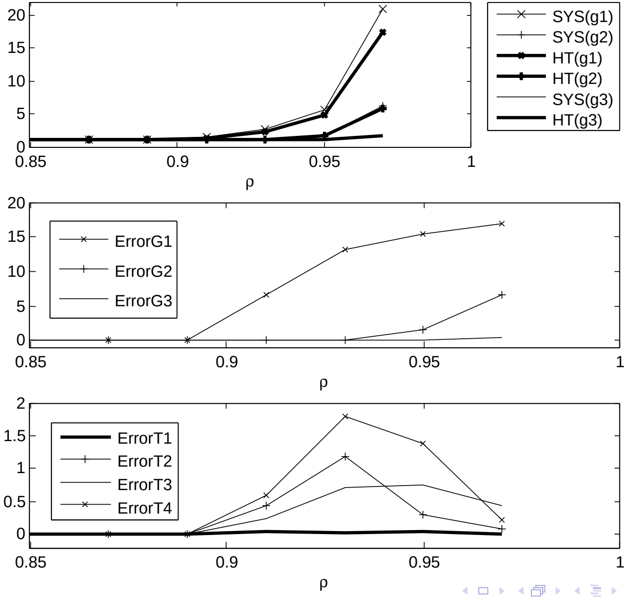

4 classes and exp serv times: heterogeneous players

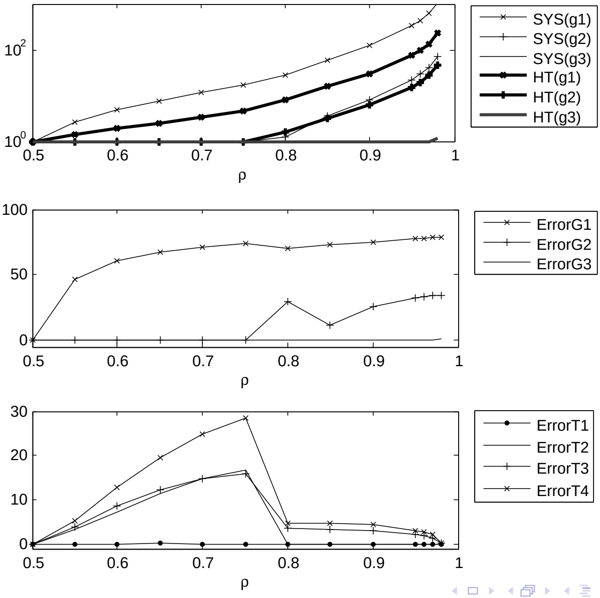

Outline

Conclusions and further work

Game theory and time-sharing systems: few previous work

Complicated model

- solved in some case
- for the rest, HT approximation

Future:

- Convergence of the Best Response
- Multiserver
- Users decreasing λ_i if $g_i > M$

Thank you

Thank you for your attention.