NUMERICAL STUDIES ON BRAESS-LIKE PARADOXES IN LOAD BALACING

Yoshihisa Hosokawa, Hisao Kameda and Odile Pourtallier

March 17, 2000

ISE-TR-00-167

H. Kameda is with the Institute of Information Sciences and Electronics, University of Tsukuba, Tsukuba Science City, Ibaraki 305-8573, Japan, Tel: +81-298-53-5539, Fax: +81-298-53-5206,
Email: kameda@is.tsukuba.ac.jp

Y. Hosokawa is with College of Information Sciences, University of Tsukuba, Tsukuba Science City, Ibaraki 305-8573, Japan, Tel & Fax: +81-298-53-5156,
Email: hosokawa@osdp.is.tsukuba.ac.jp

O. Pourtallier is with INRIA B.P. 93, 06902 Sophia Antipolis Cedex, France.
Tel: +33/0-492-387-826 Fax: +33/0-492-387-858
E-mail: Odile.Pourtallier@sophia.inria.fr.

Numerical Studies on Braess-like Paradoxes in Load Balancing

Yoshihisa Hosokawa*, Hisao Kameda† and Odile Pourtallier‡

Abstract

It has been reported that there exist a few numerical examples for the Braess-like paradox in which adding a certain capacity to the system may degrade the performance of all users. Unlike the original Braess paradox, this behavior occurs only in the case of finitely many users and not in the case of infinite number of users in the models examined. In this study, we examine a number of numerical examples around the Braess-like paradox such as above. In the numerical examples, it is observed that the Braess-like paradox is stronger (*i.e.*, the performance degradation of all users in the Brass-like paradox is larger) when the system has a higher degree of symmetry, in particular, in the parameter setting of the model whereby no forwarding of jobs occurs in the overall optimum.

Keywords: Braess paradox, Nash equilibrium, Wardrop equilibrium, performance optimization, numerical examples, distributed computer system, load balancing.

1 Introduction

We can choose between several distinct objectives for performance optimization in many systems including communication networks in distributed computer systems, transportation flow networks, *etc.* Among them, we have three typical objectives or optima :
(1) the overall optimum, system-optimum, cooperative optimum or social optimum, where a certain overall and single measure, like the total cost or the overall average response time, is to be optimized. We call it the *overall optimum* here.

*Y. Hosokawa is with Doctoral Program in Systems and Information Engineering, University of Tsukuba, Tsukuba Science City, Ibaraki 305-8573, Japan. Tel & Fax: +81-298-53-5156 E-mail: hosokawa@osdp.is.tsukuba.ac.jp .

†H. Kameda is with the Institute of Information Sciences and Electronics, University of Tsukuba, Tsukuba Science City, Ibaraki 305-8573, Japan. Tel: +81-298-53-5539 Fax: +81-298-53-5206 E-mail: kameda@is.tsukuba.ac.jp .

‡O. Pourtallier is with INRIA B.P. 93, 06902 Sophia Antipolis Cedex, France. Tel: +33/0-492-387-826 Fax: +33/0-492-387-858 E-mail: Odile.Pourtallier@sophia.inria.fr. The present work was done during her visit to the University of Tsukuba.

(2) the individual optimum, Wardrop equilibrium, or user optimum by some people; where each of infinitely many individuals, users, or jobs cannot receive any benefit by changing its own decision. Infinitely many users individually seek their own optimization. It is further assumed that decisions of a single individual have a negligible impact on the performance of other individuals. We call it the *individual optimum* or *Wardrop equilibrium* here.

(3) the class optimum, Nash non-cooperative equilibrium, or user optimum by some other people, where each of a finite number of classes, users, or players cannot receive any benefit by changing its decision. A finite number ($N(> 1)$) of players seek their own optimization non-cooperatively. We call it the *class optimum* or *Nash equilibrium* here.

Actually, (3) is reduced to (1) when the number of players reduces to 1 ($N = 1$) and approaches (2) when the number of players becomes infinitely many ($N \to \infty$) [4].

We can think that the total processing capacity of a system will increase when the capacity of a part of the system increases, and so we expect improvements in performance objectives accordingly in that case. The famous Braess paradox tells us that this is not always the case; *i.e.*, increased capacity of a part of the system may sometimes lead to the degradation in the benefits of all users in a Wardrop equilibrium [2, 3, 4]. We can expect that, in the Nash equilibrium where players seek their own optimization non-cooperatively, the similar type of paradox occurs (with large N), whenever it occurs for the Wardrop equilibrium ($N \to \infty$). Indeed, Korilis et al., [11, 12], found examples wherein the Braess-like paradox appears in a Nash equilibrium where all user classes are identical in the same topology for which the original Braess paradox (for the Wardrop equilibrium) was in fact obtained.

In [6], we presented the existence of a paradox similar to Braess's that appears in the class optimum (Nash equilibrium) but does not occur in the Wardrop equilibrium in the same environment. In this paper, we pursue the same line of investigation and show a number of numerical examples where such a paradoxe appears in class optimum but not in Wardrop. We present some tendencies or properties that can be seen from those numerical results. It is surprising that in these numerical results, for the class optimum, the Braess-like paradox appears more strongly in more symmetrical cases where no forwarding of jobs occurs in the overall and individual (Wardrop) optimum.

2 The Model and Assumptions

We consider a model consisting of two nodes (hosts) and a communication means that connects both nodes. Nodes are numbered 1 and 2. Each node consists of a single exponential server with service rate μ_i ($i = 1, 2$). We classify jobs arriving at node i into class i, $i = 1, 2$. Node i has the external Poisson arrival with rate ϕ_i, out of which the rate x_{ii} of jobs are processed at node i. The rate x_{ij} ($i \neq j$) of jobs are forwarded through the communication means to the other node j to be processed there, and the results of those jobs are returned back through the communication means to node i. Then we have $x_{ii} + x_{ij} = \phi_i$ ($i \neq j$), $x_{ij} \geq 0$, $\quad i, j = 1, 2$. We denote the vector $(x_{11}, x_{12}, x_{21}, x_{22})$ by $\boldsymbol{x}$. We denote the set of $\boldsymbol{x}$'s that satisfy the constraints by $\boldsymbol{C}$ and let $\Phi = \phi_1 + \phi_2$. Within these constraints, a set of values of x_{ij} $(i, j = 1, 2)$ are chosen to achieve optimization. The load on node i is $x_{ii} + x_{ji}$ $(i \neq j)$ and is denoted by β_i, and the expected processing (including queueing) time of a job that is processed at node i, is $1/(\mu_i - \beta_i)$ for $\beta_i < \mu_i$, and is otherwise is infinite.

As to the communication means, we consider the following two cases (A) and (B).

(A) It consists of two two-way communication lines 1 and 2. The two-way line i is used for forwarding of jobs that arrive at node i (and for sending back the processed results of these jobs).

We assume that the expected time length of forwarding and sending back a job is to be

$$G_i(x_{ij}) = \frac{t}{1 - x_{ij}t},$$

if $x_{ij}t < 1$, and is otherwise infinite.

That is, we assume each communication channel is modeled by a processor sharing server with service rate $1/t$; *i.e.*, the mean communication (without queueing) time is t, and thus, the capacity of each communication line is $1/t$.

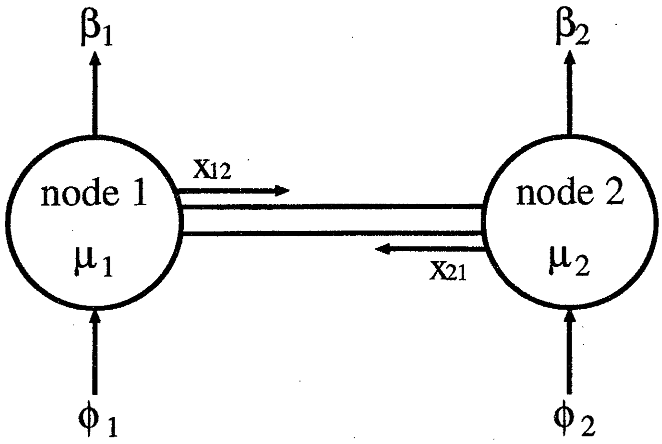

Figure 1: The system model (case (A)).

(B) It consists of a single-channel communication line that is used commonly in forwarding and sending back of jobs that arrive at both nodes.

The assumption on the line is the same as (A) except that there is only one line which is used for jobs arriving at both nodes. Thus the expected communication (with queueing) time of a job arriving at node i and being processed at node j ($\neq i$) is expressed as

$$G(\lambda) = \frac{t}{1 - (x_{12} + x_{21})t},$$

if $(x_{12} + x_{21})t < 1$, and is otherwise infinite, where $\lambda = x_{12} + x_{21}$ is the network traffic.

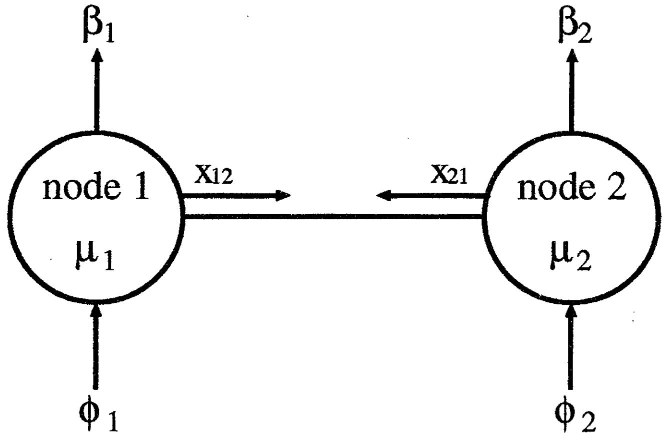

Figure 2: The system model (case (B)).

We refer to the length of time between the instant when a job arrives at a node and the instant when it leaves the node, where it has arrived, after all processing and communication (if any) are over, as *the response time for the job* arriving at the node.

Thus the expected response time of a job that arrives at node i is

$$T_i(\boldsymbol{x}) = \frac{1}{\phi_i} \sum_k x_{ik} T_{ik}(\boldsymbol{x}),$$

where

$$T_{ii}(\boldsymbol{x}) = \frac{1}{\mu_i - \beta_i},$$

if $\beta_i < \mu_i$ (and it is otherwise infinite), and for $j \neq i$,

$$\begin{aligned} T_{ij}(\boldsymbol{x}) &= \frac{1}{\mu_j - \beta_j} + \frac{t}{1 - x_{ij}t}, && \text{for case (A)}, \\ &= \frac{1}{\mu_j - \beta_j} + \frac{t}{1 - (x_{ij} + x_{ji})t}, && \text{for case (B)}. \end{aligned}$$

The above expressions hold, again, only for positive values of denominators, and are otherwise infinite.

Then, the overall expected response time of a job that arrives at the system is

$$T(\boldsymbol{x}) = \frac{1}{\Phi} \sum_i \phi_i T_i.$$

We have three optima, the overall, the individual, and the node.

(1) The overall optimum is given by such $\bar{\boldsymbol{x}}$ as satisfies the following,

$$T(\bar{\boldsymbol{x}}) = \min T(\boldsymbol{x}) \quad \text{with respect to} \quad \boldsymbol{x} \in \boldsymbol{C}. \tag{1}$$

(2) The individual optimum is given by such $\hat{\boldsymbol{x}}$ as satisfies the following for all i,

$$T_i(\hat{\boldsymbol{x}}) = \min\{T_{ii}(\hat{\boldsymbol{x}}), T_{ij}(\hat{\boldsymbol{x}})\} \ (i \neq j) \text{ such that } \hat{\boldsymbol{x}} \in \boldsymbol{C}. \tag{2}$$

(3) The class optimum is given by such $\tilde{\boldsymbol{x}}$ as satisfies the following for all i,

$$\tilde{T}_{ik} = T_{ik}(\tilde{\boldsymbol{x}}) = \min_{x_{ii}, x_{ij}} T_{ik}(\tilde{\boldsymbol{x}}_{-(ii,ij)}; x_{ii}, x_{ij}),$$
$$\text{such that } (\tilde{\boldsymbol{x}}_{-(ii,ij)}; x_{ii}, x_{ij}) \in \boldsymbol{C}. \tag{3}$$

where $(\tilde{\boldsymbol{x}}_{-(ii,ij)}; x_{ii}, x_{ij})$ denotes the $2n$ vector in which the elements corresponding to $\tilde{x}_{ii}$ and $\tilde{x}_{ij}$ has been replaced, respectively, by x_{ii} and x_{ij}.

In [9, 1] it is shown that the three problems (1), (2) and (3) are equivalent to some variational inequalities.

We are sure of the existence and the uniqueness of the overall, individual, and class optima for the model given here. For the existence and uniqueness of those optima see [5, 6].

Remark 2. 1 Note that, there should be no mutual forwarding in overall and individual optima. That is, in overall and individual optima, either one of x_{ij} $(i \neq j)$ must be zero in case (A) due to [16] or to Section 2.2.2 of [8] and in case (B) due to [13]. And thus, when one of x_{ij} $(i \neq j)$ say x_{ij} is non zero, $T_i(\boldsymbol{x})$ decreases and $T_j(\boldsymbol{x})$ increases with the change of t as shown in Theorems 2.5 and 2.7 of [8]. Consequently no paradox occurs for individual and overall optima.

Let us define three symmetries with respect to the system parameter setting.

[**Individual symmetry**] If the following condition holds

$$\frac{1}{\mu_1 - \phi_1} = \frac{1}{\mu_2 - \phi_2}, \tag{4}$$

it can be proved from the definition (2) that at the individual optimum there will be no forwarding between the two servers for any value of the communication line capacity $1/t$, for both cases (A) and (B). We say that we have an *individual symmetry* property between the two servers in this case.

[**Overall symmetry**] If we have

$$\frac{\mu_1}{(\mu_1 - \phi_1)^2} = \frac{\mu_2}{(\mu_2 - \phi_2)^2}, \tag{5}$$

then there is no forwarding when we are at the overall optimum, for both cases (A) and (B). This can be easily seen from the Kuhn-Tucker condition for the overall optimization problem (1). We say that we have an *overall symmetry* property between the two servers in this case.

[**Complete symmetry**] If both conditions (4) and (5) hold or equivalently if $\mu_1 = \mu_2$ and $\phi_1 = \phi_2$, then we say that we have a *complete symmetry* between the two servers. In the complete symmetry, no forwarding of jobs occurs both in the overall and individual optima, for both cases (A) and (B).

3 The numerical experiments

In the following numerical results, firstly we show some cases where complete symmetry holds. In these experiments we change the service rates of the servers while keeping the equalities between the two service rates and between the two arrival rates.

Secondly, we introduce some asymmetry. We start from a set of data such that the complete symmetry property holds.

- First we change both μ_1 and ϕ_1 in such a way that the property of individual symmetry, *i.e.*, condition (4) is preserved.
- Starting again from the complete symmetric case, we fix μ_2 and ϕ_2, change μ_1 and ϕ_1 in such a way that the overall property is preserved.
- Starting from the complete symmetric case, we fix μ_2, ϕ_1 and ϕ_2 and let μ_1 changes. In this case none of the property of symmetry holds anymore.

For each set of data, *i.e.*, for each vector $(\mu_1, \mu_2, \phi_1, \phi_2)$ we can find some value t^∞ (depending upon the set of data) of the mean communication time such that the communication line is not used anymore at equilibrium if the mean communication time is larger than t^∞. For each set of data $(\mu_1, \mu_2, \phi_1, \phi_2)$, we increase the mean communication time from 0 to t^∞. For each t we compute the class optimum (Nash equilibrium). In Fig. 3, we show the case where $(\mu_1, \mu_2, \phi_1, \phi_2) = (25, 25, 20, 20)$. In this case, $t^\infty = 0.8$ as is seen in the figure 3.

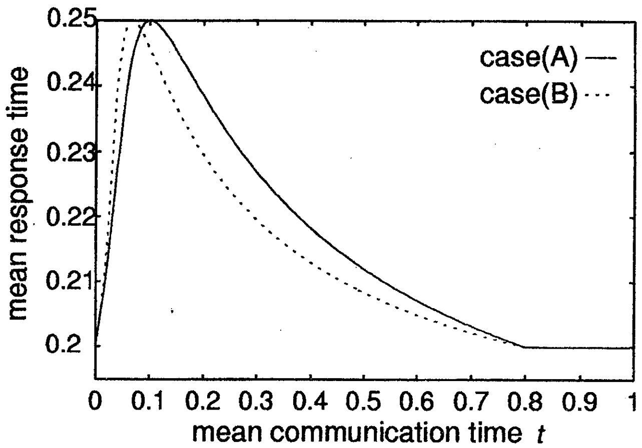

Figure 3: Mean response times at class optimum, for $(\mu_1, \mu_2, \phi_1, \phi_2) = (25, 25, 20, 20)$ (Complete symmetry) : $\tilde{T}_1 = \tilde{T}_2 = \tilde{T}$.

We use the following best reply algorithm to compute these Nash equilibria. For some fixed parameters $\phi_1, \phi_2, \mu_1, \mu_2, t$,

- Initialize $\boldsymbol{x}^0 = (x^0_{11}, x^0_{12}, x^0_{21}, x^0_{22}) \in \boldsymbol{C}$.

– Define $\boldsymbol{x}^n$ as

$$\begin{cases} (x_{11}^n, x_{12}^n) = \arg \min_{(x_{11},x_{12})} T_1(x_{11}, x_{12}, x_{21}^{n-1}, x_{22}^{n-1}) , \\ (x_{21}^n, x_{22}^n) = \arg \min_{(x_{21},x_{22})} T_2(x_{11}^n, x_{12}^n, x_{21}, x_{22}). \end{cases}$$

It can be shown that this algorithm converges for our model[8].

We focus our attention on the degradation that may occur when adding the communication line. To this aim we say that a Braess-like paradox occurs if the following holds:

$$\begin{array}{l} D_1(t_1, t_2) > 0 \quad \text{and} \quad D_2(t_1, t_2) > 0 \\ \text{for some} \quad t_1, t_2 \quad \text{such that } 0 < t_1 < t_2, \end{array} \tag{6}$$

where $D_i(t_1, t_2) = \dfrac{\tilde{T}_i(t_1) - \tilde{T}_i(t_2)}{\tilde{T}_i(t_2)}$, and $\tilde{T}_i(t)$ denotes the mean response for class i jobs, computed at the (unique) Nash equilibrium, when the mean communication time is t.

For simplicity, we only consider the case where $t_2 = t_\infty$, *i.e.*, the system has no communication means, and we denote $D_i(t, t^\infty)$ by $\Delta_i(t)$.

Thus, we define the *worst ratio of performance degradation* in the paradox $\Gamma(\mu_1, \mu_2, \phi_1, \phi_2)$ as follows:

$$\Gamma(\mu_1, \mu_2, \phi_1, \phi_2) = \max_t \min\{\Delta_1(t), \Delta_2(t)\}. \tag{7}$$

Denote t_0 the mean communication time, such that the previous maximum is attained.

Notice that we can define a weaker paradox. For example we can think of some kind of *local paradox*, such that we can observe for some interval of t, $\tilde{T}_i(t)$, for some i, increases, but nevertheless the overall mean response time $\tilde{T}(t) = (1/\Phi)(\phi_1\tilde{T}_1(t) + \phi_2(t)\tilde{T}_2)$ decreases. This weaker paradox may occur even if (6) does not hold, and has been already discussed in [7].

4 Results and Discussion

We show typical numerical examples, in the case (A) of the communication means, with changing parameter values from $\mu_1 = \mu_2 = 25$ and $\phi_1 = \phi_2 = 20$ for each of four directions: complete, overall, individual, and no symmetry maintained. For case (B), we had similar observations, a part of which is also shown below.

4.1 Complete symmetry maintained

In the following table, we show the effect of changing $\mu_1 = \mu_2$ from 25 (keeping $\phi_1 = \phi_2 = 20$ fixed) with complete symmetry maintained. "No P." means the situation where no paradox is seen for any interval of t.

μ	20.1	22	25	45	65	100	200
$\Delta_i(t_0)$ (%)	35	30	25	9	5	3	No P.
Γ(%)	35	30	25	9	5	3	No P.

As an example, we show the figure for the case where $(\mu_1, \mu_2, \phi_1, \phi_2) = (45, 45, 20, 20)$. Compare this figure 4 with Fig. 3.

The results are almost the same for both cases (A) and (B).

4.2 Overall symmetry maintained

In the following table, we show for case (A) the effect of changing μ_1 and ϕ_1 from 25 and 20, respectively (keeping $\mu_2 = 25$ and $\phi_2 = 20$ fixed) with overall symmetry maintained, *i.e.*, such that $\frac{\mu_1}{(\mu_1 - \phi_1)^2} = \frac{\mu_2}{(\mu_2 - \phi_2)^2} = 1$.

μ_1	4	9	25	64	121	400	40000
ϕ_1	2	6	20	56	110	380	39800
$\Delta_1(t_0)$ (%)	3	14	25	24	19	10	0.1
$\Delta_2(t_0)$ (%)	10	19	25	21	19	15	10
Γ (%)	3	14	25	21	19	10	0.1

As an example, we show the figure for the case where $(\mu_1, \mu_2, \phi_1, \phi_2) = (64, 25, 56, 20)$. Compare this figure 5 with Fig. 3.

The results for case (B) are quite similar to those for case (A) as shown above.

4.3 Individual symmetry maintained

In the following table, we show the effect of changing μ_1 and ϕ_1 from 25 and 20, respectively (keeping $\mu_2 = 25$ and $\phi_2 = 20$ fixed) with individual symmetry maintained, *i.e.*, such that $\frac{1}{\mu_1 - \phi_1} = \frac{1}{\mu_1 - \phi_1} = \frac{1}{5}$.

μ_1	6	15	25	65	105	125	505
ϕ_1	1	10	25	60	100	120	500
$\Delta_1(t_0)$(%)	41	32	25	9	2	1	No P.
$\Delta_2(t_0)$(%)	-8	12	25	50	62	66	
Γ(%)	-8	12	25	9	2	1	No P.

As an example, we show the figure for the case where $(\mu_1, \mu_2, \phi_1, \phi_2) = (65, 25, 60, 20)$. Compare this figure 6 with Fig. 3.

The results for case (B) are quite similar to those for case (A) as shown above.

4.4 No symmetry maintained

In the following tables, we show, for both cases (A) and (B), the effect of changing μ_1 from 25 (keeping $\mu_2 = 25$ and $\phi_1 = \phi_2 = 20$ fixed), *i.e.*, with no symmetry maintained.

Case (A)

μ_1	20.5	23	25	27	31	33
ϕ_1	20	20	20	20	20	20
Δ_1 (t_0)(%)	No P.	2	25	36	32	No P.
Δ_2 (t_0)(%)		42	25	8	-14	
Γ (%)	No P.	2	25	8	-14	No P.

Case (B)

μ_1	20.5	23	25	27	31	33
ϕ_1	20	20	20	20	20	20
Δ_1 $(t_0)(\%)$	No P.	2	25	39	45	No P.
Δ_2 $(t_0)(\%)$		42	25	6	-18	
Γ (%)	No P.	2	25	6	-18	No P.

As an example, we show the figure for the case (A) where $(\mu_1, \mu_2, \phi_1, \phi_2) = (33, 25, 20, 20)$. Compare this figure 7 with Fig. 3.

Overall tendency

In these examples, we see that the worst ratio of performance degradation in the paradox is the largest in the complete symmetry case with the arrival rate being closest to the service rate, *i.e.*, the case where $(\mu_1, \mu_2, \phi_1, \phi_2) = (25, 25, 20, 20)$. The degradation decreases most rapidly as the parameter setting departs the above-mentioned symmetric case without keeping any kind of symmetries. It decreases much slower with complete symmetry maintained and also with individual symmetry maintained. It decreases most slowly with overall symmetry maintained where no forwarding of jobs occurs in the overall optimum. It looks quite strange that in these symmetrical cases, in particular, in the overall symmetry, the chances of Braess-like paradox are greater than non-symmetrical situations.

5 Concluding remarks

We have presented a number of numerical examples for the Braess-like paradox wherein adding a resource to a system leads to the performance degradation for all users in the class optimum for load balancing. We have shown the examples in the model of load balancing between two distinct servers. it is surprising that the paradox appears the most strongly in the model of identical servers with each identical service rate in the class optimum (Nash equilibrium) while we cannot expect that such a paradox occurs neither in the overall optimum nor in the individual optimum (Wardrop equilibrium) in such situations. It is also amazing that, in the model of asymmetric servers, the paradox in the class optimum keeps to appear in the case of overall symmetry wherein no forwarding of loads occurs in the overall optimum. We think that more exhaustive research into these problems is worth pursuing in order to gain insight into the optimal design and QoS (quality of service) management of distributed computer systems, communication networks, etc.

References

[1] E. Altman and H. Kameda, "Uniqueness of Solutions for Optimal Static Routing in Multi-class Open Networks," Research report ISE-TR-98-156, Institute of Information Sciences and Electronics, University of Tsukuba, 1998.

[2] J. E. Cohen and F. P. Kelly, "A Paradox of Congestion in a Queuing Network," *J. Appl. Prob. 27*, pp.730-734, 1990.

[3] J. E. Cohen and C. Jeffries, "Congestion Resulting from Increased Capacity in Single-Server Queueing Networks," *IEEE/ACM Trans. on Networking 5*, 2, April pp.1220-1225, 1997.

[4] A. Haurie and P. Marcotte, "On the Relationship between Nash-Cournot and Wardrop Equilibria," *Networks 15*, 1985, pp.295-308.

[5] H. Kameda, E. Altman and T. Kozawa, "Braess-like Paradoxes of Nash Equilibria for Load Balancing in Distributed Computer Systems," ISE Technical Report ISE-TR-98-157, Institute of information Sciences and Electronics, University of Tsukuba, 1998.

[6] H. Kameda, E. Altman and T. Kozawa, "A Case where a Paradox Like Braess's Occurs in the Nash Equilibrium but Does Not Occur in the Wardrop Equilibrium — A Situation of Load Balancing in Distributed Computer Systems —," *Proc. IEEE Conference on Decision & Control (Arizona)*, 1999, pp 4014-4019.

[7] H. Kameda, T. Kozawa and J. Li, "Anomalous Relations among Various Performance Objectives in Distributed Computer Systems," *Proc. 1st World Congress on Systems Simulation*, IEEE, 1997, pp.459-465.

[8] H. Kameda, J. Li, C. Kim and Y. Zhang, *Optimal Load Balancing in Distributed Computer Systems*, Springer, 1997.

[9] H. Kameda and Y. Zhang, "Uniqueness of the Solution for Optimal Static Routing in Open BCMP Queueing Networks," *Mathematical and Computer Modelling 22*, 10-12, pp.119-130, 1995.

[10] C. Kim and H. Kameda, "An Algorithm for Optimal Static Load Balancing in Distributed Computer Systems," *IEEE Trans. Compt. 41*, 3, 1990, pp.381-384.

[11] Y. A. Korilis, A. A. Lazar and A. Orda, "Architecting Noncooperative networks", *IEEE Journal on Selected Areas in Communications 13*, 7, September 1995, pp.1241-1251. Noncooperative Routing", *IEEE Trans. Automatic Control 42*, 3, March 1997, pp.309-871.

[12] Y. A. Korilis, A. A. Lazar and A. Orda, "Avoiding the Braess Paradox in Noncooperative Networks", *Proc. IEEE Conference on Decision & Control (San Diego)*, 1997, pp.864-878.

[13] J. Li and H. Kameda, "Load Balancing Problems for Multiclass Jobs in Distributed/Parallel Computer Systems," *IEEE Trans. Compt. 47*, 3, 1998, pp.322-1998.

[14] A. Orda, R. Rom and N. Shimkin, "Competitive Routing in Multiuser Communication Networks," *IEEE/ACM Trans. on Networking 1*, 1993, pp.614-627.

[15] J. B. Rosen, "Existence and Uniqueness of Equilibrium Points for Concave N-person Games", *Econometrica 33*, pp. 153-163, 1965.

[16] A. N. Tantawi and D. Towsley, "Optimal Static Load Balancing in Distributed Computer Systems," *J. ACM 32*, 2, pp.445-465, Apr, 1985.

[17] Y. Zhang, H. Kameda and K. Shimizu, "Parametric Analysis of Optimal Load Balancing in Distributed Computer Systems," *Journal of Information Processing 14*, 4, 1992, pp.433-441.

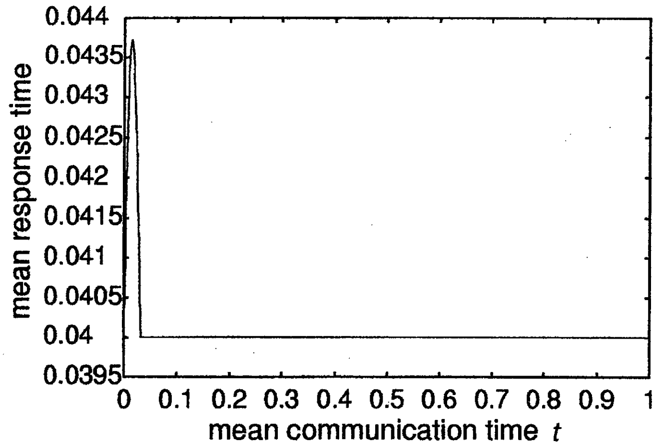

Figure 4: Mean response times at class optimum, for $(\mu_1, \mu_2, \phi_1, \phi_2) = (45, 45, 20, 20)$: Complete symmetry. $T_1 = T_2 = T$.

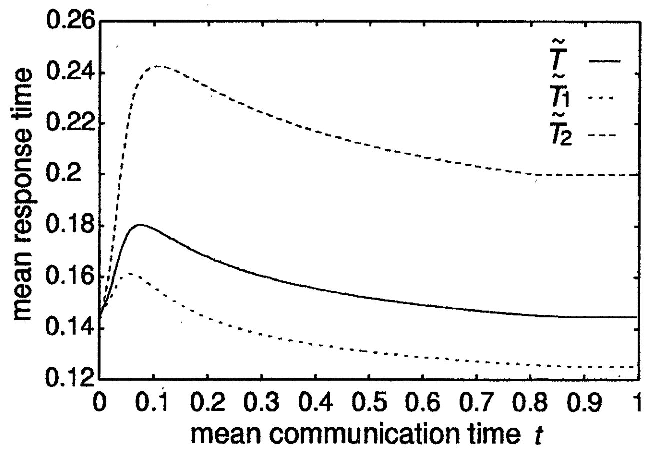

Figure 5: Mean response times at class optimum, for $(\mu_1, \mu_2, \phi_1, \phi_2) = (64, 25, 56, 20)$: Overall symmetry.

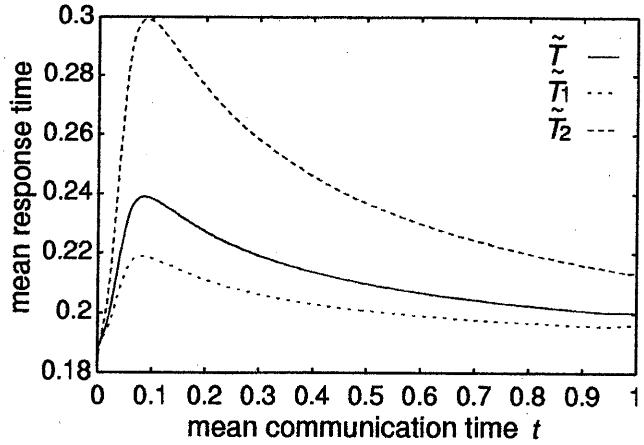

Figure 6: Mean response times at class optimum, for $(\mu_1, \mu_2, \phi_1, \phi_2) = (65, 25, 60, 20)$: Individual symmetry.

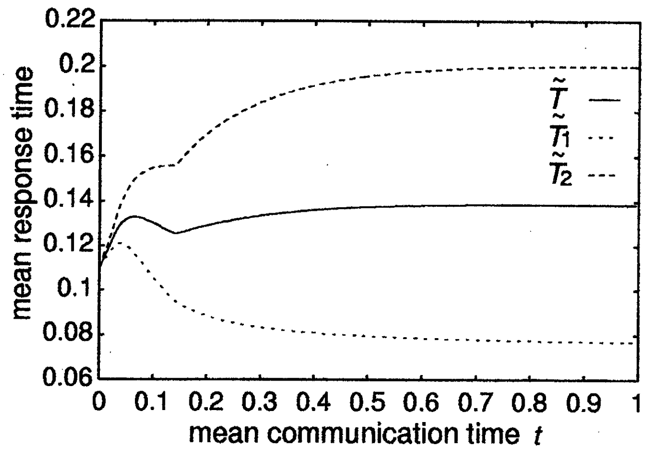

Figure 7: Mean response times at class optimum, for $(\mu_1, \mu_2, \phi_1, \phi_2) = (33, 25, 20, 20)$: No symmetry.