J. Appl. Prob. **36,** 211–222 (1999)
Printed in Israel

AVOIDING THE BRAESS PARADOX IN NON-COOPERATIVE NETWORKS

YANNIS A. KORILIS,* *Lucent Technologies*
AUREL A. LAZAR,** *Columbia University*
ARIEL ORDA,*** *Technion*

Abstract

The exponential growth of computer networking demands massive upgrades in the capacity of existing networks. Traditional capacity design methodologies, developed with the single-class networking paradigm in mind, overlook the non-cooperative structure of modern networks. Consequently, such design approaches entail the danger of degraded performance when resources are added to a network, a phenomenon known as the Braess paradox.

The present paper proposes methods for adding resources efficiently to a non-cooperative network of general topology. It is shown that the paradox is avoided when resources are added across the network, rather than on a local scale, and when upgrades are focused on direct connections between the sources and destinations. The relevance of these results for modern networks is demonstrated.

Keywords: Braess Paradox; computer communication networks; non-cooperative games; routeing.

AMS 1991 Subject Classification: Primary 60K30
Secondary 60M10; 90B12

1. Introduction

The exponential growth of computer networking, in terms of number of users and components, traffic volume and diversity of services, demands massive upgrades in capacity for existing networks. Traditionally, capacity design methodologies have been developed with the single-class networking paradigm in mind. This approach overlooks the non-cooperative structure of modern networks and entails, as will be explained in the following, the danger of degraded performance when resources are added to a network. The term *non-cooperative* is used to characterize networks operated according to a decentralized control paradigm, where control decisions are made by each user independently, according to its own individual performance objectives. The term 'user' may refer to a network user itself or, if the user's traffic consists of multiple connections, to individual connections that are controlled independently. The most common example of a non-cooperative network is the Internet. In the current TCP flow control mechanism, each user adjusts its transmission window—the maximum number

Received 3 January 1997; revision received 4 March 1998.
* Postal address: Bell Laboratories, Lucent Technologies, Holmdel, NJ 07733 USA.
** Postal address: Department of Electrical Engineering, Columbia University, New York, NY 10027 USA.
*** Postal address: Department of Electrical Engineering, Technion, Haifa 32000, Israel.
Email address: ariel@ee.technion.ac.il

of unacknowledged packets that the user can have circulating in the network—independently, based on some feedback information about the level of congestion in the network (detected as packet loss). Moreover, the Internet Protocol provides the option of source routeing, that enables the user to determine the path(s) its flow follows from source to destination.

Game theory [16] provides the systematic framework to study and understand the behavior of non-cooperative networks; see [1, 13, 14, 18, 19] and references therein. The operating points of a non-cooperative network are the *Nash equilibria* of the underlying game, that is, the points where unilateral deviation does not help any user to improve its performance. Non-cooperative equilibria are generically Pareto inefficient [9], as is traditionally illustrated in the well-known prisoner's dilemma [16]. In the context of non-cooperative networks, this inefficiency manifests itself as the potential degradation of performance when resources are added to the network. This non-intuitive behavior (under the single-class networking approach) is typically referred to as the Braess paradox [5]. Like many other paradoxes, however, the Braess paradox is a paradox in name only as will be explained in Section 3. Throughout the paper, the term 'paradox' is used in a non-strict sense.

The Braess paradox indicates that traditional practices, which overlook the non-cooperative structure of networks, can tentatively lead to degraded performance. Design methodologies for a class of non-cooperative networks have been studied in [14], using routeing as a control paradigm. They focus on networks consisting of non-interfering paths, establishing that the paradox cannot occur in such topologies, and obtaining the optimal solution for the capacity allocation problem. The problem of efficiently adding resources to general topology networks, however, is still open.

The present paper proposes methods for adding resources to a general network that guarantee an improvement in performance, thus establishing a methodology for efficiently coping with the Braess paradox in non-cooperative networks. Although our study was originally motivated by design problems in the field of computer networking, the results may be applied to other types of networks in which the paradox has also been observed, e.g. transportation networks [8]. Some of the results presented in this paper were briefly outlined, without proof, in a survey of our research on non-cooperative networks [12].

The paper is structured as follows. In Section 2 we present the non-cooperative network model and formulate the problem. The Braess paradox, in the context of non-cooperative routeing, is presented in Section 3. The proposed methods for capacity addition are studied in Section 4. Finally, concluding remarks are presented in Section 5.

2. Model and preliminaries

We consider a network $(\mathcal{V}, \mathcal{L})$, where $\mathcal{V}$ is a finite set of nodes and $\mathcal{L} \subseteq \mathcal{V} \times \mathcal{V}$ is a set of directed links. For simplicity of notation and without loss of generality, we assume that at most one link exists between each pair of nodes (in each direction). For any link $l = (u, v) \in \mathcal{L}$, define $S(l) = u$ and $D(l) = v$. Considering a node $v \in \mathcal{V}$, let $\mathrm{In}(v) = \{l : D(l) = v\}$ denote the set of its in-going links, and $\mathrm{Out}(v) = \{l : S(l) = v\}$ the set of its out-going links. Let c_l be the capacity of link l, where $\boldsymbol{c} = (c_l)_{l \in \mathcal{L}}$ is called the *capacity configuration* of the network.

A set $\mathcal{I} = \{1, 2, \ldots, I\}$ of users share the network $(\mathcal{V}, \mathcal{L})$. We shall assume that all users ship flow from a common source s to a common destination d. Each user i has a throughput demand that is some process with average rate r^i. User i splits its demand r^i among the paths connecting the source to the destination, so as to optimize some individual performance objective. Let f_l^i denote the expected flow that user i sends on link l. The user flow configuration $\boldsymbol{f}^i = (f_l^i)_{l \in \mathcal{L}}$ is called a routeing *strategy* of user i. The set of strategies for user i that satisfy the user's demand and preserve its flow at all nodes is called the strategy space of user i and denoted F^i, that is

$$F^i = \{\boldsymbol{f}^i \in \mathbb{R}^{|\mathcal{L}|} : 0 \le f_l^i \le c_l,\ l \in \mathcal{L};\ \sum_{l \in \mathrm{Out}(v)} f_l^i = \sum_{l \in \mathrm{In}(v)} f_l^i + r_v^i,\ v \in \mathcal{V}\},$$

where $r_s^i = r^i$, $r_d^i = -r^i$ and $r_v^i = 0$ for $v \neq s, d$. The system flow configuration $\boldsymbol{f} = (\boldsymbol{f}^1, \ldots, \boldsymbol{f}^I)$ is called a routeing *strategy profile* and takes values in the product strategy space $F = \otimes_{i \in \mathcal{I}} F^i$.

The grade of service that the flow of user i receives is quantified by means of a cost function $J^i : F \to \mathbb{R}$. $J^i(\boldsymbol{f})$ is the cost of user i under strategy profile $\boldsymbol{f}$; where a higher $J^i(\boldsymbol{f})$ means that a lower grade of service is provided to the flow of the user. We consider cost functions that are the sum of link cost functions

$$J^i(\boldsymbol{f}) = \sum_{l \in \mathcal{L}} f_l^i T_l(f_l), \tag{2.1}$$

where $\boldsymbol{f}_l = (f_l^1, \ldots, f_l^I)$, and $T_l(f_l)$ is the average delay on link l, and depends only on the total flow $f_l = \sum_{i \in \mathcal{I}} f_l^i$ on that link. The average delay should be interpreted as a general *congestion cost* per unit of flow, that encapsulates the dependence of the quality of service provided by a finite capacity resource on the total load f_l offered to it. In the present paper, we concentrate on congestion costs of the form

$$T_l(f_l) = \begin{cases} (c_l - f_l)^{-1}, & f_l < c_l, \\ \infty, & f_l \ge c_l, \end{cases} \tag{2.2}$$

that are typical of various practical routeing algorithms [4]. Note that Equation (2.2) describes the $M/M/1$ delay function. Therefore, if we assume that the delay characteristics of each link can be approximated by an $M/M/1$ queue, then $J^i(\boldsymbol{f})/r^i$ is the average time-delay that the flow of user i experiences under strategy profile $\boldsymbol{f}$.

User i aims to find a strategy $\boldsymbol{f}^i \in F^i$ that minimizes its cost. This optimization problem depends on the routeing decisions of the other users, described by the strategy profile $\boldsymbol{f}^{-i} = (\boldsymbol{f}^1, \ldots, \boldsymbol{f}^{i-1}, \boldsymbol{f}^{i+1}, \ldots, \boldsymbol{f}^I)$, since J^i is a function of the system flow configuration $\boldsymbol{f}$. A *Nash equilibrium* of the routeing game is a strategy profile from which no user finds it beneficial to unilaterally deviate. Hence, $\boldsymbol{f} \in F$ is a Nash equilibrium if

$$\boldsymbol{f}^i \in \arg\min_{\boldsymbol{g}^i \in F^i} J^i(\boldsymbol{g}^i, \boldsymbol{f}^{-i}), \quad i \in \mathcal{I}. \tag{2.3}$$

For the cost function $J^i(\boldsymbol{f})$ given by (2.1) and (2.2), the existence of a Nash equilibrium has been established in [18]. For the routeing problem to be meaningful, we assume that the network can accommodate the total offered load, i.e. there is a routeing strategy profile $\boldsymbol{f} \in F$ that is stable, in the sense that $f_l < c_l$ holds at all links $l \in \mathcal{L}$. The existence of such a stable strategy may be verified by solving a standard multi-commodity flow problem through linear programming techniques. Equations (2.3) and (2.2) then guarantee that, at any Nash equilibrium, we have $f_l < c_l$ for all $l \in \mathcal{L}$, and the costs of all users are finite.

Given a strategy profile $\boldsymbol{f}^{-i}$ for the other users, the cost of user i, as defined by (2.1) and (2.2), is a convex function of its strategy $\boldsymbol{f}^i$. Hence, the minimization problem in (2.3) has a unique solution. The Kuhn-Tucker optimality conditions [15], then, imply that $\boldsymbol{f}^i$ is the optimal response of user i to $\boldsymbol{f}^{-i}$ if and only if there exist (Lagrange multipliers) $(\lambda_u^i)_{u\in\mathcal{V}}$, such that

$$\lambda_u^i = f_{uv}^i T'_{uv} + T_{uv} + \lambda_v^i, \quad \text{if } f_{uv}^i > 0, \quad (u,v) \in \mathcal{L}, \tag{2.4}$$

$$\lambda_u^i \le f_{uv}^i T'_{uv} + T_{uv} + \lambda_v^i, \quad \text{if } f_{uv}^i = 0, \quad (u,v) \in \mathcal{L}, \tag{2.5}$$

$$\lambda_d^i = 0. \tag{2.6}$$

(The conditions (2.4)–(2.6) are presented in a form that is standard in the networking literature; for completeness, their derivation is presented in the Appendix.) Therefore, a strategy profile $\boldsymbol{f} \in F$ is a Nash equilibrium if and only if there exist λ_u^i, such that the optimality conditions (2.4)–(2.6) are satisfied for all $i \in \mathcal{I}$. Note that λ_s^i is, in fact, the marginal cost of user i at the optimality point. In accordance with the economics terminology, λ_s^i will be referred to as the *price* of user i [17].

We consider the following design problem. Given a network with some initial capacity configuration, the network designer can distribute some additional capacity allowance among the network links. The aim of the designer is to come up with a new capacity configuration that improves performance at the corresponding Nash equilibrium according to certain criteria. The problem is well-defined if the Nash equilibrium, under any capacity configuration, is unique. Whether this property holds in general topologies is an open question. (In [18], an example of a general topology with multiple Nash equilibria is presented, however, the cost functions are not of the form specified in (2.1) and (2.2).) Thus, we shall concentrate on some special cases of interest (presented in the following) for which uniqueness has been established.

The designer may have different measures for characterizing the efficiency of a capacity configuration. We shall concentrate on measures that are expressed by means of either the user prices or costs. Although the user's cost is a direct measure of its level of satisfaction, the prices may be a more important measure from the system's point of view, since they account for the level of congestion as seen by users and are the direct indication of how each user could accommodate fluctuations in the system's state. The designer can consider various ways of combining either the prices or the costs of the users. We shall concentrate on *user* optimization, i.e. trying to reduce the price or cost of each and every user:

Definition 2.1. Consider two capacity configurations $\boldsymbol{c}$ and $\hat{\boldsymbol{c}}$ and let J^i and $\hat{J}^i$ (and correspondingly, λ_s^i and $\hat{\lambda}_s^i$) be the cost (or price) of user i at the respective equilibrium. Configuration $\hat{\boldsymbol{c}}$ is said to be *user cost (price) efficient* relative to configuration $\boldsymbol{c}$, if $\hat{J}^i \le J^i$ ($\hat{\lambda}_s^i \le \lambda_s^i$), for all $i \in \mathcal{I}$.

As mentioned above, we will focus on cases for which the uniqueness of the Nash equilibrium point has been established. The first is that of 'identical' users, defined in the following.

Definition 2.2. Users are said to be *identical* if all their demands are equal, that is, $r^i = r^j$ for all $i, j \in \mathcal{I}$.

The following result has been established [18]:

Lemma 2.3. *The routeing game in a general topology network with identical users has a unique Nash equilibrium. This equilibrium is symmetrical, i.e.* $f_l^i = f_l^j = f_l/I$ *for all* $i, j \in \mathcal{I}$ *and* $l \in \mathcal{L}$.

Another case of special interest is that of 'simple users', defined as follows.

Definition 2.4. A user is said to be *simple* if all of its flows are routed through paths of minimal delay.

Users often route their flows according to the 'simple' scheme due to practical considerations. Many typical routeing algorithms send flows through the shortest paths, without accounting for derivatives or bifurcating flows. The corresponding necessary and sufficient conditions require the existence of some $(\lambda_u)_{u\in\mathcal{V}}$, such that

$$\lambda_u = T_{uv} + \lambda_v, \quad \text{if } f^i_{uv} > 0, \quad (u, v) \in \mathcal{L}, \tag{2.7}$$

$$\lambda_u \leq T_{uv} + \lambda_v, \quad \text{if } f^i_{uv} = 0, \quad (u, v) \in \mathcal{L}, \tag{2.8}$$

$$\lambda_d = 0. \tag{2.9}$$

We note that when (2.7)–(2.9) are satisfied, λ_s is the sum of delays on the links of any minimum-delay path. Since simple users send their flow exclusively over such paths, the equilibrium cost of user i is given by

$$J^i = r^i \lambda_s. \tag{2.10}$$

We shall refer to the value of λ_s as the price of the simple users. From (2.7)–(2.9), it is easy to see that users routeing according to the optimality conditions (2.4)–(2.6) become simple users as their population grows to infinity and their individual demands become infinitesimally small, while their total demand remains R. This is the typical scenario in a transportation network. The following result follows from [18].

Lemma 2.5. *In a general topology with simple users there exists a Nash equilibrium, which is unique with respect to the total link flows.*

3. The Braess paradox

The Braess paradox, originally introduced for traffic flows, describes the existence of non-intuitive equilibrium points in networks of various kinds ranging from transportation networks [8], to electrical circuits, to mechanical networks of springs and strings, to hydraulic systems (see [6] and references therein), to queueing networks [5, 7], to loss networks [2], and to distributed computational systems [10].

In this section we present an example that adapts the Braess paradox to the routeing game. The example demonstrates that an addition of capacity may, in general, increase both the price and the cost of each and every user. We also explain that, despite its non-intuitive behavior, there is nothing paradoxical about the 'paradox,' which is, in fact, an instance of the Pareto inefficiency of non-cooperative equilibria.

Consider the network depicted in Figure 1. Links $(1, 2)$ and $(3, 4)$ each have capacity c_1. Link $(1, 3)$ represents a path of n tandem links, each with capacity c_2. Similarly, links $(2, 4)$ and $(2, 3)$ are paths of n consecutive links each with capacities c_2 and c_3 respectively. There are I identical users, each with an average throughput demand r, sending flow from node 1 to node 4. For $c_2 \gg Ir$, each of the paths $(1, 3)$ and $(2, 4)$ approximates a link with non-negligible delay that has low sensitivity to flow changes. Such constructions are required in order to reproduce the classical Braess paradox in a queueing setting [5].

As explained in Section 2, this system has a unique and symmetrical Nash equilibrium, that is, the flows (and thus, the costs and prices) of the users at equilibrium are equal. Figures 2 and 3 show, correspondingly, the user price and cost as functions of c_3, for $c_1 = 2.7$, $c_2 = 27$,

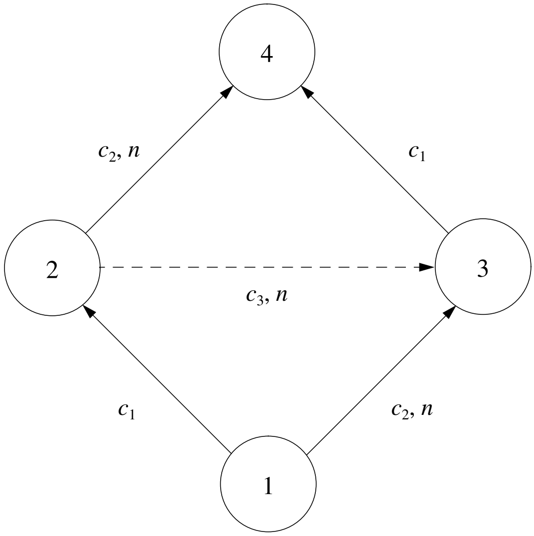

FIGURE 1: Network paradox.

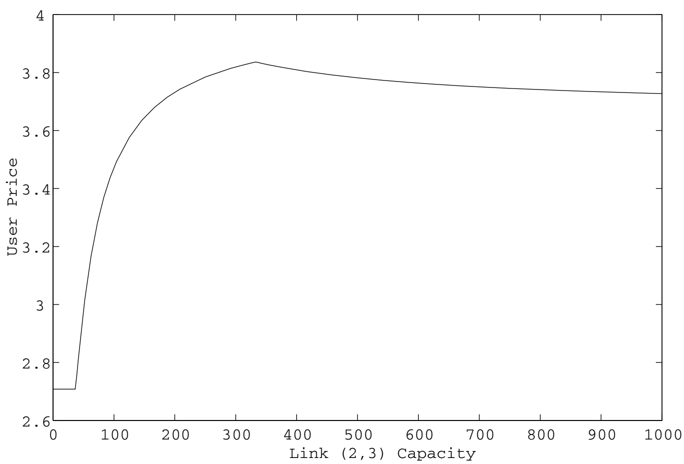

FIGURE 2: User price as a function of link capacity c_3.

$n = 54$, $I = 10$ and $r = 0.2$. The figures indicate that, for any $c_3 > 0$, both the price and the cost of each user are higher than for $c_3 = 0$, i.e., eliminating the path (2, 3) leads to an improvement in performance for all users. More surprisingly, it can be verified that this behavior persists even if $c_3 = \infty$, that is, if nodes 2 and 3 are merged into a single node.

Like most paradoxes, the Braess paradox is a paradox only in name. Adding capacity to link (2, 3)—or adding the link itself—augments the strategy space F^i of each user $i \in \mathcal{I}$ and thus the product strategy space F. Let $\hat{F}^i$ and $\hat{F}$, respectively, denote the resulting strategy spaces. If there were a single user, the routeing game would be an optimization problem, and thus performance could not deteriorate (since we would be optimizing over the augmented

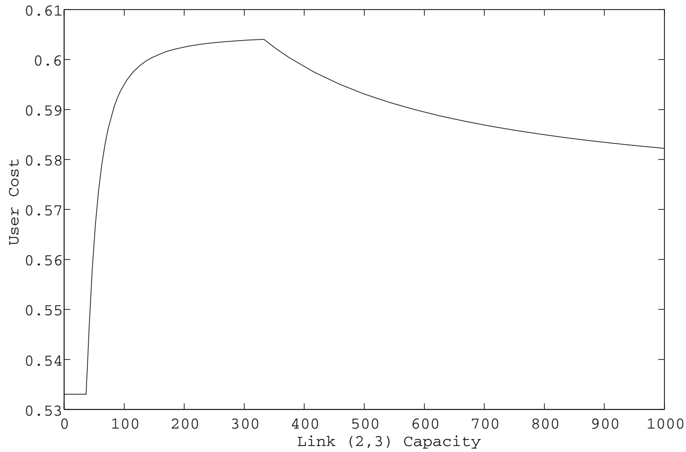

FIGURE 3: User cost as a function of link capacity c_3.

strategy space $\hat{F}$). In the presence of multiple users, however, the network is not operating at optimum but at equilibrium and there are no guarantees that augmenting the strategy space will lead to improved performance. In the prisoner's dilemma [16], for example, eliminating link (2, 3) corresponds to not giving the suspects the option to confess, thus resulting in the universally optimal outcome.

The Braess paradox is, in fact, a projection of the prisoner's dilemma onto the networking domain. In other words, if the equilibrium of the routeing game were Pareto efficient, then the paradox would not occur. To see this, let $\boldsymbol{f} \in F$ be the equilibrium before the addition of link (2,3) and $\hat{\boldsymbol{f}} \in \hat{F}$ be the equilibrium after the link is added. If $\hat{\boldsymbol{f}}$ were Pareto efficient in $\hat{F}$, then there would exist no point $\tilde{\boldsymbol{f}} \in \hat{F}$ such that $J^i(\tilde{\boldsymbol{f}}) \leq J^i(\hat{\boldsymbol{f}})$ for all $i \in \mathcal{I}$ and $J^j(\tilde{\boldsymbol{f}}) < J^j(\hat{\boldsymbol{f}})$ for some user j, thus the paradox would not occur. In [11], it is shown that Braess' and a number of other 'paradoxes' are, in fact, structurally equivalent to the prisoner's dilemma.

4. Avoiding the paradox

The example presented in Section 3 demonstrates that adding capacity to a network, even in infinite amounts, may result in an increase of both the price and the cost of each and every user. This indicates that an upgrade of a general network, in terms of capacity and link addition, should be carried out in a cautious way. In this section we devise methods for upgrading a general network, so that the Braess paradox does not occur.

Consider an upgrade that is achieved by multiplying the capacity of each link by some constant factor $\alpha > 1$. That is, from a capacity configuration $\boldsymbol{c} = (c_l)_{l \in \mathcal{L}}$ we obtain an augmented capacity configuration $\hat{\boldsymbol{c}} = (\hat{c}_l)_{l \in \mathcal{L}}$, such that $\hat{c}_l = \alpha c_l$ for all $l \in \mathcal{L}$. We say that $\hat{\boldsymbol{c}}$ is an *α-product* of $\boldsymbol{c}$. We then have the following.

Proposition 4.1. *If $\hat{\mathbf{c}}$ is an α-product of $\mathbf{c}$ then:*

1. *for identical users, $\hat{\boldsymbol{c}}$ is user price efficient relative to $\boldsymbol{c}$,*
2. *for simple users, $\hat{\boldsymbol{c}}$ is user price and cost efficient relative to $\boldsymbol{c}$.*

Proof. Consider first the case of identical users. Let $\boldsymbol{f}$ and $\hat{\boldsymbol{f}}$ be the Nash equilibria under capacity configurations $\boldsymbol{c}$ and $\hat{\boldsymbol{c}}$, respectively. Consider now the same network, but with the initial capacity configuration $\boldsymbol{c}$ and a set of I users with demands $\tilde{r}^i = r^i/\alpha$, $i \in \mathcal{I}$. It is easy to verify that, for such a capacity configuration and users, the optimality conditions (2.4)–(2.6) are satisfied by the system flow configuration $\tilde{\boldsymbol{f}}$, where $\tilde{f}_l^i = \hat{f}_l^i/\alpha$, for all users i and links l, and the corresponding Lagrange multipliers are $\tilde{\lambda}_u^i = \alpha\hat{\lambda}_u^i$, for all users i and nodes u. Therefore, $\tilde{\boldsymbol{f}}$ is the unique Nash equilibrium of the new network.

We construct a directed network $(\mathcal{V}', \mathcal{L}')$, whose set of nodes is identical to that of $(\mathcal{V}, \mathcal{L})$ (i.e. $\mathcal{V}' = \mathcal{V}$) and the set of links $\mathcal{L}'$ is constructed as follows:

- for each link $l = (u, v) \in \mathcal{L}$, such that $f_l^i \geq \tilde{f}_l^i$, we have a link $l' = (u, v) \in \mathcal{L}'$; to such a link l' we assign a (flow) value $x_{l'} = f_l^i - \tilde{f}_l^i = (f_l - \tilde{f}_l)/I$.
- for each link $l = (u, v) \in \mathcal{L}$, such that $f_l^i < \tilde{f}_l^i$, we have a link $l' = (v, u) \in \mathcal{L}'$; to such a link l' we assign a (flow) value $x_{l'} = \tilde{f}_l^i - f_l^i = (\tilde{f}_l - f_l)/I$.

In other words, we redirect links according to the relation between f_l^i and $\tilde{f}_l^i$. It is easy to verify that the values $x_{l'}$ constitute a nonnegative, directed flow in the network. Since $r^i > \tilde{r}^i$, $x_{l'}$ must carry some flow (an amount $r^i - \tilde{r}^i$) from the source s to the destination d. Thus, there is some (simple) path p in $(\mathcal{V}', \mathcal{L}')$, such that $x_{l'} > 0$ for all $l' \in p$.

Consider now a link $l' = (u, v) \in p$. Since $x_{l'} > 0$, either $f_{uv}^i > \tilde{f}_{uv}^i$ or else $\tilde{f}_{vu}^i > f_{vu}^i$. In the case where $f_{uv}^i > \tilde{f}_{uv}^i \geq 0$, we have

$$\lambda_u^i - \lambda_v^i = f_{uv}^i T'_{uv} + T_{uv} > \tilde{f}_{uv}^i \tilde{T}'_{uv} + \tilde{T}_{uv} \geq \tilde{\lambda}_u^i - \tilde{\lambda}_v^i, \tag{4.1}$$

where the first transition follows from the optimality conditions, since $f_{uv}^i > 0$; the second is due to $f_{uv}^i > \tilde{f}_{uv}^i$, which implies $f_{uv} > \tilde{f}_{uv}$ (since the users are identical) and, thus, $T_{uv} > \tilde{T}_{uv}$ and $T'_{uv} > \tilde{T}'_{uv}$ (since $c_{uv} = \tilde{c}_{uv}$); the third transition is again due to the optimality conditions. In the case where $\tilde{f}_{vu}^i > f_{vu}^i \geq 0$, we have, by symmetry, that

$$\tilde{\lambda}_v^i - \tilde{\lambda}_u^i > \lambda_v^i - \lambda_u^i. \tag{4.2}$$

Note that the results of equations (4.1) and (4.2) are, in fact, identical.

Denote $\Delta\lambda_u = \lambda_u^i - \tilde{\lambda}_u^i$, for all $u \in \mathcal{V}$. From (4.1) and (4.2) we conclude that, for $l' = (u, v)$, $x_{l'} > 0$ implies that $\Delta\lambda_u > \Delta\lambda_v$. This means that along the path p described above we have a monotonically decreasing sequence of $\Delta\lambda$'s (starting from the source s). Since $\lambda_d = \tilde{\lambda}_d = 0$, we conclude that $\tilde{\lambda}_s^i < \lambda_s^i$. Hence: $\hat{\lambda}_s^i = \tilde{\lambda}_s^i/\alpha < \tilde{\lambda}_s^i < \lambda_s^i$, thus proving the first part of the lemma.

Consider now the case of simple users. Let $(f_l)_{l\in\mathcal{L}}$ and $(\hat{f}_l)_{l\in\mathcal{L}}$ be the total flow vectors at the Nash equilibria corresponding to $\boldsymbol{c}$ and $\hat{\boldsymbol{c}}$. Consider the same network, but with the initial capacity configuration $\boldsymbol{c}$ and a set of simple users with demands $\tilde{r}^i = r^i/\alpha$, for all users i. For this network, the optimality conditions (2.7)–(2.9) are satisfied by $(\tilde{f}_l)_{l\in\mathcal{L}}$, with $\tilde{f}_l = \hat{f}_l/\alpha$, $l \in \mathcal{L}$, and Lagrange multipliers $\tilde{\lambda}_u = \alpha\hat{\lambda}_u$, $u \in \mathcal{V}$. Proceeding as in the first part of the proof, we obtain $\tilde{\lambda}_s < \lambda_s$, thus: $\hat{\lambda}_s = \tilde{\lambda}_s/\alpha < \tilde{\lambda}_s < \lambda_s$, and the claim on the prices follows. The claim on the costs is immediate from equation (2.10).

The following proposition gives a sufficient condition on the product parameter α for obtaining cost efficiency in the case of identical users.

Proposition 4.2. *In a general topology network with identical users, a capacity configuration $\hat{\boldsymbol{c}}$ that is an α-product of a capacity configuration $\boldsymbol{c}$ is user cost efficient relative to $\boldsymbol{c}$, for $\alpha > I$.*

Proof. It can be shown [18] that for a given capacity configuration $\boldsymbol{c}$, the vector of total link flows at the Nash equilibrium corresponds to the unique minimum of the function

$$H((f_l)_{l\in\mathcal{L}}) = \sum_{l\in\mathcal{L}} \frac{f_l}{c_l - f_l} - (I-1)\sum_{l\in\mathcal{L}} \ln(c_l - f_l).$$

Denote by $(f_l)_{l\in\mathcal{L}}$ and $(\hat{f}_l)_{l\in\mathcal{L}}$ the vectors of the total link flows at the Nash equilibria corresponding to $\boldsymbol{c}$ and $\hat{\boldsymbol{c}}$. Since $(\hat{f}_l)_{l\in\mathcal{L}}$ minimizes the H-function that corresponds to configuration $\hat{\boldsymbol{c}} = \alpha\boldsymbol{c}$, we have

$$\sum_{l\in\mathcal{L}} \frac{\hat{f}_l}{\alpha c_l - \hat{f}_l} - (I-1)\sum_{l\in\mathcal{L}} \ln(\alpha c_l - \hat{f}_l) \le \sum_{l\in\mathcal{L}} \frac{f_l}{\alpha c_l - f_l} - (I-1)\sum_{l\in\mathcal{L}} \ln(\alpha c_l - f_l),$$

and thus

$$\begin{aligned}\sum_{l\in\mathcal{L}} \frac{\hat{f}_l}{\alpha c_l - \hat{f}_l} &\le \sum_{l\in\mathcal{L}} \frac{f_l}{\alpha c_l - f_l} - (I-1)\sum_{l\in\mathcal{L}} \ln\left(\frac{\alpha c_l - f_l}{\alpha c_l - \hat{f}_l}\right)\\ &\le \sum_{l\in\mathcal{L}} \frac{f_l}{\alpha c_l - f_l} - (I-1)\sum_{l\in\mathcal{L}} \ln\left(1 - \frac{f_l}{\alpha c_l}\right). \end{aligned} \tag{4.3}$$

Hence, in order to prove that $\sum_{l\in\mathcal{L}} \hat{f}_l/(\alpha c_l - \hat{f}_l) < \sum_{l\in\mathcal{L}} f_l/(c_l - f_l)$, it is enough to show that

$$Q(f_l) = \frac{f_l}{c_l - f_l} - \frac{f_l}{\alpha c_l - f_l} + (I-1)\ln\left(1 - \frac{f_l}{\alpha c_l}\right) > 0,$$

or, since $Q(0) = 0$:

$$\frac{\mathrm{d}Q}{\mathrm{d}f_l} = \frac{c_l}{(c_l - f_l)^2} - \frac{\alpha c_l}{(\alpha c_l - f_l)^2} - \frac{I-1}{\alpha c_l - f_l} > 0.$$

In view of $\alpha > I$ and after some algebraic manipulation one can see that it suffices to show

$$(2\alpha + 1)(I-1)c_l^2 + (I-1)f_l^2 > f_l c_l(I\alpha + 2I - 3),$$

or $(\alpha I - 2\alpha + I - 1)c_l > f_l(2I - 3)$, which trivially holds. Since $f_l^i = f_l/I$ (and $f_l^i h = f_l h/I$),

$$\hat{J}^i = \frac{1}{I}\sum_{l\in\mathcal{L}} \frac{f_l h}{\alpha c_l - f_l h} < \frac{1}{I}\sum_{l\in\mathcal{L}} \frac{f_l}{c_l - f_l} = J^i,$$

and this concludes the proof.

Propositions 4.1 and 4.2 indicate that capacity should be added across the network, rather than on a local scale. The next result suggests that yet another good design practice is to focus the upgrades on direct connections between sources and destinations.

Consider an upgrade that is achieved by adding a link between the source s and the destination d. Denote by $\boldsymbol{c}$ and $\hat{\boldsymbol{c}}$, respectively, the capacity configurations before and after this addition. We say that $\hat{\boldsymbol{c}}$ is a *direct augmentation* of $\boldsymbol{c}$. We then have the following.

Proposition 4.3. *In a general topology network, consider two capacity configurations $\hat{\boldsymbol{c}}$ and $\boldsymbol{c}$, such that $\hat{\boldsymbol{c}}$ is a direct augmentation of $\boldsymbol{c}$.*

1. *If the users are identical, then $\hat{\boldsymbol{c}}$ is user price efficient relative to $\boldsymbol{c}$.*

2. *If the users are simple, then $\hat{\boldsymbol{c}}$ is user price and cost efficient relative to $\boldsymbol{c}$.*

Proof. Consider first the case of identical users. Let $\boldsymbol{f}$ and $\hat{\boldsymbol{f}}$ be the Nash equilibria under configurations $\boldsymbol{c}$ and $\hat{\boldsymbol{c}}$, respectively. Denote by $\hat{l}$ the link added between s and d. Then $\hat{\mathcal{L}} = \mathcal{L} \cup \{\hat{l}\}$. If $\hat{f}_{\hat{l}} = 0$, then $\hat{f}_l = f_l$ for all other links $l \in \hat{\mathcal{L}}$, and the prices of each user are equal under both configurations.

Assume, then, that $\hat{f}_{\hat{l}} > 0$. Consider now the same network, but with the initial capacity configuration $\boldsymbol{c}$ and a set of I identical users, with demands $\tilde{r}^i = r^i - \hat{f}_{\hat{l}}/I$, for all users i. For this network it is easy to verify that the optimality conditions (2.4)–(2.6) are satisfied by the system flow configuration $\tilde{\boldsymbol{f}}$, with $\tilde{f}_l^i = \hat{f}_l^i$, for all users i and links $l \in \mathcal{L}$, and the Lagrange multipliers $\tilde{\lambda}_u^i = \hat{\lambda}_u^i$, $i \in \mathcal{I}$ and $u \in \mathcal{V}$. Thus, $\tilde{\boldsymbol{f}}$ is the unique Nash equilibrium corresponding to such a network. Since this network has the initial capacity configuration, but the demand of each user is now less ($\tilde{r}^i < r^i$), following the proof of Proposition 4.1, one can show that:

$$\hat{\lambda}_s^i = \tilde{\lambda}_s^i < \lambda_s^i, \quad i \in \mathcal{I},$$

and this completes the proof for the case of identical users. The proof for simple users follows similar lines and the details are omitted.

Remark. The last proposition holds also for the case of adding capacity to an already existing direct link $\hat{l} = (s, d)$ between source and destination. The proof is omitted.

5. Conclusions

We have investigated the problem of efficiently adding resources to general topology networks where users non-cooperatively implement their optimal routeing strategies. We have derived a set of methodologies for upgrading the network so that the Braess paradox does not occur. One indication from our results is that capacity should be added across the network, rather than on a local (e.g. single link) scale. This fits well with common engineering practice, where common folklore suggests that local improvement may only result in transferring the problem somewhere else in the system. Another indication is that upgrades should be aimed at direct connections between the source and the destination. This provides additional evidence of the potential benefit of modern networking practices, which tend to decouple complex structures by pre-allocating resources to various non-interfering routeing paths.

The practical application of our findings should be facilitated by exploiting additional structure of real-world scenaria. This is the case, for example, with the global Internet, which has a characteristic hierarchical topology. Indeed, due to this structure it should be sufficient to uniformly upgrade just the ‘backbone’ links, and still avoid the Paradox; conversely, it should be sufficient to uniformly upgrade the links of a ‘stub’ network, without requiring any changes in other Internet links. Similarly, a ‘direct’ connection should be interpreted in the hierarchical sense: it may refer, for example, to a backbone link between the regional networks with which the source and destination routers are associated, rather than to a link between the routers themselves.

Appendix

A. Derivation of the Kuhn-Tucker conditions

Proposition A.1. *Assume that there is a routeing strategy profile $\boldsymbol{f} \in F$ that is stable, in the sense that $f_l < c_l$ holds at all links $l \in \mathcal{L}$. If the cost function of a user $i \in \mathcal{I}$ is defined by (2.1) and (2.2), then $\boldsymbol{f}^i$ is the optimal response of user i to $\boldsymbol{f}^{-i}$ if and only if there exist (Lagrange multipliers) $(\lambda_u^i)_{u \in \mathcal{V}}$, such that (2.4)–(2.6) are satisfied.*

Proof. F^i is convex, and, as observed, given a strategy profile $\boldsymbol{f}^{-i}$ of the other users, the cost of user i, as defined by (2.1) and (2.2), is a convex function of its strategy $\boldsymbol{f}^i$. Consequently, any local minimum of J^i is also a global minimum. Also, since J^i increases to infinity as f_l^i increases to c_l, the stability constraint $f_l^i \leq c_l$ can be treated as absent.

The stability assumption, namely that there is a routeing strategy profile $\boldsymbol{f} \in F$ that is stable, implies that Slater's 'interiority' condition [3] is satisfied. The Kuhn-Tucker saddle point theorem then guarantees the existence of the respective (finite) Lagrange multipliers. Therefore, if we form the Lagrangian

$$G^i_{\boldsymbol{f}^{-i}}(\boldsymbol{f}^i, \lambda^{\mathbf{i}}, \eta^{\mathbf{i}}) = \sum_{u \in \mathcal{V}} \sum_{v \in \mathrm{Out}(u)} f^i_{uv} T_{uv} + \sum_{u \in \mathcal{V}} \lambda^i_u \left(\sum_{v \in \mathrm{In}(u)} f^i_{vu} - \sum_{v \in \mathrm{Out}(u)} f^i_{uv} + r^i_u \right) + \sum_{u \in \mathcal{V}} \sum_{v \in \mathrm{Out}(u)} \eta^i_{uv} f^i_{uv}$$

where, for all $u \in \mathcal{V}$ and $v \in \mathrm{Out}(u)$,

$$\eta^i_{uv} \leq 0 \tag{A.1}$$

and

$$\eta^i_{uv} f^i_{uv} = 0, \tag{A.2}$$

a flow $\boldsymbol{f}^i$ minimizes $J^i(\cdot, \boldsymbol{f}^{-i})$ if and only if it is stationary with respect to $G^i_{\boldsymbol{f}^{-i}}$ [15]. Taking derivatives and equating to zero one obtains

$$0 = \frac{\partial G^i_{\boldsymbol{f}^{-i}}}{\partial f^i_{uv}} = f^i_{uv} T'_{uv} + T_{uv} + \lambda^i_v - \lambda^i_u + \eta^i_{uv}. \tag{A.3}$$

Finally, (A.1)–(A.2) imply that (A.3) is equivalent to (2.4)–(2.6), hence establishing the result.

Acknowledgements

This work was carried out in part while Yannis A. Korilis was with the Center for Telecommunications Research at Columbia University.

Ariel Orda was supported by the Lowengart Research Fund and by the fund for the promotion of research at the Technion. This work was carried out in part while Ariel Orda was visiting with the Center for Telecommunications Research at Columbia University.

References

[1] Altman, E. (1994). Flow control using the theory of Zero-Sum Markov games. *IEEE Trans. Automatic Control* **39,** 814–818.

[2] Bean, N. G., Kelly, F. P. and Taylor, P. G. (1997). Braess' paradox in a loss network. *J. Appl. Prob.* **34,** 155–159.

[3] Ben-Israel, A., Ben-Tal, A. and Zlobec, S. (1981). *Optimality in Nonlinear Programming: a Feasible Directions Approach*. Wiley, New York.

[4] Bertsekas, D. and Gallager, R. (1992). *Data Networks*, 2nd edn. Prentice-Hall, Englewood Cliffs, NJ.

[5] Cohen, J. E. and Kelly, F. P. (1990). A paradox of congestion in a queuing network. *J. Appl. Prob.* **27,** 730–734.

[6] Cohen, J. E. and Horwitz, P. (1991). Paradoxical behaviour of mechanical and electrical networks. *Nature* **352,** 699–701.

[7] Cohen, J. E. and Jeffries, C. (1997). Congestion resulting from increased capacity in single-server queueing networks. *IEEE/ACM Trans. Networking* **5,** 305–310.

[8] Dafermos, S. and Nagurney, A. (1984). On some traffic equilibrium theory paradoxes. *Transportation Research B* **18,** 101–110.

[9] Dubey, P. (1986). Inefficiency of Nash equilibria. *Math. Operat. Res.* **11,** 1–8.

[10] Glance, N. S. and Hogg, T. (1995). Dilemmas in computational societies. In *Proc. 1st International Conference on Multiagent Systems*, San Francisco, CA, June 1995, pp. 117–124.

[11] Irvine, A. D. (1993). How Braess' paradox solves Newcomb's problem. *Internat. Stud. Philos. Sci.* **7,** 141–160.

[12] Korilis, Y. A., Lazar, A. A. and Orda, A. (1995). Architecting non-cooperative networks. *IEEE J. Sel. Areas Commun.* **13,** 1241–1251.

[13] Korilis, Y. A., Lazar, A. A. and Orda, A. (1997). Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Trans. Networking* **5,** 161–173.

[14] Korilis, Y. A., Lazar, A. A. and Orda, A. (1997). Capacity allocation under non-cooperative routing. *IEEE Trans. Automatic Control* **42,** 309–325.

[15] Luenberger, D. G. (1984). *Linear and Nonlinear Programming*, 2nd edn. Addison Wesley, Reading, MA.

[16] Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, MA.

[17] Nicholson, W. (1994). *Intermediate Microeconomics and its Applications*, 6th edn. Dryden Press, Fort Worth, TX.

[18] Orda, A., Rom, R. and Shimkin, N. (1993). Competitive routing in multiuser communication networks. *IEEE/ACM Trans. Networking* **1,** 510–521.

[19] Zhang, Z. and Douligeris, C. (1992). Convergence of synchronous and asynchronous greedy algorithms in a multiclass telecommunications environment. *IEEE Trans. Commun.* **40,** 1277–1281.