

On the Existence of Equilibria in Noncooperative Optimal Flow Control

YANNIS A. KORILIS AND AUREL A. LAZAR

Columbia University, New York, New York

Abstract. The existence of Nash equilibria in noncooperative flow control in a general product-form network shared by K users is investigated. The performance objective of each user is to maximize its average throughput subject to an upper bound on its average time-delay. Previous attempts to study existence of equilibria for this flow control model were not successful, partly because the time-delay constraints couple the strategy spaces of the individual users in a way that does not allow the application of standard equilibrium existence theorems from the game theory literature. To overcome this difficulty, a more general approach to study the existence of Nash equilibria for decentralized control schemes is introduced. This approach is based on directly proving the existence of a fixed point of the best reply correspondence of the underlying game. For the investigated flow control model, the best reply correspondence is shown to be a function, implicitly defined by means of K interdependent linear programs. Employing an appropriate definition for continuity of the set of optimal solutions of parametrized linear programs, it is shown that, under appropriate conditions, the best reply function is continuous. Brouwer's theorem implies, then, that the best reply function has a fixed point.

Categories and Subject Descriptors: C.2.3 [**Computer-Communication Networks**]: Network Operations—*network management*; D.4.8 [**Operating Systems**]: Performance—*queueing theory*; G.1.6 [**Mathematics of Computing**]: Optimization—*constrained optimization, linear programming*

General Terms: Performance, Theory

Additional Key Words and Phrases: Fixed points, flow control, game theory, Nash equilibria

1. Introduction

Flow control algorithms allow each user of a telecommunication network to regulate the traffic it sends into the network and thereby satisfy certain performance objectives. Typically, as the load of traffic a user offers to the network increases, the packet queues at the nodal buffers build up and may eventually lead to degradation in the quality of service provided to the user, for example, to excessive delays and even to packet loss due to buffer overflow. The main objective of flow control is to achieve an efficient trade-off between high network utilization (e.g., average throughput) and guaranteed quality of service (e.g., low average time-delay and small probability of packet loss).

The research reported here by the Office of Naval Research under Contract N00014-90-J-1289. Authors' address: Department of Electrical Engineering and Center for Telecommunications Research, Columbia University, New York, NY 10027-6699, e-mail: john@ctr.columbia.edu, aurel@ctr.columbia.edu

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 1995 ACM 0004-5411/95/0500-0584 \$03.50

The complexity of modern high-speed wide-area networks calls for decentralized congestion control algorithms, where the control tasks are performed “locally” by the appropriate network entities. A decentralized flow control scheme, under which each user determines the load of traffic it offers to the network so that it satisfies its own performance objective, can be modeled as a *game* [Fudenberg and Tirole 1992; Myerson 1991]. In this setting, flow control can either be cooperative, if the users are allowed to make joint decisions on their strategies, or noncooperative, if joint decisions are not possible. An extensive survey on cooperative and noncooperative optimal flow control can be found in Korilis and Lazar [1992]. In the present paper, the focus will be on noncooperative flow control.

In a noncooperative flow control scheme, each user selects among its available strategies those that optimize its own performance objective. Since the performance of a user depends, in general, on the strategies of the other users, this mode of operation results in a dynamic behavior. Understanding the dynamics of the network when the users implement their optimal control strategies is of fundamental importance. One of the key questions is whether the network converges to an equilibrium operating point, such that no user is willing to *unilaterally* modify its flow control strategy. In the game-theory language such a point is called a *Nash equilibrium*.

The game theoretic models that arise in congestion control schemes are *constrained* games, due to the presence of stability and/or quality of service constraints. In this type of games the set of strategies that a user is allowed to choose from is affected by the actions of the other users. The subset of the product strategy space defined by the constraints is called the *joint* strategy space. Existence of Nash equilibria of constrained games is typically established in the congestion control literature based on a theorem by Rosen [1965]. The theorem asserts that a game has an equilibrium point if:

- (i) The joint strategy space is convex and compact.
- (ii) The objective function that each player seeks to maximize is concave in its own strategy and continuous at every point in the product strategy space.

The “greedy algorithm” for an $M/1$ queue was one of the first noncooperative flow control schemes to be studied in a game-theoretic framework [Bovopoulos and Lazar 1988; Douligeris and Mazumdar 1988, 1992]. Each user maximizes its power (average throughput over average time-delay) by regulating its (state independent) input rate. The properties of the power criterion indicate that the stability constraint of the system does not become effective, provided that the users behave optimally. Consequently, this constraint may be treated as absent, resulting in an unconstrained game, that was shown to have a unique Nash equilibrium. Convergence of synchronous and asynchronous implementations of the algorithm to the unique equilibrium point was examined in Zhang and Douligeris [1992]. The greedy algorithm for a “store-and-forward” network of $M/1$ queues was studied in Bovopoulos and Lazar [1988]. Existence was established using Rosen’s theorem; again, the underlying model was an unconstrained one. The same theorem was also used in Orda et al. [1993] to show existence of Nash equilibria in a noncooperative routing scheme.

State-dependent flow control strategies for a network of $M/1$ queues were analyzed in Hsiao and Lazar [1991]. The traffic generated by each user enters

the network at a rate that depends on the number of packets that the user has outstanding in the network. The performance objective of the user is to maximize its average throughput subject to an upper bound on its average time-delay. The structure of the optimal strategies was found by means of a simplified equivalent network model, that allows a transformation of the optimization problem associated with the user to a linear program. Existence of Nash equilibria was shown for a special class of networks that satisfy certain monotonicity properties.

In the present paper, we study the existence of Nash equilibria of the flow control model introduced in Hsiao and Lazar [1991] for a general product-form network where the monotonicity assumptions of that reference do not necessarily hold. Simple examples show that the joint strategy space—defined by the time-delay constraints of all users—does not satisfy the convexity condition posed by Rosen’s theorem. On the other hand, the conventional game theoretic approach of incorporating the time-delay constraint of each user as a penalty component in its utility results in utility functions that either are not continuous, or do not satisfy certain second order monotonicity properties required by the standard equilibrium existence theorems (e.g., quasi-concavity).

To overcome these difficulties, we take a more general approach to the problem of equilibrium existence in noncooperative flow control algorithms. The concept of the *best reply correspondence* of the underlying game plays a central role in this approach. Finding a noncooperative equilibrium of a decentralized control scheme (“game”) amounts to simultaneously solving a number of interdependent individual constraint optimization problems. Given a choice of strategies for all but one of the controllers, the set of optimal strategies for the remaining controller can be found. This way, a point-to-set correspondence, that maps the product strategy space of the game into itself, is defined. This is the best reply correspondence of the game [Border 1985]. Any Nash equilibrium is a *fixed point* of this correspondence [Nash 1951].

Instead of using standard game-theoretic results that concentrate on the properties of the utility functions and the joint strategy space, we focus on the best reply correspondence of the flow control game and examine the existence of fixed points of this correspondence. From a technical standpoint, the merits of this approach are clear. The requirements in Rosen’s theorem—and other similar results from game theory—guarantee that the best reply correspondence satisfies the sufficient conditions of some fixed-point theorem. Nevertheless, the best reply correspondence might satisfy these conditions, while the requirements on the utility functions and the joint strategy space are not met. Conceptually, on the other hand, the best reply correspondence is the natural representation of a decentralized control algorithm, for which objective functions or quality of service constraints might not even be explicitly defined.

The outline of the paper is the following: The flow control model and its best reply correspondence are presented in Section 2. A fundamental result from Hsiao and Lazar [1991] on the structure of the best reply correspondence is briefly analyzed in Section 3. This result is based on the observation that each user can determine its best reply to any choice of strategies of the other users by means of a linear program. In Section 4, we provide a further characterization of the best reply correspondence by showing that it is a function. The continuity properties of the best reply function are examined in Section 5. In studying these properties, we adopt a general definition and derive sufficient




FIG. 1. Flow control model of a general product-form network.

conditions for continuity of parametrized linear programs. The related analysis is presented in Appendix A and can be read independently of the rest of the paper. Using these sufficient conditions, we establish, under a certain regularity assumption, continuity of the best reply function and, thus, existence of Nash equilibria.

2. The Model and Game Theoretic Formulation

The flow control model under consideration is depicted in Figure 1. There are K users sharing a network of I quasi-reversible queues [Kelly 1979]. This is a generalization of the model studied in Hsiao and Lazar [1991], where a network of $M/1$ -FCFS queues was considered. Node i can be of any of the following types:

- (i) A single server FCFS (first-come-first-served) queue with class independent service times that are exponentially distributed with mean $1/\mu^i$, that is, an $M/1$ -FCFS queue.
- (ii) A single server queue where the service times for class k packets are independent and identically distributed (iid) following a general distribution with mean $1/\mu^{ki}$ and the service discipline is either LCFS-PR (last-come-first-served, preemption/resume) or PS (processor sharing), that is, either a $GI/1$ -LCFS or a $GI/1$ -PS queue.
- (iii) An infinite server queue where the service times for class k packets are independent and identically distributed following a general distribution with mean $1/\mu^{ki}$, that is, a GI/∞ queue.

This is one of the most general analytically tractable network models [Walrand 1988] and it is well-suited for modeling a broad class of practical networking environments. It allows for nodal schedulers that employ not only the traditional FCFS policy, but also a wide class of scheduling algorithms that approximate the processor sharing service discipline and have been receiving increasing attention in recent years due to their fairness properties and their ability to provide quality of service guarantees [Demers et al. 1989; Parekh and Gallager 1993]. Moreover, the model allows for non-negligible propagation delays that arise in high-speed, wide-area networks, and can be represented as GI/∞ queues.

Let \mathcal{S}_1 denote the set of nodes that belong to types (i) and (ii) and \mathcal{S}_2 the set of type (iii) nodes. We denote the average service rate of class k packets in queue i by μ^{ki} , with the understanding that it is class independent whenever i is an $M/1$ -FCFS queue.

The traffic generated by user k enters the network as a conditional Poisson stream, with rate λ_k^k when the number of packets that the user has outstanding in the network is l_k . In practice, the information l_k is provided to the user by means of an acknowledgment protocol. The maximum number of packets that user k can have outstanding is N_k and the maximum rate at which it can send packets into the network is c^k , a positive constant.

Routing is probabilistic: Upon service completion, class k packets are routed from node i to node j with probability r^{kij} , $1 \leq i, j \leq I$. Class k packets enter the network at node j with probability r^{k0j} and exit through node i with probability $r^{ki0} = 1 - \sum_{j=1}^I r^{kij}$. Let θ^{ki} denote the visit ratio of class k packets to node i . The visit ratios can be determined by the following linear system:

$$\theta^{ki} = r^{k0i} + \sum_{j=1}^I \theta^{kj} r^{kji}, \quad 1 \leq i \leq I, \quad 1 \leq k \leq K.$$

Let the state of queue i be described by $n_i = (n_{1i}, \dots, n_{Ki})$, where n_{ki} is the number of class k packets in the queue. The state of the network is described by $n = (n_1, \dots, n_I)$. The network has a product-form equilibrium distribution that is given by Kelly [1979]:

$$p(n) = p(0) \left(\prod_{i=1}^I \lambda_{i1}^1 \cdots \prod_{i=1}^I \lambda_{iK}^K \right) \prod_{i=1}^I p_i(n_i), \tag{2.1}$$

where $l_k = \sum_{i=1}^I n_{ki}$ is the number of class k packets in the network when its state is n , $p(0)$ is the normalization constant, and

$$p_i(n_i) = \begin{cases} \frac{(n_{1i} + \dots + n_{Ki})!}{n_{1i}! \cdots n_{Ki}!} \prod_{k=1}^K \left(\frac{\theta^{ki}}{\mu^{ki}} \right)^{n_{ki}}, & \text{if } i \in \mathcal{S}_1 \\ \frac{(n_{1i} + \dots + n_{Ki})!}{n_{1i}! \cdots n_{Ki}!} \prod_{k=1}^K \left(\frac{\theta^{ki}}{\mu^{ki}} \right)^{n_{ki}} \frac{1}{n_{ki}!}, & \text{if } i \in \mathcal{S}_2. \end{cases} \tag{2.2}$$

Let us now introduce the following:

Definition 2.1. For given $N_k \in \mathbb{N}$ and $c^k \in \mathbb{R}_+$, the set of *admissible* controls $\lambda^k = (\lambda_{i1}^k)_{i=0}^{N_k-1}$ for user k is defined by:

$$\Lambda^k = \left\{ \lambda^k \in \mathbb{R}^{N_k} : 0 \leq \lambda_{i1}^k \leq c^k, \quad 0 \leq l_k < N_k \right\}.$$

Λ^k is also called the *strategy space* of user k , and a $\lambda^k \in \Lambda^k$ is called a (flow control) *strategy* of user k . The Cartesian product $\Lambda = \otimes_{k=1}^K \Lambda^k$ of the individual strategy spaces is termed the *product strategy space*. Strategy K -tuples $\lambda = (\lambda^1, \dots, \lambda^K) \in \Lambda$ are called *strategy profiles*.

The performance objective of user k is to find an admissible flow control strategy λ^{*k} in Λ^k that maximizes its average throughput $E\gamma^k$, subject to an upper bound $T^k > 0$ on its average time-delay $E\tau^k$. This optimization problem depends on the strategies of the other users, described by the strategy profile $\lambda^{-k} = (\lambda^1, \dots, \lambda^{k-1}, \lambda^{k+1}, \dots, \lambda^K)$, since $E\gamma^k$ and $E\tau^k$ depend on these strategies as well. Note that user k can trivially satisfy its time-delay constraint, by using the admissible strategy $\lambda_{l_k}^k = 0, l_k = 0, \dots, N_k - 1$, that is, by sending no flow into the network, since in that case $E\tau^k = 0 < T^k$, independently of the actions of the other users. Under this strategy, the throughput of the user is also zero, that is, $E\gamma^k = 0$. Let Λ^{-k} denote the product strategy space of all users except the k th. To each strategy profile λ^{-k} in Λ^{-k} we can assign the (possibly empty) set $R^k(\lambda^{-k})$ of optimal strategies of user k , defining this way a *point-to-set correspondence* [Border 1985] from Λ^{-k} to Λ^k .

Definition 2.2. The point-to-set correspondence $R^k: \Lambda^{-k} \rightarrow \Lambda^k$, that maps each $\lambda^{-k} \in \Lambda^{-k}$ to the set of all $\lambda^{*k} \in \Lambda^k$ that achieve the maximum:

$$\max_{E\tau^k \leq T^k} E\gamma^k \quad (2.3)$$

is called the *best reply correspondence* of user k . The correspondence $R: \Lambda \rightarrow \Lambda$, with:

$$R(\lambda) = \bigotimes_{k=1}^K R^k(\lambda^{-k}), \quad \lambda \in \Lambda, \quad (2.4)$$

is called the best reply correspondence of the flow control game.

In Section 3.2, it will be shown that $R^k(\lambda^{-k})$ is nonempty for all $\lambda^{-k} \in \Lambda^{-k}$, that is, the user always has an optimal response to any strategy profile of the other users.

Under the typical game-theoretic approach, the preferences (in terms of strategy profiles in Λ) of user k are expressed by means of a utility function, $U^k: \Lambda \rightarrow \mathbb{R}$, that assigns to each strategy profile $\lambda \in \Lambda$ the “payoff” $U^k(\lambda)$ to the user when this profile is in effect. A *Nash equilibrium* of the game is a strategy profile $\lambda^* \in \Lambda$, from which no user finds it beneficial to unilaterally deviate, such that:

$$U^k(\lambda^{*k}, \lambda^{*, -k}) = \max_{\lambda^k \in \Lambda^k} U^k(\lambda^k, \lambda^{*, -k}), \quad 1 \leq k \leq K.$$

The presence of the time-delay constraints does not allow a direct application of this conventional approach to the flow control game. However, the time-delay constraint of each user can be enforced as a penalty on its utility function, such that the constraint is not violated at any Nash equilibrium. In particular, the performance objective of user k can be represented by the following utility function:

$$U^k(\lambda) = \begin{cases} E\gamma^k, & \text{if } E\tau^k \leq T^k \\ -\epsilon, & \text{if } E\tau^k > T^k \end{cases}, \quad \lambda \in \Lambda, \quad (2.5)$$

where $\epsilon > 0$ and $E\gamma^k$ and $E\tau^k$ are the average throughput and time-delay of the user that correspond to the strategy profile λ . The assignment of negative utility for the case $E\tau^k > T^k$ implies that the optimal response of user k to

any strategy profile $\lambda^{-k} \in \Lambda^{-k}$ of the other users never leads to a violation of its time-delay constraint, since the user can receive a zero payoff by sending no flow into the network. Hence:

$$\max_{\lambda^k \in \Lambda^k} U^k(\lambda^k, \lambda^{-k}) = \max_{E\tau^k \leq T^k} E\gamma^k. \quad (2.6)$$

Therefore, the best reply correspondence of the game (defined by means of the user utility functions (2.5)) is precisely the one given by Definition 2.2. It is easy to see that a strategy profile $\lambda^* \in \Lambda$ is a Nash equilibrium if and only if it is a fixed point of the best reply correspondence, that is, if and only if $\lambda^* \in R(\lambda^*)$. Note that eq. (2.6) indicates that any Nash equilibrium is a strategy profile in the subset Λ' of the product strategy space where the time-delay constraints of all users are satisfied. Λ' is called the *joint strategy space*.

Let us now consider the problem of existence of equilibrium points of the flow control game. Since the average throughput of user k is always nonnegative, the utility function of the user, as defined in (2.5), is not continuous at any $\lambda \in \Lambda$, such that the user's time-delay holds tight, that is, $E\tau^k = T^k$.¹ Therefore, the classical equilibrium existence theorems for infinite games (games with strategy spaces that are not countable sets), for example, the theorems by Debreu, Glicksberg, and Fan (see, e.g., Fudenberg and Tirole [1992] and references therein), cannot be applied, because one of their main conditions requires that the utility functions of the users be continuous at every point in the product strategy space. It is worth mentioning that, typically, in a telecommunication network the average time-delay encountered by a user increases with its average throughput, and thus the Nash equilibria (if any) of the flow control game typically appear at the discontinuity points where the time-delay constraints of the users hold tight.

The existence of equilibria in games with discontinuous payoffs has been investigated by Dasgupta and Maskin [1986]. The main existence theorem of that reference relaxes the continuity requirement to two weaker conditions. It retains, however, the quasi-concavity assumption of the theorems by Debreu, Glicksberg and Fan. More precisely, it requires that the user's utility function be quasi-concave in its strategy.² In Section 3.2, we provide a simple example showing that enforcing the constraint of the user as in eq. (2.5) results in a utility function U^k that might fail to be quasi-concave in λ^k .

Hence, the conventional game-theoretic approach of incorporating the time-delay constraint of each user as an additional penalty to its payoff results in utility functions that do not satisfy, in general, the requirements of the standard existence theorems. However, since any Nash equilibrium lies in the joint strategy space, we can extend the definition $U^k(\lambda) = E\gamma^k$ for all $\lambda \in \Lambda$,³ but restrict the set of admissible strategy profiles to the joint strategy space Λ' . Under this approach, the time-delay constraints introduce a coupling of the individual strategy spaces, since the set of strategies that a user is allowed to

¹ These discontinuities cannot be removed, even if we define $U^k(\lambda) = 0$ for the case $E\tau^k > T^k$. As can be seen by the analysis in Section 3, $E\gamma^k > 0$ unless $\lambda_{l_k}^k = 0$ for all $l_k \geq 0$. Hence, for any λ such that $E\tau^k = T^k > 0$, we have $E\gamma^k > 0$.

² A function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is quasi-concave if and only if for any $x, y \in \mathbb{R}^n$ and any $\alpha \in [0, 1]$, we have $f(\alpha x + (1 - \alpha)y) \geq \min\{f(x), f(y)\}$.

³ This extension makes U^k continuous at every $\lambda \in \Lambda$, since $E\gamma^k$ is a continuous function of λ , as can be seen from the analysis in Section 3.

choose from is a subset of its strategy space and depends on the actions of the other users. Existence of equilibria for constrained games of this type is typically established based on Rosen's theorem [Rosen 1965], which requires that: (i) the joint strategy space be convex and compact, and (ii) the utility function of each user be continuous at every $\lambda \in \Lambda$ and concave in λ^k . For the flow control game, simple examples can be constructed (see Section B.1 in Appendix B) to demonstrate that, in general, the joint strategy space is not convex. Thus, Rosen's theorem cannot be applied to deduce existence of Nash equilibria.

The conditions that the existence theorems impose on the individual utility functions and/or the joint strategy space are *sufficient* conditions that guarantee—via the Kakutani fixed-point theorem [Border 1985]—that the best reply correspondence has a fixed point. Since the conditions of these theorems are not satisfied by the flow control model under consideration, we take a more general approach and study directly the existence of fixed points of its best reply correspondence. It is essential to realize that this correspondence encapsulates the constraints (see (2.3)), leaving, this way, the individual strategy spaces uncoupled. The structure of the best reply correspondence has been characterized in Hsiao and Lazar [1991] and is briefly presented in the following section.

3. The Structure of the Best Reply Correspondence

Suppose that the flow control strategies of all users except the k th are fixed and described by the profile λ^{-k} . The best reply of user k to this strategy profile is the set (Definition 2.2):

$$R^k(\lambda^{-k}) = \arg \max_{\lambda^k \in \Lambda^k, E\tau^k \leq T^k} E\gamma^k. \quad (3.1)$$

To find the structure of $R^k(\lambda^{-k})$, we need analytic expressions for the average throughput and time-delay of user k . The properties of the average throughput are presented in Subsection 3.1. These properties are used in Subsection 3.2 to transform the optimization problem on the right-hand side of eq. (3.1) into a linear program and characterize the structure of the user's best reply correspondence.

3.1. THE STRUCTURE OF THE AVERAGE THROUGHPUT. Let $\nu_{l_k}^k$ denote the conditional expectation of the rate at which class k packets are served by the network, given that there are l_k such packets in it. If $p_{l_k}^k$ is the probability that the number of class k packets in the network is l_k , then the average throughput $E\gamma^k$ of user k and the average number of packets EQ^k that the user has in the network are:

$$E\gamma^k = \sum_{l_k=1}^{N_k} \nu_{l_k}^k p_{l_k}^k = \sum_{l_k=0}^{N_k-1} \lambda_{l_k}^k p_{l_k}^k, \quad EQ^k = \sum_{l_k=1}^{N_k} l_k p_{l_k}^k. \quad (3.2)$$

The average time-delay of the user is given by Little's formula [Kelly 1979; Schwartz 1987]:

$$E\tau^k = \frac{EQ^k}{E\gamma^k} = \sum_{l_k=1}^{N_k} l_k p_{l_k}^k \left[\sum_{l_k=1}^{N_k} \nu_{l_k}^k p_{l_k}^k \right]^{-1}. \quad (3.3)$$

In the sequel, we derive expressions for $\nu_{i_k}^k$ and $p_{i_k}^k$ and we present a structural result called *the separation principle*. This principle allows the computation of $E\gamma^k$ by means of a simplified equivalent queueing model.

Define $\mathcal{N}(l_1, \dots, l_K)$ to be the set of all network states n , for which the total number of class k packets in the network is l_k , $1 \leq k \leq K$, that is:

$$\mathcal{N}(l_1, \dots, l_K) = \left\{ n = (n_1, \dots, n_I) : \sum_{i=1}^I n_{ki} = l_k, \quad 1 \leq k \leq K \right\}.$$

The probability that there are l_k class k packets in the network is then:

$$p_{i_k}^k = \sum_{\substack{l_m=0 \\ m \neq k}}^{N_m} \sum_{n \in \mathcal{N}(l_1, \dots, l_K)} p(n), \quad l_k = 0, \dots, N_k. \quad (3.4)$$

The expected rate of class k packets leaving the network through node i , given that the state of the node is n_i , is:

$$q^{ki}(n_i) = \begin{cases} \mu^{ki} \left(n_{ki} / \sum_m n_{mi} \right) 1 \left\{ \sum_m n_{mi} > 0 \right\} r^{ki0}, & \text{if } i \in \mathcal{S}_1 \\ \mu^{ki} n_{ki} r^{ki0}, & \text{if } i \in \mathcal{S}_2. \end{cases} \quad (3.5)$$

Equation (3.5) is immediate for the ./GI/1-PS and ./GI/ ∞ cases. For an ./M/1-FCFS or a ./GI/1-LCFS queue, given that its state is n_i , the packet classes in the queue are independent [Kelly 1979] and, provided that the queue is not empty, $n_{ki} / \sum_m n_{mi}$ is the probability that the packet in service is of class k . Note that $\sum_i q^{ki}(n_i)$ is the expected service rate of class k packets, given that the network is in state $n = (n_1, \dots, n_I)$.

The conditional expectation of the rate at which class k packets are served by the network, given that there are l_k ($1 \leq l_k \leq N_k$) such packets, is:⁴

$$\begin{aligned} \nu_{i_k}^k &= \sum_{\substack{l_m=0 \\ m \neq k}}^{N_m} \sum_{n \in \mathcal{N}(l_1, \dots, l_K)} \sum_{i=1}^I q^{ki}(n_i) \frac{p(n)}{p_{i_k}^k} \\ &= \sum_{\substack{l_m=0 \\ m \neq k}}^{N_m} \prod_{\substack{l_i=0 \\ i \neq k}}^{l_i-1} \lambda_{i_i}^v \sum_{n \in \mathcal{N}(l_1, \dots, l_K)} \sum_{i=1}^I q^{ki}(n_i) \prod_{j=1}^I p_j(n_j) \\ &\quad \times \left[\sum_{\substack{l_m=0 \\ m \neq k}}^{N_m} \prod_{\substack{l_i=0 \\ i \neq k}}^{l_i-1} \lambda_{i_i}^v \sum_{n \in \mathcal{N}(l_1, \dots, l_K)} \prod_{j=1}^I p_j(n_j) \right]^{-1}, \end{aligned} \quad (3.6)$$

where we have made use of eqs. (3.4) and (2.1), and $p_i(n_i)$ and $q^{ki}(n_i)$ are given by eqs. (2.2) and (3.5), respectively.

From eq. (3.6), we observe that $\nu_{i_k}^k$ does not depend on the flow control strategy of user k and is only a function of λ^{-k} . In a Markovian network, the marginal probabilities $p_{i_k}^k$ and the conditional estimates $\nu_{i_k}^k$ satisfy [Hsiao and Lazar 1989]:

$$\lambda_{i_k}^k p_{i_k}^k = \nu_{i_{k+1}}^k p_{i_{k+1}}^k, \quad l_k = 0, \dots, N_k - 1. \quad (3.7)$$

⁴ By convention, null summations are equal to zero and null products equal to one.




FIG. 2. Equivalent network model associated with user k .

This property can be verified for the network model of Figure 1, using eqs. (3.4) and (3.6).⁵ Note that eqs. (3.7) are precisely the detailed balance equations of the birth-death process that describes the state of an M/M/1 queue with state-dependent service rate $\nu_{l_k}^k$ and arrival rate $\lambda_{l_k}^k$. Therefore, the average throughput of user k is equal to the throughput of the queueing model depicted in Figure 2, which is usually called *Norton's equivalent* [Schwartz 1987]. By abuse of terminology, the term, *Norton's equivalent*, will be used to denote both the equivalent model in Figure 2 and the vector of conditional estimates $\nu^k = (\nu_{l_k}^k)_{l_k=1}^{N_k}$.

This structural result, called *the separation principle*, was used in Hsiao and Lazar [1991] to transform the optimization problem in eq. (3.1) into a linear program. The linear programming formulation will be presented in the following subsection.

Observe that the expression in the denominator of $\nu_{l_k}^k$ in eq. (3.6) is positive for all λ^{-k} , since the term corresponding to $l_m = 0$ for all $m \neq k$ is a positive constant (independent of λ^{-k}). Thus, $\nu_{l_k}^k$ is a *continuous* function of λ^{-k} . Continuity of $\nu_{l_k}^k$ over the compact set Λ^{-k} implies that it attains its minimum and maximum values:

$$\underline{\nu}_{l_k}^k = \min_{\lambda^{-k} \in \Lambda^{-k}} \nu_{l_k}^k > 0, \quad \bar{\nu}_{l_k}^k = \max_{\lambda^{-k} \in \Lambda^{-k}} \nu_{l_k}^k. \quad (3.8)$$

Positivity of $\underline{\nu}_{l_k}^k$ follows from the observation that the numerator of the expression in eq. (3.6) is positive for all λ^{-k} .

3.2. THE STRUCTURE OF THE BEST REPLY CORRESPONDENCE. In this section, we show that the best reply $R^k(\lambda^{-k})$ of user k to any strategy profile λ^{-k} of the other users is a set of “band-bang” control strategies.

Using eqs. (3.2) and (3.3), the best of user k can be written as:

$$R^k(\lambda^{-k}) = \arg \max_{\lambda^k \in \Lambda^k} \sum_{l_k=0}^{N_k-1} \lambda_{l_k}^k p_{l_k}^k, \quad \text{subject to } \sum_{l_k=1}^{N_k} (l_k - T^k \nu_{l_k}^k) p_{l_k}^k \leq 0 \quad (3.9)$$

⁵ The proof is identical to the one in Hsiao and Lazar [1991], for a network of M/M/1 queues, and is not repeated here.

where:

$$p_{l_k}^k = \prod_{i=0}^{l_k-1} \frac{\lambda_i^k}{\nu_{i+1}^k} \left[1 + \sum_{j=1}^{N_k} \prod_{i=0}^{j-1} \frac{\lambda_i^k}{\nu_{i+1}^k} \right]^{-1}, \quad l_k = 0, \dots, N_k, \quad (3.10)$$

by virtue of eq. (3.7). Note that $R^k(\lambda^{-k})$ depends on λ^{-k} only through the conditional estimates $\nu_{l_k}^k$. These estimates are independent of the k th user's flow control strategy (see eq. (3.6)) and, therefore, can be obtained in practice by means of an appropriate estimation technique.

The constrained maximization problem on the right-hand side of eq. (3.9) is nonlinear in λ^k . Nonetheless, if we define the equilibrium flows [Hsiao and Lazar 1991]:

$$y_{l_k}^k = \lambda_{l_k}^k p_{l_k}^k, \quad 0 \leq l_k < N_k, \quad (3.11)$$

this nonlinear optimization problem can be transformed into the following linear program [Gale 1960; Luenberger 1984] with unknown vector⁶ $(y^k, p^k) \in \mathbb{R}^{2N_k+1}$, where $y^k = (y_0^k, \dots, y_{N_k-1}^k)$, $p^k = (p_0^k, \dots, p_{N_k}^k)$:

$$\max y_1^k + \dots + y_{N_k-1}^k \quad (3.12)$$

subject to:

$$y_{l_k}^k = \nu_{l_k+1}^k p_{l_k+1}^k, \quad l_k = 0, \dots, N_k - 1, \quad (3.13)$$

$$p_0^k + p_1^k + \dots + p_{N_k}^k = 1, \quad (3.14)$$

$$y_{l_k}^k - c^k p_{l_k}^k \leq 0, \quad l_k = 0, \dots, N_k - 1, \quad (3.15)$$

$$\sum_{l_k=1}^{N_k} (l_k - T^k \nu_{l_k}^k) p_{l_k}^k \leq 0, \quad (3.16)$$

$$p_{l_k}^k \geq 0, \quad l_k = 0, \dots, N_k, \quad (3.17)$$

$$y_{l_k}^k \geq 0, \quad l_k = 0, \dots, N_k - 1. \quad (3.18)$$

Constraints (3.13) represent the detailed balance equations (3.7). Constraint (3.15) is the peak constraint on the control which, together with the nonnegativity constraints, assures that $\lambda^k \in \Lambda^k$. Thus, the linear program is equivalent to the maximization problem in (3.9).

The linear program (3.12)–(3.18) is parametrized by λ^{-k} , the strategy profile of the other users, through the Norton's equivalent ν^k . Let $\mathcal{L}^k: \Lambda^{-k} \rightarrow \mathbb{R}^{2N_k+1}$ denote the point-to-set correspondence that maps each strategy profile λ^{-k} to the set of optimal solutions of the linear program. If (y^k, p^k) is a point in $\mathcal{L}^k(\lambda^{-k})$, then λ^k , with:

$$\lambda_{l_k}^k = \begin{cases} \frac{y_{l_k}^k}{p_{l_k}^k}, & \text{if } p_{l_k}^k > 0 \\ 0, & \text{if } p_{l_k}^k = 0 \end{cases}, \quad 0 \leq l_k < N_k, \quad (3.19)$$

is a strategy in the k th user's best reply $R^k(\lambda^{-k})$.⁷

⁶ Throughout this paper, all vectors are assumed to be column-vectors. Since transposition of column-vectors, when necessary, is implied by the context, it is not denoted explicitly.

⁷ Note that the definition $\lambda_{l_k}^k = 0$ if $p_{l_k}^k = 0$, although arbitrary, is a natural one, since user k never has l_k (or more) packets in the network.

A linear program either possesses an optimal solution, or is infeasible, or has an unbounded objective value. Note that the linear program described by (3.12)–(3.18) is feasible for any λ^{-k} : $y^k = (0, \dots, 0)$, $p^k = (1, 0, \dots, 0)$ is a feasible solution, since $T^k > 0$. Furthermore, its objective value (the average throughput of user k) is bounded between 0 and c^k . Therefore, the linear program always has an optimal solution, that is, user k has a best reply to any strategy profile of the other users. The following lemma [Hsiao and Lazar 1991] characterizes the structure of any such optimal strategy.

LEMMA 3.2.1. *Let λ^k be a best reply of user k to a given strategy profile λ^{-k} of the other users. Then there exists some integer L_k , $1 \leq L_k \leq N_k$, such that:*

$$\lambda_{l_k}^k = \begin{cases} 0, & L_k \leq l_k < N_k \\ \lambda_{m_k}^k \in (0, c^k), & \text{at most one } m_k, 0 \leq m_k < L_k, \\ c^k, & 0 \leq l_k < L_k, l_k \neq m_k \end{cases} \quad (3.20)$$

that is, λ^k is of a generalized window-type, with at most one “intermediate” point, at some m_k . In particular, the intermediate point appears only if the time-delay achieved by λ^k is equal to T^k .

Remark 3.2.2. Let λ^k be as in eq. (3.20). Then the peak constraints (3.15) of the linear program for $0 \leq l_k < L_k$, $l_k \neq m_k$, are binding (hold as equalities), while the peak constraint for $l_k = m_k$ is nonbinding. Thus, the last assertion of Lemma 3.2.1 implies that, at optimality, at most one of the peak constraints (3.15) for $l_k = 0, \dots, L_k - 1$, or the time-delay constraint (3.16) can be nonbinding.

If a Nash equilibrium exists, then it is a profile of window-type strategies, as described by (3.20). Existence of equilibria was established in Hsiao and Lazar [1991] only for the special class of product-form networks, where the equivalent service rate $\nu_{l_k}^k$ is concave increasing in l_k , for every user k . Under this assumption, the user’s optimal strategy is unique [Lazar 1983], and thus its best reply correspondence is a point-to-point mapping, that is, a function. Furthermore, any intermediate point of the user’s best reply appears at the end of the window ($m_k = L_k - 1$). Nevertheless, simple examples show that these monotonicity properties do not necessarily characterize product-form networks. Consider, for instance, a single $M/1$ -FCFS queue with service rate $\mu = 1$ shared by two users, and take $N_1 = N_2 = 2$ and $\lambda^2 = (1/19, 2)$. Then, eq. (B.2), in Appendix B, gives $\nu_1^1 = 0.814 > 0.794 = \nu_2^1$.

This lack of monotonicity is inherited by the average throughput and time-delay of the user and does not allow the application of standard existence results from the game theory literature to the flow control game. The following example shows that the utility function of user k , as defined in (2.5), might fail to be quasi-concave in λ^k .

Example 3.2.3. Let $N_k = 3$, $c^k = 3$ and $T^k = 1.51$. Fix a strategy profile λ^{-k} of the other users, such that, $\nu_1^k = \nu_2^k = 1$ and $\nu_3^k = 2$. Consider now the following strategies of user k : $\lambda^k(1) = (3, 1, 0)$ and $\lambda^k(2) = (3, 3, 0.65)$. The respective average throughputs and time-delays of user k are $E\gamma^k(1) = 0.857$, $E\tau^k(1) = 1.5$ and $E\gamma^k(2) = 1.194$, $E\tau^k(2) = 1.504$, as can be seen by eqs. (3.2), (3.3), and (3.10). Then, by eq. (2.5), we have $U^k(\lambda^k(1), \lambda^{-k}) = 0.857$ and

$U^k(\lambda^k(2), \lambda^{-k}) = 1.194$. Consider the strategy $\lambda^k(3) = 0.5(\lambda^k(1) + \lambda^k(2))$, for which $E\gamma^k(3) = 1.028$, $E\tau^k(3) = 1.548$, thus $U^k(\lambda^k(3), \lambda^{-k}) = -\epsilon < 0$. Therefore, $U^k(\lambda^k(3), \lambda^{-k}) < \min\{U^k(\lambda^k(1), \lambda^{-k}), U^k(\lambda^k(2), \lambda^{-k})\}$, that is, U^k is not quasi-concave in λ^k . A closer investigation of this example reveals that the average throughput itself is concave in the linear segment $(\lambda^k(1), \lambda^k(2))$. The time-delay, however, first increases to a maximum and then decreases in that segment resulting, this way, in the non-quasi-concavity of U^k .

In the following section, the uniqueness result of Lazar [1983] is extended to the general case where no monotonicity assumptions are made on the equivalent service rate. This result will be used in subsequent sections to substantially simplify the overall proof of existence of a Nash equilibrium point.

4. The Best Reply Correspondence Is a Function

In this section, we show that the best reply of user k to any strategy profile of the other users is unique and, therefore, its best reply correspondence is a function. This amounts to showing that the optimal flow control strategy of the single-class model depicted in Figure 2 is unique. Throughout this section, we temporarily drop the sub/superscripts that denote the particular user to simplify the notation.

PROPOSITION 4.1. *The optimal flow control strategy for the model of Figure 2 is unique.*

PROOF. Suppose that there exist two distinct optimal flow control strategies, say $\lambda(1)$ and $\lambda(2)$. Let $(y(1), p(1))$ and $(y(2), p(2))$ be the corresponding optimal solutions of the linear program. Then, by linearity, $(y(3), p(3)) = 0.5(y(1), p(1)) + 0.5(y(2), p(2))$ is also an optimal solution. Let $E\tau(i)$ denote the average time-delay achieved by $(y(i), p(i))$, $i = 1, 2, 3$. All optimal solutions achieve the same average throughput. If L_i is the window size (see Lemma 3.2.1) resulting from the optimal solution $(y(i), p(i))$, it is easy to see that $L_3 = \max\{L_1, L_2\}$.

If a peak constraint is nonbinding for one of the optimal solutions, say $(y(1), p(1))$, the same constraint is nonbinding for $(y(3), p(3))$. To see this, assume that $y_m(1) - cp_m(1) < 0$, for some $m < L_1$. Then, $y_m(3) - cp_m(3) = 0.5(y_m(1) - cp_m(1)) + 0.5(y_m(2) - cp_m(2)) < 0$. Similarly, if the time-delay constraint is nonbinding for one of $(y(i), p(i))$, $i = 1, 2$, it is nonbinding for $(y(3), p(3))$ as well. We have to consider the following cases:

- (i) *At least one of $E\tau(1)$ and $E\tau(2)$ is less than T .*

Then $E\tau(3) < T$. Therefore, both $\lambda(1)$ and $\lambda(2)$ must have *full* windows, with no intermediate points.⁸ Since these optimal solutions are distinct, their windows are not equal. Suppose that $L_1 < L_2 \leq N$. Then:

$$\begin{aligned} y_{L_1}(3) - cp_{L_1}(3) &= 0.5(y_{L_1}(1) - cp_{L_1}(1)) + 0.5(y_{L_1}(2) - cp_{L_1}(2)) \\ &= -0.5cp_{L_1}(1) < 0, \end{aligned}$$

since:

$$y_{L_1}(1) = 0, \quad p_{L_1}(1) > 0, \quad y_{L_1}(2) = cp_{L_1}(2).$$

⁸ If one of them, say $\lambda(1)$, had an intermediate point, $(y(1), p(1))$ would lead to a nonbinding peak constraint of the linear program and the same constraint would be nonbinding for $(y(3), p(3))$, contradicting this way optimality of $(y(3), p(3))$: see Remark 3.2.

This implies that both the time-delay constraint and the peak constraint for $l = L_1 < L_3$ are non-binding at $(y(3), p(3))$. Thus, according to Remark 3.2.2, $(y(3), p(3))$ cannot be an optimal solution of the linear program.

(ii) $E\tau(1) = E\tau(2) = T$.

Then $E\tau(3) = T$. First suppose that $\lambda(1)$ has an intermediate point at $l = m$, $0 \leq m < N$. If $\lambda(2)$ had an intermediate point at a different l , $(y(3), p(3))$ would lead to two nonbinding peak constraints of the linear program and would not be an optimal solution. Therefore, any intermediate point of $\lambda(2)$ should appear at $l = m$. A similar argument shows that the corresponding window sizes must be equal: $L_1 = L_2 = L$. Then, $\lambda(1) \neq \lambda(2)$ implies:

$$\lambda_m(1) \neq \lambda_m(2). \quad (4.1)$$

Define $\phi_l = \nu_1 \cdots \nu_l$, for $l \geq 1$, and $\phi_0 = 1$. Then:

$$p_l(i) = \begin{cases} p_0(i) \cdot \frac{c^l}{\phi_l}, & l = 1, \dots, m \\ p_0(i) \lambda_m(i) \cdot \frac{c^{l-1}}{\phi_l}, & l = m + 1, \dots, L \end{cases}, \quad i = 1, 2,$$

and the average queue length resulting from $\lambda(i)$ is:

$$EQ(i) = \left[\sum_{l=0}^m l \frac{c^l}{\phi_l} + \lambda_m(i) \sum_{l=m+1}^L l \frac{c^{l-1}}{\phi_l} \right] \times \left[\sum_{l=0}^m \frac{c^l}{\phi_l} + \lambda_m(i) \sum_{l=m+1}^L \frac{c^{l-1}}{\phi_l} \right]^{-1}. \quad (4.2)$$

Since $E\tau(1) = E\tau(2) = T$ and $E\gamma(1) = E\gamma(2)$, then $EQ(1) = EQ(2)$. In view of (4.1) and (4.2), the latter can hold only if:

$$\sum_{l=0}^m l \frac{c^l}{\phi_l} \sum_{j=m+1}^L \frac{c^{j-1}}{\phi_j} = \sum_{l=0}^m \frac{c^l}{\phi_l} \sum_{j=m+1}^L j \frac{c^{j-1}}{\phi_j},$$

or equivalently:

$$\sum_{l=0}^m \sum_{j=m+1}^L \frac{c^l}{\phi_l} \frac{c^{j-1}}{\phi_j} (j - l) = 0. \quad (4.3)$$

Recalling that $c > 0$, we conclude that eq. (4.3) is a contradiction, since the expression on its left-hand side is always positive.

Finally, consider the case where both $\lambda(1)$ and $\lambda(2)$ have full windows and assume that $L_1 < L_2$. A similar argument, based on the equality of the average queue lengths, can be used to obtain a contradiction to $\lambda(1) \neq \lambda(2)$. \square

Proposition 4.1, together with Lemma 3.2.1, generalizes the results of Lazar [1983] for the flow control model of Figure 2, to the case where no assumptions are made on the monotonicity properties of the state-dependent service rate. The optimal flow control strategy is *unique* and has a generalized window

structure with at most one intermediate point. Under the assumptions of Lazar [1983] the intermediate point can only appear at the end of the window. The following example shows that if the Norton's equivalent is not concave increasing, the optimal strategy may have an intermediate point that does not occur at the end of the window. Thus, the structure of the optimal strategy is completely characterized by Lemma 3.2.1 and Proposition 4.1.

Example 4.2. Let $N = 5$, $c = 3$, $T = 1$, and $\nu_1 = 2$, $\nu_2 = 1$, $\nu_3 = 3$, $\nu_4 = \nu_5 = 2$. Using the linear programming formulation, one can see that the optimal flow control strategy is $\lambda_0 = \lambda_2 = 3 = c$, $\lambda_1 = 1 < c$, $\lambda_3 = \lambda_4 = 0$. The time-delay constraint is binding at this optimal solution, that is, $E\tau = 1 = T$.

Extending Proposition 4.1 to the multiclass model of Figure 1, we conclude that the best reply correspondence of user k is a *best reply function* $R^k: \Lambda^{-k} \rightarrow \Lambda^k$ that maps each strategy profile λ^{-k} to the unique best reply of user k . Consequently, the best reply correspondence of the flow control game is also a function $R: \Lambda \rightarrow \Lambda$ and its definition in (2.4) can be rewritten as:

$$R(\lambda) = (R^1(\lambda^{-1}), \dots, R^K(\lambda^{-K})), \quad \lambda \in \Lambda. \quad (4.4)$$

A Nash equilibrium is a fixed point of this best reply function, that is, a strategy profile $\lambda^* \in \Lambda$, such that $\lambda^* = R(\lambda^*)$.

5. Continuity of the Best Reply Function

Sufficient conditions for a function to have a fixed point are provided by Brouwer's theorem [Border 1985]. The theorem asserts that a *continuous* function mapping a convex and compact set into itself has a fixed point. The product strategy space Λ satisfies the convexity and compactness requirements of Brouwer's theorem. Therefore, if the best reply function $R: \Lambda \rightarrow \Lambda$ is continuous, it has a fixed point which is a Nash equilibrium of the flow control game. In this section we show that, under a certain regularity condition, R is indeed continuous.

The best reply function of user k is defined implicitly, by means of the linear program (3.12)–(3.18). Specifically, λ^k is the best reply to λ^{-k} , if and only if the equilibrium distribution p^k and the equilibrium flows y^k resulting from $(\lambda^k, \lambda^{-k})$ are the optimal solution of the linear program. Let $f^k: \Lambda \rightarrow \mathbb{R}^{2N_k+1}$ denote the function that assigns to each $(\lambda^k, \lambda^{-k})$ the corresponding equilibrium distribution and equilibrium flows, according to eqs. (3.6), (3.10), and (3.11). Then, for all $(\lambda^k, \lambda^{-k}) \in \Lambda$:

$$\lambda^k = R^k(\lambda^{-k}) \Leftrightarrow f^k(\lambda^k, \lambda^{-k}) = \mathcal{L}^k(\lambda^{-k}). \quad (5.1)$$

Clearly, the continuity properties of R^k are closely related to those of \mathcal{L}^k .⁹

The continuity of \mathcal{L}^k is treated in a general setting in Appendix A. In that appendix we consider a parametrized linear program and investigate the continuity properties of the set of its optimal solutions. We employ a definition of continuity for the general case, that includes the possibility of multiple optimal solutions, and derive sufficient conditions for continuity in Theorem A.3. These conditions require that the *dual* linear program exhibits certain properties.

⁹ Throughout this section \mathcal{L}^k —defined as a point-to-set correspondence in Section 3.2—will be treated as a *function*, by virtue of Proposition 4.1.

By abuse of terminology, the dual linear program associated with (3.12)–(3.18) will be referred to as the dual flow control problem. The dual problem is also parametrized by the strategy profile λ^{-k} of the other users through the Norton's equivalent ν^k . Let $\mathcal{D}^k(\lambda^{-k})$ denote the set of its optimal solutions. Then, Corollary A.4 in Appendix A implies that \mathcal{D}^k is continuous if the following conditions are satisfied:

- (i) There exists a compact set $\Theta^k \subset \mathbb{R}^{2N_k+1}$, such that $\mathcal{D}^k(\lambda^{-k}) \in \Theta^k$, for all $\lambda^{-k} \in \Lambda^{-k}$.
- (ii) There exists a compact set $W^k \subset \mathbb{R}^{2N_k+2}$, such that $\mathcal{D}^k(\lambda^{-k}) \cap W^k \neq \emptyset$, for all $\lambda^{-k} \in \Lambda^{-k}$.

Any feasible solution (y^k, p^k) of the linear program described by (3.12)–(3.18) satisfies:

$$\begin{aligned} 0 \leq p_{l_k}^k \leq 1, & \quad 0 \leq l_k \leq N_k, \\ 0 \leq y_{l_k} \leq c^k, & \quad 0 \leq l_k < N_k, \end{aligned}$$

for all λ^{-k} and, by defining $\Theta^k = [0, c^k]^{N_k} \otimes [0, 1]^{N_k+1}$, condition (i) above automatically holds. In the sequel, we will derive a sufficient condition for the dual flow control problem to satisfy property (ii).

The dual flow control problem for user k is a minimization problem with unknown vector $w^k \in \mathbb{R}^{2N_k+2}$ described by:

$$\min w_{N_k+1}^k \quad (5.2)$$

subject to:

$$w_i^k + w_{i+N_k+1}^k \geq 1, \quad i = 1, \dots, N_k, \quad (5.3)$$

$$w_{N_k+1}^k - c^k w_{N_k+2}^k \geq 0, \quad (5.4)$$

$$- \nu_i^k w_i^k + w_{N_k+1}^k - c^k w_{i+N_k+2}^k + (i - \nu_i^k T^k) w_{2N_k+2}^k \geq 0,$$

$$i = 1, \dots, N_k - 1, \quad (5.5)$$

$$- \nu_{N_k}^k w_{N_k}^k + w_{N_k+1}^k + (N_k - \nu_{N_k}^k T^k) w_{2N_k+2}^k \geq 0, \quad (5.6)$$

$$w_1^k, \dots, w_{N_k+1}^k: \text{unrestricted},$$

$$w_{N_k+2}^k, \dots, w_{2N_k+2}^k \geq 0. \quad (5.7)$$

LEMMA 5.1. *If the upper bound T^k on the average time-delay of user k satisfies:*

$$\frac{1}{\nu_1^k} < T^k \quad (5.8)$$

then, there exists a compact set $W^k \subset \mathbb{R}^{2N_k+2}$, such that $\mathcal{D}^k(\lambda^{-k}) \subset W^k$, for all $\lambda^{-k} \in \Lambda^{-k}$.

PROOF. Fix λ^{-k} in Λ^{-k} and let w^k be an optimal solution in $\mathcal{D}^k(\lambda^{-k})$.

The objective values of the primal and the dual linear programs are equal at optimality. Therefore, $w_{N_k+1}^k$ is equal to the optimal throughput of user k , given that the strategy profile of the other users is λ^{-k} . Thus:

$$0 \leq w_{N_k+1}^k \leq c^k. \quad (5.9)$$

Then, (5.4) implies:

$$w_{N_k+2}^k \leq 1, \quad (5.10)$$

which, together with (5.3) for $i = 1$, gives:

$$w_1^k \geq 0. \quad (5.11)$$

Inequality (5.5) for $i = 1$ gives:

$$(\nu_1^k T^k - 1)w_{2N_k+2}^k \leq w_{N_k+1}^k \leq c^k, \quad (5.12)$$

since w_1^k and $w_{N_k+3}^k$ are nonnegative. Under condition (5.8), we have:

$$\frac{1}{\nu_1^k} \leq \frac{1}{\underline{\nu}_1^k} < T^k \Rightarrow \nu_1^k T^k - 1 \geq \underline{\nu}_1^k T^k - 1 > 0,$$

and (5.12) gives:

$$w_{2N_k+2}^k \leq \frac{c^k}{\underline{\nu}_1^k T^k - 1} \equiv \bar{w}_{2N_k+2}^k. \quad (5.13)$$

Note that $\underline{\nu}_1^k$, and therefore $\bar{w}_{2N_k+2}^k$, does not depend on λ^{-k} (see eq. (3.8)).

Multiplying inequality (5.3) by ν_i^k and adding it with (5.5), we obtain:

$$\nu_i^k w_{i+N_k+1}^k - \nu_i^k + w_{N_k+1}^k - c^k w_{i+N_k+2}^k + (i - \nu_i^k T^k)w_{2N_k+2}^k \geq 0,$$

which, in view of (5.7) and (5.13), implies:

$$w_{i+N_k+2}^k \leq 1 + \frac{\bar{\nu}_i^k}{c^k} w_{i+N_k+1}^k + \frac{i}{c^k} \bar{w}_{2N_k+2}^k, \quad i = 1, \dots, N_k - 1. \quad (5.14)$$

Since $w_{N_k+2}^k \leq 1$, we can recursively obtain upper bounds $\bar{w}_{i+N_k+2}^k$ on $w_{i+N_k+2}^k$ using (5.14), for $i = 1, \dots, N_k - 1$. These bounds do not depend on λ^{-k} , as $\bar{\nu}_i^k$ is independent of λ^{-k} .

Now, by inequality (5.3), we have:

$$w_i^k \geq 1 - \bar{w}_{i+N_k+1}^k, \quad i = 2, \dots, N_k.$$

Finally, (5.5) and (5.6), together with (5.9), give:

$$w_i^k \leq \frac{c^k}{\underline{\nu}_i^k} + \frac{i}{\underline{\nu}_i^k} \bar{w}_{2N_k+2}^k, \quad i = 1, \dots, N_k.$$

Thus, for all $i = 1, \dots, 2N_k + 2$, there exist $\underline{w}_i^k, \bar{w}_i^k$, independent of λ^{-k} , such that $\underline{w}_i^k \leq w_i^k \leq \bar{w}_i^k$. Defining:

$$W^k = \bigotimes_{i=1}^{2N_k+2} [\underline{w}_i^k, \bar{w}_i^k],$$

the result follows. \square

Note that the result of the lemma is stronger than the sufficient condition $\mathcal{D}^k(\lambda^{-k}) \cap W^k \neq \emptyset$, for all $\lambda^{-k} \in \Lambda^{-k}$. Using this result, continuity of \mathcal{L}^k is established in the following:

LEMMA 5.2. *If $1/\nu_1^k < T^k$, then $\mathcal{L}^k: \Lambda^{-k} \rightarrow \Theta^k$ is continuous.*

PROOF. The linear program (3.12)–(3.18) is *continuously* parametrized (Appendix A) by λ^{-k} , since the Norton's equivalent ν^k is a continuous function of λ^{-k} (Section 3.1). Furthermore, Θ^k is compact and in view of Lemma 5.1 the sufficient conditions of Corollary A.4 in Appendix A are satisfied. Thus, $\mathcal{L}^k: \Lambda^{-k} \rightarrow \Theta^k$ is continuous. \square

Let us define the set:

$$\mathcal{T}^k = \left\{ T^k \in \mathbb{R}_+ : \frac{1}{\underline{\nu}_1^k} < T^k \right\}.$$

We are ready to prove the following:

THEOREM 5.3. *If $T^k \in \mathcal{T}^k$, then the best reply function of user k , $R^k: \Lambda^{-k} \rightarrow \Lambda^k$, is continuous. Furthermore, if $F^k(\lambda^{-k})$ is the optimal throughput of user k when the other users' strategy profile is λ^{-k} , then $F^k: \Lambda^{-k} \rightarrow [0, c^k]$ is also continuous.*

PROOF. Since Λ^k is compact, according to Lemma A.2 in Appendix A, $R^k: \Lambda^{-k} \rightarrow \Lambda^k$ is continuous if and only if for every convergent sequence $\{(\lambda^k(n), \lambda^{-k}(n)), n \geq 0\}$ in Λ , with:

$$\lambda^k(n) = R^k(\lambda^{-k}(n)), \quad n \geq 0 \quad (5.15)$$

and

$$\lim_{n \rightarrow \infty} (\lambda^k(n), \lambda^{-k}(n)) = (\lambda^k, \lambda^{-k}),$$

we have:

$$\lambda^k = R^k(\lambda^{-k}).$$

From (5.1) and eq. (5.15), we obtain:

$$f^k(\lambda^k(n), \lambda^{-k}(n)) = \mathcal{L}^k(\lambda^{-k}(n)), \quad n \geq 0. \quad (5.16)$$

From eqs. (3.10) and (3.11) and continuity of the Norton's equivalent ν^k in λ^{-k} , it is easy to see that $f^k: \Lambda \rightarrow \mathbb{R}^{2N_k+1}$ is continuous. Furthermore, \mathcal{L}^k is continuous, by Lemma 5.2. Taking the limit as $n \rightarrow \infty$ on both sides of eq. (5.16), we get:

$$f^k(\lambda^k, \lambda^{-k}) = \mathcal{L}^k(\lambda^{-k}).$$

Then, (5.1) implies $\lambda^k = R^k(\lambda^{-k})$ and continuity of R^k follows.

Continuity of F^k in λ^{-k} can be proven using the continuity of the best reply function and the Norton's equivalent associated with user k . The details are left to the reader. \square

In Appendix B.2, we give a simple example of an ./M/1 queue shared by two users showing that the best reply function of user k might be discontinuous, if the upper bound on its average time-delay does not satisfy the condition $1/\underline{\nu}_1^k < T^k$. Therefore, this is a tight sufficient condition for continuity of R^k .

Remark 5.4. Recall, from eq. (3.8), that $\underline{\nu}_1^k = \min_{\lambda^{-k}} \nu_1^k > 0$ is the minimum rate at which user k is served in the network, given that it has one packet in it. Hence, $1/\underline{\nu}_1^k$ is the maximum time-delay of the user with one packet in the network. The sufficient condition (5.8) for continuity of R^k ensures that if user k sends just one packet in the network, that is, if $\lambda_0^k > 0$ and $\lambda_{i_k}^k = 0$ for

all $l_k \geq 1$, then its time-delay constraint is not violated, independently of the actions of the other users. As explained in Section 3.2, user k has a best reply to any strategy profile $\lambda^{-k} \in \Lambda^{-k}$ of the other users. Condition (5.8), then, implies that this best reply is such that the user sends at least one packet in the network. In other words, it is a regularity condition guaranteeing that the flow control problem has always the structure of a K -person game: there is no strategy profile of the other users, that can force user k to refrain from utilizing network resources, in order to satisfy its time-delay constraint. In fact, if user k can be forced by the actions of the other users to send no flow into the network, its best reply function is discontinuous. More specifically, using eq. (3.10) in eqs. (3.2) and (3.3), it can be easily verified that the average throughput of the user is increasing in λ_0^k , while its average time-delay does not depend on λ_0^k . Thus, as long as the user can satisfy its time-delay constraint, λ_0^k is equal to c^k , whereas inability of the user to satisfy its constraint implies $\lambda_0^k = 0$. Therefore, if the user is forced to “withdraw” from the game, λ_0^k drops from c^k to zero and its best reply function is discontinuous.

A condition similar to (5.8) has been postulated in Lazar [1983] for the time-delay constraint of the user. A condition that secures that no user can be forced by the actions of the others to withdraw from the game has also been imposed on the routing game of Orda et al. [1993]. It is important to notice that condition (5.8) does not imply that the time-delay constraint of the user will not hold tight at an equilibrium. Typically, the average time-delay of the user increases with the number of packets it sends into the network.¹⁰ Therefore, although the constraint is not effective if the user sends just one packet, it will eventually become tight as the user sends more packets in the network, in order to maximize its throughput.

Since the individual best reply functions are continuous (Theorem 5.3), the best reply function of the flow control game $R: \Lambda \rightarrow \Lambda$, defined by eq. (4.4), is also continuous and, by Brouwer’s theorem, it has a fixed point. This point is a Nash equilibrium of the flow control game. We have, therefore, proven the following:

THEOREM 5.5. *A Nash equilibrium point exists for the flow control problem of K users sharing a network of quasi-reversible queues, when the control of user k ($1 \leq k \leq K$) is state-dependent and bounded and its performance objective is to maximize its individual average throughput subject to an upper bound $T^k \in \mathcal{F}^k$ on its average time-delay.*

Remark 5.6. An immediate question that arises from the existence result of Theorem 5.5, is whether the flow control game has a unique or multiple Nash equilibria. In Appendix B.3, we provide a simple example of a two-user flow control game with multiple equilibria. The presence of multiple Nash equilibria is to be contrasted with the uniqueness results for the greedy algorithm in Bovopoulos and Lazar [1988] and Douligeris and Mazumdar [1988].

¹⁰ In single-class networks of $M/1$ queues, the average time-delay increases with the number of packets [Walrand 1988]. Although there is no proof of this monotonicity result for general multiclass queueing systems, this behavior is typical in telecommunication networks.

It is important to realize that, although the best reply correspondence of the flow control model considered in this paper is a function, uniqueness of the best reply for each user is not the key property that leads to the existence result of Theorem 5.5. Rather, the *continuity* of the best reply function plays the fundamental role in the existence of Nash equilibria. In general, the best reply correspondence of a congestion control scheme might not be a function. In such cases, Kakutani's theorem provides sufficient conditions for a point-to-set correspondence, mapping a convex and compact set into itself, to have a fixed point. Continuity in Brouwer's theorem is replaced by *closedness* (Definition A.1 in Appendix A) of the correspondence in Kakutani's theorem; there is also an additional condition that requires that the best reply of each user to any strategy profile of the other users be a nonempty and convex set. Since the set of optimal solutions of a linear program is convex, given that it is not empty, Theorem A.3 provides sufficient conditions for existence of equilibrium points in decentralized control schemes, where the control parameters are determined by means of a number of interconnected (mutually parametrized) linear programs.

6. Conclusions

In the present paper, we investigated the existence of Nash equilibria in noncooperative flow control of a general product-form network. Existence of equilibria was established, under a regularity condition guaranteeing that no user can be forced by the other users to withdraw from utilizing network resources. The network model is one of the most general analytically tractable models: a network of quasi-reversible queues. High-speed networks with propagation delays and nodal schedulers that approximate the processor sharing service discipline belong to this category. The user performance objective, defined as maximizing the utilization of network resources without violating certain quality of service constraints, is well suited for realistic networking environments. This is to be contrasted with other game theoretic studies in the congestion control literature, where it is a priori assumed that the users' utility functions satisfy first and second order monotonicity properties that allow a direct application of standard equilibrium existence results.

We proposed a general method to study the existence of equilibria in noncooperative congestion control schemes that is based on the best reply correspondence of the underlying game. This method relies on the most general mathematical techniques available to show existence of noncooperative equilibria, namely fixed-point theorems. Thus, it can be applied to constrained game models that do not satisfy the conditions of standard game theoretic results, for example, Rosen's theorem.

The problem of continuity of the optimal solution correspondence of a parametrized linear program was addressed. Theorem A.3 provides sufficient conditions for this correspondence to be closed and is of interest on its own right. More importantly, it can be readily applied to study the existence of equilibria in decentralized control schemes where the control parameters are determined via a collection of interdependent linear programs.

Several topics for future research arise from the existence result, such as characterizing the set of Nash equilibria and their properties (e.g., stability) and examining convergence of synchronous and asynchronous algorithms to an

equilibrium point. The formulation of the flow control problem around its best reply function provides a means to systematically study these topics. Moreover, the properties of the best reply function established in the present paper might prove to be an important starting point in this direction.

Appendix A. Linear Programming and Continuity

In this appendix, we study the problem of continuity in linear programming. Consider the linear program in standard form:

$$\max cy$$

subject to:

$$Ay = b, \quad y \geq 0,$$

where $y \in \mathbb{R}^m$ is the unknown vector and $c \in \mathbb{R}^m$, $b \in \mathbb{R}^k$ and A is a $k \times m$ matrix. Now suppose that the “constants” of the problem, that is, vectors c and b and matrix A , are functions of a parameter x that takes values in some set X . Denote the parametrized linear program by:

$$\max c(x)y \tag{A.1}$$

subject to:

$$A(x)y = b(x), \quad y \geq 0. \tag{A.2}$$

If the entries of matrix A and vectors c and b are continuous functions of x , we say that the linear program (A.1)–(A.2) is *continuously parametrized* by x .

Each value of the parameter x leads to a different linear program, according to (A.1)–(A.2). The set of optimal solutions of this linear program depends on the parameter x and is denoted by $\mathcal{L}(x)$. This way, we define a point-to-set correspondence $\mathcal{L}: X \rightarrow \rightarrow \mathbb{R}^m$, that maps each $x \in X$ to the set of optimal solutions of the linear program. Therefore, to address the issue of continuity in linear programming, we need an appropriate definition of continuity for point-to-set correspondences. Conceptually, continuity means that a small perturbation in the parameter x will result in a small change in the set of optimal solutions of the linear program. There are several definitions that generalize the concept of continuity from functions to correspondences [Takayama 1985]. In the present paper, we adopt the concept of a *closed* correspondence.

Definition A.1. Let X and Y be subsets of two (finite dimensional) Euclidean spaces. A point-to-set correspondence $R: X \rightarrow \rightarrow Y$ is said to be *closed* if its *graph*

$$\mathcal{G}R = \{(x, y) \in X \otimes Y: y \in R(x)\}$$

is a closed subset of $X \otimes Y$, that is, if for any convergent sequence $\{(x^n, y^n), n \geq 0\}$ in $\mathcal{G}R$:

$$\lim_{n \rightarrow \infty} (x^n, y^n) = (x, y) \in \mathcal{G}R.$$

A relation between closedness for correspondences and continuity for functions is given in the following lemma [Takayama 1985, Theorem 2.D.12]:

LEMMA A.2. A function $f: X \rightarrow Y$, where Y is a compact set, is closed if and only if f is continuous.

In the sequel, we derive sufficient conditions for a continuously parametrized linear program to have a closed graph. Continuity considerations are scarce in the linear-programming literature. Most of the related studies treat cases where either the objective vector c , or the right-hand-side vector b is varied linearly with the parameter and concentrate on the implications of this perturbation on the optimal basis and the simplex method; see, for example, Bazaraa et al. [1990, Chapter 6] and references therein. In a more general context, sufficient conditions for continuity of the set of solutions to a parametrized maximization problem are given by the “maximum theorem,” due to Berge [1963], which is one of the fundamental theorems employed in mathematical economics [Border 1985; Takayama 1985]. However, as will be explained in the sequel, the maximum theorem does not have the precise form needed in this paper.

Consider the parametrized linear program (A.1)–(A.2) and fix x in X . By the theory of duality in linear programming [Gale 1960; Luenberger 1984], $y \in \mathbb{R}^m$ is an optimal solution of the linear program, if and only if there exists some $w \in \mathbb{R}^k$, such that:

$$A(x)y = b(x), \quad (\text{A.3})$$

$$wA(x) \geq c(x), \quad (\text{A.4})$$

$$wb(x) = c(x)y, \quad (\text{A.5})$$

$$y \geq 0. \quad (\text{A.6})$$

If such a w exists, it is an optimal solution of the *dual* linear program:

$$\min wb(x)$$

subject to:

$$wA(x) \geq c(x), \quad w: \text{unrestricted.}$$

Define the point-to-set correspondence $\mathcal{D}: X \rightarrow \rightarrow \mathbb{R}^k$ that maps each $x \in X$ to the set of optimal solutions of the dual linear program. If there exists a pair of y and w that satisfy the optimality conditions (A.3)–(A.6), then $y \in \mathcal{L}(x)$ and $w \in \mathcal{D}(x)$. We are now ready to prove the following:

THEOREM A.3. *Consider a linear program continuously parametrized by $x \in X$, a subset of a finite dimensional Euclidean space. If there exists a compact set $W \subset \mathbb{R}^k$, such that:*

$$\mathcal{D}(x) \cap W \neq \emptyset, \quad (\text{A.7})$$

for all $x \in X$ satisfying $\mathcal{D}(x) \neq \emptyset$, then the point-to-set correspondence $\mathcal{L}: X \rightarrow \rightarrow \mathbb{R}^m$ is closed.

PROOF. Let $\{(x^n, y^n), n \geq 0\}$ be a convergent sequence in \mathcal{GL} , with limit (x, y) . We have to show that $(x, y) \in \mathcal{GL}$.

For all n , we have $y^n \in \mathcal{L}(x^n)$. Thus, there exists a w^n , such that the optimality conditions (A.3)–(A.6) are satisfied:

$$\begin{aligned} A(x^n)y^n &= b(x^n), & w^n A(x^n) &\geq c(x^n), \\ w^n b(x^n) &= c(x^n)y^n, & y^n &\geq 0. \end{aligned}$$

Condition (A.7) implies that, for all $n \geq 0$, w^n can be chosen to belong to W . Hence, $\{w^n\}$ is a sequence in a compact set and it has a convergent subsequence $\{w^{k_n}, n \geq 0\}$ with limit $w \in W$. Note that for all n :

$$\begin{aligned} A(x^{k_n})y^{k_n} &= b(x^{k_n}), & w^{k_n}A(x^{k_n}) &\geq c(x^{k_n}), \\ w^{k_n}b(x^{k_n}) &= c(x^{k_n}), & y^{k_n} &\geq 0. \end{aligned} \quad (\text{A.8})$$

By continuity of $A(\cdot)$, $b(\cdot)$, and $c(\cdot)$, taking the limit as $n \rightarrow \infty$ in (A.8), we get:

$$A(x)y = b(x), \quad wA(x) \geq c(x), \quad wb(x) = c(x)y, \quad y \geq 0.$$

These are precisely the optimality conditions that imply that $y \in \mathcal{L}(x)$, or equivalently that $(x, y) \in \mathcal{GL}$. \square

Condition (A.7) requires that whenever the primal linear program has an optimal solution, at least one optimal solution of the dual belongs to W . The set of optimal solutions of the dual is not necessarily a subset of W , but its intersection with W must be nonempty. The following example shows that if this sufficient condition is not satisfied, \mathcal{L} might fail to be closed.

Example A.4. Consider the following linear program, with parameter $x \geq 0$:

$$\max y_1 + y_2$$

subject to:

$$x(y_1 + y_2) \leq 0, \quad 2y_1 + y_2 \leq 1, \quad y_1, y_2 \geq 0.$$

The dual linear program is:

$$\min w_2$$

subject to:

$$xw_1 + 2w_2 \geq 1, \quad xw_1 + w_2 \geq 1, \quad w_1, w_2 \geq 0.$$

A little thought shows that:

$$\mathcal{L}(x) = \begin{cases} \{(0, 0)\}, & x > 0 \\ \{(0, 1)\}, & x = 0 \end{cases} \quad (\text{A.9})$$

and

$$\mathcal{D}(x) = \begin{cases} \{(\alpha, 0) : \alpha \geq 1/x\}, & x > 0 \\ \{(\alpha, 1) : \alpha \geq 0\}, & x = 0. \end{cases} \quad (\text{A.10})$$

From (A.9), one can see that \mathcal{L} does not have a closed graph (take $x^n \rightarrow 0$). On the other hand, for every compact subset W of \mathbb{R}^2 , there exists some $x \geq 0$, sufficiently small, such that $\mathcal{D}(x) \cap W = \emptyset$, as can be seen from (A.10). This situation is depicted in Figure 3.

*Remark A.5.*¹¹ The problem of continuity in linear programming is a special case of the problem addressed in the ‘‘maximum theorem’’ [Berge 1963], which

¹¹ The discussion in this remark requires familiarity with topological notions, such as upper and lower semi-continuity of correspondences [Border 1985; Takayama 1985]. The definitions of which will not be presented here. For the purposes of this discussion, suffice it to say that $R: X \rightarrow Y$ is: (i) upper semi-continuous if it is closed and Y is compact, (ii) lower semi-continuous, if and only if, for every $x \in X$, $x^n \rightarrow x$ and $y \in R(x)$ imply that there is a sequence $y^n \in R(x^n)$ with $y^n \rightarrow y$, and (iii) continuous if it is both upper and lower semi-continuous.




FIG. 3. Optimal solution set of the dual linear program.

states that the set of solutions to the problem of maximizing a continuous function over a compact set that varies continuously with some parameter is an upper semi-continuous correspondence. More precisely, let X and Y be subsets of two (finite dimensional) Euclidean spaces, $\mathcal{E}: X \rightarrow Y$ a compact-valued correspondence and $f: X \otimes Y \rightarrow \mathbb{R}$ a continuous function. Define $\Gamma: X \rightarrow Y$ by $\Gamma(x) = \arg \max_{y \in \mathcal{E}(x)} f(x, y)$, and $F: X \rightarrow \mathbb{R}$ by $F(x) = \max_{y \in \mathcal{E}(x)} f(x, y)$. The maximum theorem asserts that if \mathcal{E} is continuous at $x \in X$, then Γ is closed and upper semi-continuous at x , and F is continuous at x [Border 1985; Theorem 12.1].

Let us now investigate the relation of Theorem A.3 with the maximum theorem. Consider the continuously parametrized linear program (A.1)–(A.2) and define $\mathcal{E}(x) = \{y \in \mathbb{R}^m: A(x)y = b(x), y \geq 0\}$ and $f(x, y) = c(x)y$, for all $x \in X, y \in \mathbb{R}^m$. We concentrate on the case that the constraint set $\mathcal{E}(x)$ is a subset of a compact set $Y \subset \mathbb{R}^m$, since this is the structure of the linear program (3.12)–(3.18). Evidently, f is continuous in $X \otimes Y$ and $\mathcal{E}(x)$ is a compact set for all $x \in X$, as it is required by the maximum theorem. Furthermore, it is easy to see that $\mathcal{E}: X \rightarrow Y$ is closed and, since Y is compact, it is also upper semi-continuous. Nonetheless, \mathcal{E} is not lower semi-continuous in general, as can be seen by the example following Theorem A.3. In that example, we have:

$$\mathcal{E}(x) = \begin{cases} \{(0, 0)\}, & x > 0 \\ \{y_1, y_2 \geq 0: 2y_1 + y_2 \leq 1\}, & x = 0 \end{cases}$$

Taking a sequence $x^n > 0$ with $x^n \rightarrow 0$, and $y = (0, 1) \in \mathcal{E}(0)$, there is no sequence $y^n \in \mathcal{E}(x^n) = \{(0, 0)\}$ such that $y^n \rightarrow y$; thus, \mathcal{E} is not lower semi-continuous.

Theorem A.3 relaxes the requirement of continuity of the constraint set itself to continuity of $A(\cdot)$ and $b(\cdot)$, but it introduces the additional condition (A.7) on the solution set of the dual linear program. It is worth noting that this condition does not imply lower semi-continuity of the constraint set of the primal. Consider, for example, the following linear program with parameter $x \geq 0$:

$$\max y_2$$

subject to:

$$xy_1 + y_2 \leq 0, \quad 2y_1 + y_2 \leq 1, \quad y_1, y_2 \geq 0.$$

The dual linear program is:

$$\min w_2$$

subject to:

$$xw_1 + 2w_2 \geq 0, \quad w_1 + w_2 \geq 1, \quad w_1, w_2 \geq 0.$$

Then, $\mathcal{D}(x) = \{(\alpha, 0) : \alpha \geq 1\}$, for all $x \geq 0$, and \mathcal{D} satisfies the sufficient condition of Theorem A.3. However, $\mathcal{E}(x) = \{(0, 0)\}$ for all $x > 0$, while $\mathcal{E}(0) = \{(\alpha, 0) : 0 \leq \alpha \leq 1/2\}$; therefore, \mathcal{E} is not lower semi-continuous at $x = 0$.

Combining the results of Lemma A.2 and Theorem A.3, we arrive at the following:

COROLLARY A.6. *Consider a continuously parametrized linear program as in Theorem A.3. If $\mathcal{L}(x)$ is a singleton for all $x \in X$ and in addition to the sufficient condition (A.7) of Theorem A.3, there exists a compact set $Y \subset \mathbb{R}^m$, such that $\mathcal{L}(x) \subset Y$, for all $x \in X$, then the function $\mathcal{L}: X \rightarrow Y$ is continuous.*

Appendix B. The Two-User Case: Examples

In this appendix we use a simple flow control model to construct three examples, the first of which demonstrates that the joint strategy space is nonconvex, the second that if the sufficient condition of Theorem 5.3 is relaxed, the best reply function of user k might be discontinuous, and the third that a flow control game might, in general, have multiple Nash equilibria.

Consider the flow control model of Figure 1 with two users and the network consisting of a single $M/1$ -FCFS queue with service rate μ . Let (l_1, l_2) be the state of the network queue, where l_k is the number of class k packets in it. Using eqs. (2.1) and (2.2), the equilibrium distribution of the network queue is:

$$p_{l_1, l_2} = p_{00} \binom{l_1 + l_2}{l_1} \prod_{i_1=0}^{l_1-1} \frac{\lambda_{i_1}^1}{\mu} \prod_{i_2=0}^{l_2-1} \frac{\lambda_{i_2}^2}{\mu}, \quad 0 \leq l_1 \leq N_1, \quad 0 \leq l_2 \leq N_2. \quad (\text{B.1})$$

The conditional estimate of the rate at which class 1 packets are served in the network queue, given that there are l_1 such packets in it, is:

$$\begin{aligned} v_{l_1}^1 &= \sum_{l_2=0}^{N_2} \mu \frac{l_1}{l_1 + l_2} p_{l_1, l_2} \\ &= \mu \sum_{l_2=0}^{N_2} \binom{l_1 + l_2 - 1}{l_1 - 1} \prod_{i_2=0}^{l_2-1} \frac{\lambda_{i_2}^2}{\mu} \left[\sum_{l_2=0}^{N_2} \binom{l_1 + l_2}{l_1} \prod_{i_2=0}^{l_2-1} \frac{\lambda_{i_2}^2}{\mu} \right]^{-1}, \quad (\text{B.2}) \end{aligned}$$

The conditional estimates $\nu_{l_2}^2$ for user 2 are given by a similar expression.

B.1. The Joint Strategy Space is Nonconvex

Let us now provide an example that shows that the joint strategy space defined as the subset of Λ in which the time-delay constraints of all users are satisfied, that is, $\Lambda' = \{\lambda \in \Lambda: E\tau^k \leq T^k, k = 1, \dots, K\}$ is not convex. For the two-user model under consideration, the average throughput and average queue length for user 1 are respectively:

$$E\gamma^1 = \sum_{l_1=0}^{N_1-1} \lambda_{l_1} \sum_{l_2=0}^{N_2} p_{l_1 l_2}, \quad EQ^1 = \sum_{l_1=1}^{N_1} l_1 \sum_{l_2=0}^{N_2} p_{l_1 l_2}, \quad (\text{B.3})$$

Using Little's formula, the time-delay constraint of user 1 can be written as $EQ^1 - T^1 E\gamma^1 \leq 0$.

Let $N_1 = 2$ and $N_2 = 1$. Using eqs. (B.1) and (B.3), the time-delay constraint of user 1 is given by:

$$\left(\frac{1}{\mu} - T^1\right) + \lambda_1^1 \left(\frac{2}{\mu^2} - \frac{T^1}{\mu}\right) + \lambda_0^2 \left(\frac{2}{\mu^2} - \frac{T^1}{\mu}\right) + \lambda_1^1 \lambda_0^2 \left(\frac{6}{\mu^3} - \frac{2T^1}{\mu^2}\right) \leq 0. \quad (\text{B.4})$$

Similarly, the time-delay constraint for user 2 is:

$$\left(\frac{1}{\mu} - T^2\right) + \lambda_0^1 \left(\frac{2}{\mu^2} - \frac{T^2}{\mu}\right) + \lambda_0^1 \lambda_1^1 \left(\frac{3}{\mu^3} - \frac{T^2}{\mu^2}\right) \leq 0. \quad (\text{B.5})$$

Now, let $\mu = 1$, $c^1 = c^2 = 3$, $T^1 = 1.5$ and $T^2 = 1.6$. The joint strategy space Λ' defined by $0 \leq \lambda_{l_k}^k \leq c^k$, $0 \leq l_k < N_k$, $k = 1, 2$, and the time-delay constraints (B.4) and (B.5) is depicted in Figure 4 and can be easily seen not to be convex.

B.2. Discontinuity of the Best Reply Function

Let us now construct an example that demonstrates that the best reply function of a user might be discontinuous if its time-delay bound does not satisfy the sufficient condition (5.8) of Theorem 5.3.

Consider the two-user model and restrict Λ^k ($k = 1, 2$) to strategies that have the window structure described by eq. (3.20) with the "intermediate" point, if any, at the end of the window, that is, strategies of the form:

$$\lambda_{l_k}^k = \begin{cases} 0, & L_k \leq l_k < N_k \\ \lambda_{L_k-1}^k \in (0, c^k], & l_k = L_k - 1 \\ c^k, & 0 \leq l_k < L_k - 1 \end{cases}, \quad (\text{B.6})$$

where L_k is some integer between 1 and N_k .

Suppose that user 2 employs a strategy $\lambda^2 \in \Lambda^2$. Then, $\nu_{l_1}^1$ is concave increasing in l_1 [Hsiao and Lazar 1984]. In that case, according to Lazar [1983], the best reply of user 1 will have the structure described by eq. (B.6). Therefore, if one of the users chooses a strategy as in eq. (B.6), all subsequent




FIG. 4. The joint strategy space Λ' is not convex.

best replies of both users will be of the same type. This observation allows us to concentrate only on strategies $\lambda^k \in \Lambda^k$. However, in doing so, we do not consider all the admissible strategies of the users and we partially observe the behavior of the flow control game.

Due to its special structure, any strategy $\lambda^k \in \Lambda^k$ can be represented by a single scalar, instead of a vector. Specifically, let us define $s^k = \sum_{i_k=0}^{N_k-1} \lambda_{i_k}^k$. Then, any $s^k \in [0, N_k c^k]$ corresponds to a unique strategy $\lambda^k \in \Lambda^k$. This observation allows us to represent the best replies of the users in two-dimensional plots.

Using standard differentiation techniques, it is a matter of simple algebra to verify that $v_{l_1}^1$, given in eq. (B.2), is strictly decreasing in $\lambda_{l_2}^2$, $0 \leq l_2 < N_2$. Therefore, $v_{l_1}^1$ attains its minimum value when $\lambda_{l_2}^2 = c^2$, for all l_2 . Now, let $N_1 = N_2 = 5$, $\mu = 1$, $c^1 = c^2 = 2$ and $T^1 = 5$. From eq. (B.2) with $\lambda_{l_2}^2 = c^2$, we have $1/y_1^1 \simeq 5.092 > T^1$, that is, the sufficient condition (5.8) for continuity of R^1 is not satisfied. The best reply function of user 1 versus s^2 is shown in Figure 5. From that figure, it is evident that the best reply function of user 1 is not continuous. In particular, at $s^2 = 9.625$ the user's best reply drops from 2 to 0; note that this indicates that R^1 is not continuous in λ_4^2 , when $\lambda_0^2 = \dots = \lambda_3^2 = 2$. Using the same example, we have observed that the best reply function is indeed discontinuous whenever T^1 does not satisfy (5.8); as T^1 is increased toward 5.092, the discontinuity point moves toward $s^2 = 10$.

B.3. Existence of Multiple Nash Equilibria

We now proceed with an example that demonstrates that the flow control game might have multiple Nash equilibria.

Consider the two-user model of the previous example. As in that example, we concentrate on strategies λ^k of the type described by eq. (B.6), that can be represented by a single scalar $s^k \in [0, N_k c^k]$. If we plot the best reply curves of both users in the same $s^1 - s^2$ coordinate system, it is easy to see that the points where these two curves intersect are Nash equilibria of the flow control game. Note, however, that by using this method we can determine only the equilibria that have the structure described by eq. (B.6).

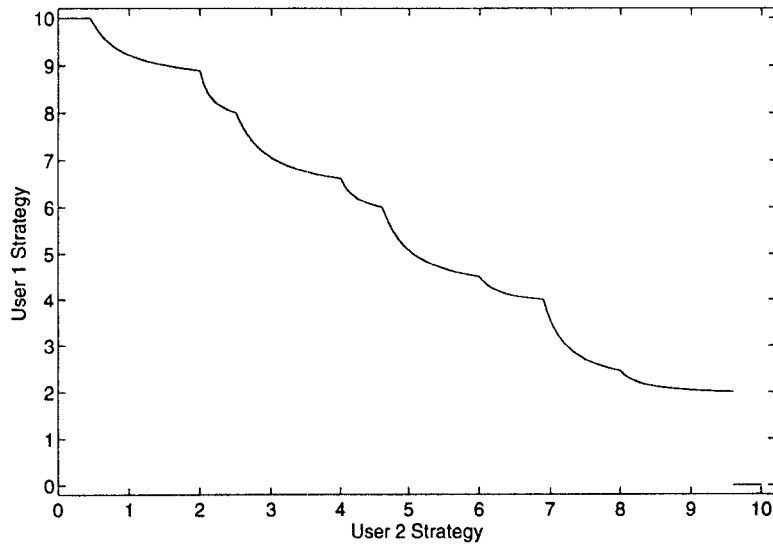


FIG. 5. The best reply function of user 1 is discontinuous.

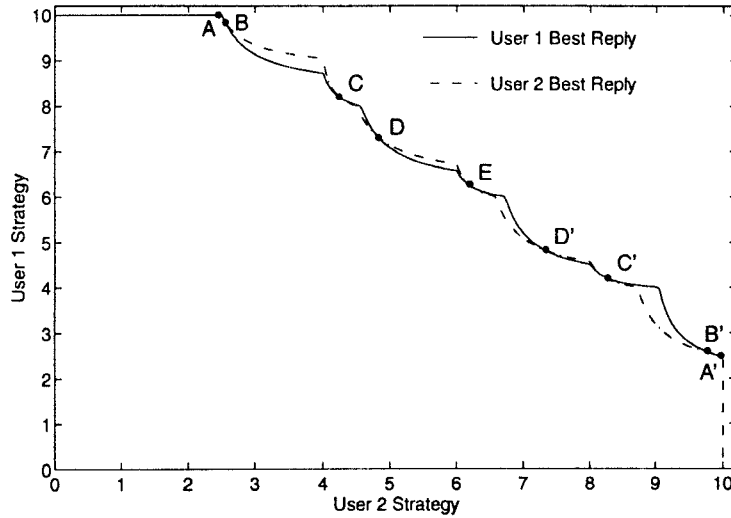


FIG. 6. Nash equilibria of the flow control game.

Now, let $N_1 = N_2 = 5$, $\mu = 1$, $c^1 = c^2 = 2$ and $T^1 = T^2 = 6$. As in the previous example, it can be seen that $1/\rho^k = 5.092 < T^k$, $k = 1, 2$, that is, the sufficient condition of Theorem 5.4 is satisfied. The best reply curves of the users are depicted in Figure 6. In that figure, one can identify nine Nash equilibria. Note that there exists an equilibrium (E) for which $s^1 = s^2$, while the remaining equilibria can be partitioned into symmetric pairs (A and A' , B and B' , and so forth), something that should be expected, because of the symmetry inherent in the model.

Using the two-user model, we have also observed the existence of multiple Nash equilibria in flow control games with users that are not symmetrical.

ACKNOWLEDGMENTS. The authors would like to thank the referees for their insightful comments and suggestions. Exploring the connection between the continuity of parametrized linear programs and the maximum theorem has been suggested by one of the referees, who provided also a pointer to the following references: [Berge 1963], [Gale 1960], and [Takayama 1985].

REFERENCES

- BAZARAA, M. S., JARVIS, J. J., AND SHERALI, H. D. 1990. *Linear Programming and Network Flows*, 2nd ed. Wiley, New York.
- BERGE, C. 1963. *Topological Spaces*. Macmillan, New York.
- BORDER, K. C. 1985. *Fixed-Point Theorems with Applications to Economics and Game Theory*. Cambridge University Press, New York.
- BOVOPOULOS, A. D., AND LAZAR, A. A. 1988. Asynchronous iterative algorithms for optimal load balancing. In *Proceedings of the 22nd Annual Conference on Information Sciences and Systems* (Princeton, N.J., Mar.). Princeton Univ., Princeton, N.J., pp. 1051–1057.
- DASGUPTA, P., AND MASKIN, E. 1986. The existence of equilibrium in discontinuous economic games. I: Theory. *Rev. Econ. Stud.* 53, 1–26.
- DEMERS, A., KESHAV, S., AND SHENKER, S. 1989. Analysis and simulation of a fair queueing algorithm. In *Proceedings of the ACM SIGCOM'89* (Austin, Tex., Sept. 19–22). ACM, New York, pp. 1–12.
- DOULIGERIS, C., AND MAZUMDAR, R. 1988. User optimal flow control in an integrated environment. In *Proceedings of the Indo-US Workshop on Signals and Systems* (Bangalore, India, Jan.).
- DOULIGERIS, C., AND MAZUMDAR, R. 1992. A game-theoretic approach to flow control in an integrated environment. *J. Franklin Inst.* 329, 3 (Mar.), 383–402.
- FUDENBERG, D., AND TIROLE, J. 1992. *Game Theory*. MIT Press, Cambridge, Mass.
- GALE, D. 1960. *The Theory of Linear Economic Models*. McGraw-Hill, New York.
- HSIAO, M.-T. T., AND LAZAR, A. A. 1984. Bottleneck modeling and decentralized optimal flow control—I. Global objective. In *Proceedings of the 18th Conference on Information Sciences and Systems* (Princeton, N.J., Mar.). Princeton Univ., Princeton, N.J., pp. 169–173.
- HSIAO, M.-T. T., AND LAZAR, A. A. 1989. An extension to Norton's equivalent. *Queueing Syst. Theory Appl.* 5, 401–412.
- HSIAO, M.-T. T., AND LAZAR, A. A. 1991. Optimal decentralized flow control of Markovian queueing networks with multiple controllers. *Perf. Eval.* 13, 3, 181–204.
- KELLY, F. P. 1979. *Reversibility and Stochastic Networks*. Wiley, New York.
- KORILIS, Y. A., AND LAZAR, A. A. 1992. Why is flow control hard: Optimality, fairness, partial and delayed information. CTR Technical Report 332-93-11, Center for Telecommunications Research, Columbia Univ., New York.
- LAZAR, A. A. 1983. Optimal control of a class of queueing networks in equilibrium. *IEEE Trans. Automatic Control AC-28*, 11 (Nov.), 1001–1007.
- LUENBERGER, D. G. 1984. *Linear and Nonlinear Programming*, 2nd ed. Addison-Wesley, Reading, Mass.
- MYERSON, R. B. 1991. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, Mass.
- NASH, J. 1951. Noncooperative games. *Ann. Math.* 54, 2 (Sept.), 286–295.
- ORDA, A., ROM, R., AND SHIMKIN, N. 1993. Competitive routing in multiuser communication networks. *IEEE/ACM Trans. Netw.* 1, 5 (Oct.), 510–521.
- PAREKH, A. K., AND GALLAGER, R. G. 1993. A generalized processor sharing approach to flow control in integrated services networks: The single-node case. *IEEE/ACM Trans. Netw.* 1, 3 (June), 344–357.
- ROSEN, J. B. 1965. Existence and uniqueness of equilibrium points for concave n -person games. *Econometrica* 33, 3 (July), 520–534.
- SCHWARTZ, M. 1987. *Telecommunication Networks: Protocols, Modeling and Analysis*. Addison-Wesley, Reading, Mass.
- TAKAYAMA, A. 1985. *Mathematical Economics*, 2nd ed. Cambridge University Press, New York.

WALRAND, J. 1988. *Introduction to Queueing Networks*. Prentice-Hall, New York.

ZHANG, Z., AND DOULIGERIS, C. 1992. Convergence of synchronous and asynchronous greedy algorithms in a multiclass telecommunications environment. *IEEE Trans. Commun.* 40, 8 (Aug.), 1277–1281.

RECEIVED JUNE 1993; REVISED NOVEMBER 1994; ACCEPTED DECEMBER 1994