Incentive Compatible Pricing Strategies for QoS Routing

Yannis A. Korilis
Dept. of Systems Engineering
University of Pennsylvania
Philadelphia, PA 19104-6315, USA
korilis@seas.upenn.edu

Ariel Orda
Dept. of Electrical Engineering
Technion
Haifa 32000, Israel
ariel@ee.technion.ac.il

Corresponding author:
Ariel Orda
Department of Electrical Engineering
Technion
Haifa 32000, Israel
email: ariel@ee.technion.ac.il

Abstract

QoS routing mechanisms allow users identify paths that can accommodate their performance requirements and reserve the necessary resources. An important problem is how to conduct such resource allocation efficiently, not only from the single-connection, but also from the *network* point of view. We propose the use of pricing mechanisms as a means to regulate the users' decisions in a networkwide efficient manner. Focusing on QoS architectures that employ rate-based schedulers, we formulate a congestion-based pricing scheme. We establish the structure of the corresponding user-optimal response, i.e., a path selection algorithm that satisfies the user's requirements at minimal cost. We show that the underlying non-cooperative game among users has a unique equilibrium, for any particular choice of price functions. Then, we establish the existence of *incentive compatible* price functions, which drive the network into an equilibrium point that coincides with the optimum of a social function. Specifically, these price functions are the derivatives of the social function. We then extend our results to scenarios in which users can identify only sub-optimal paths, as is often the case with multi-constrained path optimization.

Keywords: QoS Routing, Networking Games, Pricing, Rate-Based Schedulers.

Broadband integrated services networks are expected to support multiple and diverse applications, with various quality of service (QoS) requirements. Accordingly, a key issue in the design of broadband architectures is how to provide resources in order to meet the requirements of each connection, and, moreover, how to meet that goal in a networkwide efficient manner. The establishment of efficient QoS routing schemes is one of the major building blocks in such architectures. QoS routing has been the subject of several studies and proposals (see, e.g., Apostolopoulos et al. (1999), Crawley et al. (1998), Ma and Steenkiste (1997), Orda (1999), The PNNI Specification (1996), Pornavalai et al. (1997), Wang and Crowcroft (1996), and references therein).

One of the major problems in the establishment of a connection with QoS guarantees arises from the need to map end-to-end requirements, such as delay and/or jitter, onto local (nodal) requirements, which would indicate how to reserve resources along the route. The ability to derive such a mapping depends to a large extent on the scheduling policy and service discipline employed at the nodes. Such disciplines are characterized by bounds on the maximal delay that any node can incur, and hence a corresponding bound on the end-to-end delay can be derived. This way, the routing problem can be formulated as identifying the path that has the best performance according to that bound and with respect to the QoS requirement. Several studies have proposed schedulers that map delay guarantees into rate requirements and have each node advertise its residual rate (Georgiadis et al. (1996), Parekh and Gallager (1994), Zhang (1995)). The Guaranteed Service Class proposed for the Internet (Shenker et al. (1997)) is based on such rate-based principles. Based on the properties of such schedulers, several studies (Guerin and Orda (1999), Guerin et al. (1997), Ma and Steenkiste (1997), Orda (1999), Pornavalai (1997)) have analyzed algorithms for computing paths that satisfy end-to-end delay bounds. In particular, it has been shown in Guerin and Orda (1999) that for a given connection with an end-to-end delay constraint, the existence and identity of a feasible path can be obtained through up to M executions of a standard shortest path algorithm, where M is the number of network links. In Ma and Steenkiste (1997) and Pornavalai (1997), it was shown that, through this scheme, one can accommodate additional connection requirements such as jitter.

An important problem that has not been sufficiently addressed in the literature on QoS routing is that of efficient allocation of resources, namely "rates" or "bandwidth", not only from the single-connection, but also from the *network* point of view. In particular, while each connection[1] can choose the path between source and destination along with the corresponding bandwidth reservations, the network provider/manager typically aims at an allocation of resources that is deemed efficient with respect to the overall network performance. The underlying assumption in previous studies on QoS routing is that efficient usage of network resources can be enforced through appropriate choice of pricing strategies. What constitutes "efficient" resource utilization and how it can be achieved through pricing mechanisms

[1]The terms "connection" and "user" are used interchangeably.

are still open problems. Both these open problems are addressed in the present study.

Pricing as an allocation mechanism that makes decentralized decisions compatible with overall efficiency has been studied in the context of queueing systems, e.g., Mendelson and Whang (1990). In computer networks, pricing has been receiving increasing attention from both the research and the corporate world, mostly due to the explosive growth of the Internet, which is evolving from a heavily subsidized network to a commercial enterprise, e.g., Edell (1995), Jordan and Jiang (1995), Kelly (1993), Korilis et al. (1998), Low and Varaiya (1993), MacKie-Mason and Varian (1995), Orda and Shimkin (2000), Sairamesh et al. (1995). The research community has taken a normative approach, proposing usage-based pricing mechanisms that will motivate the users to adopt a social behavior, e.g., by regulating their traffic, or by requesting lower grade of service. Network providers, on the other hand, will benefit from using network-efficient pricing schemes, but at the same time are interested in mechanisms that first generate profit and second are more appealing to the end users than the mechanisms of their competitors.

The interface between the performance- and the market-oriented approach to pricing, is one of the main factors that will define the future evolution of public networks – networks where access is not restricted to members of an enterprise – such as the Internet. In an attempt to delineate this interface, various debates have arisen within both communities. There seems to be a consensus, though, that some type of usage-constraining pricing is necessary mainly due to congestion considerations. In the present study we do not attempt to enter these debates. Rather, within the framework of the normative approach, we demonstrate how pricing strategies can be used to drive the network to an operating point that is deemed efficient with respect to the overall network performance.

More specifically, we consider a general network with nodal schedulers that belongs to the rate-based class as described in the specification of the Guaranteed Service for IP. Each connection is characterized by its source-destination nodes, maximal packet size, maximal burst, and an upper bound on the end-to-end delay. The properties of rate-based schedulers allow the derivation of an upper bound on the end-to-end delay of a connection when it is routed over a given path at a given reserved rate.

There is a cost associated with reserving a unit of rate over a link, which is the price of the link. Focusing on congestion pricing, we assume that the price of a link is a function of the aggregate rate reserved at the link. Each connection is established so as to minimize the total usage cost while satisfying its end-to-end delay constraint. The interaction among the various connections that decide independently on their individual routing strategies can be modelled as a *game* (Fudenberg and Tirole (1992), Myerson (1991)). Any operating point of the network is a *Nash equilibrium* of that game, that is, a collection of routing strategies from which no user has an incentive to deviate unilaterally.

Link price functions are determined by the network provider/manager. The goal of the manager is to drive the users to a Nash equilibrium that is efficient from the network's point of view. More specifically, we assume that efficiency is defined as minimizing a global (social) cost function that quantifies the overall network performance and is the sum of link cost functions. The manager seeks a pricing strategy that enforces a unique Nash equilibrium that minimizes this social cost function. Any such pricing strategy is called *incentive compatible*.

We investigate the structure of the QoS-routing game and show that, for any given set of link price functions (conforming to a set of general assumptions), it has a unique Nash equilibrium. Moreover, we establish a set of necessary and sufficient conditions for a feasible (link) flow vector to be the equilibrium of the game. Having established these results, we turn our attention to the problem of incentive compatible pricing strategies. We show that if the network manager imposes link price functions that are equal to the derivatives of the link cost functions, the unique equilibrium of the QoS-routing game coincides with the network optimum. We note that this type of price functions have been known to enforce the network optimum when the users implement a much simpler class of optimal routing strategies, such as in transportation networks (Dafermos and Sparrow (1971)).

In the sequel, we turn our attention to connections that conduct *multi-constrained* path optimization. A typical setting is to identify a path that minimizes some target function, e.g., administrative costs, while observing two or more constraints, such as end-to-end delay and jitter. For this setting we show that the previous results about the routing game and the incentive compatible prices still apply. These results are based on the assumption that the users are able to determine optimal routes that provide both delay and jitter guarantees. However, such multi-constrained path optimization problems are, in general, NP-complete (Garey and Johnson (1979)), therefore optimal routing solutions are prohibitively complex. On the other hand, there are efficient approximation schemes which provide ϵ-optimal solutions within polynomial time complexity (see, e.g., Hassin (1992)).

This means that users can be expected to make not self-optimizing but only sub-optimal decisions. This situation presents a harder challenge for network management, as the response of users to management schemes becomes unpredictable. An important question is, then, whether there is still a pricing scheme that drives the network to an efficient operating point. We indicate that the answer is affirmative. Moreover, we show that the required prices are exactly those that correspond to the standard scenario of self-optimizing users.

The rest of the paper is structured as follows. In Section 1 we present the QoS-routing model and formulate the problem of incentive compatible pricing. Focusing on end-to-end delay constraints, in Section 2 we investigate the structure of the QoS-routing game and study the problem of incentive compatible pricing. The multi-constrained path optimization case is considered in Section 3. Conclusions are presented in Section 4. Due to space limits, proofs and some details are not included in this version and can be found in Korilis and Orda (2002).

1 Model and Problem Formulation

We consider a network $\mathcal{G}(\mathcal{V}, \mathcal{L})$, where $\mathcal{V}$ is the set of nodes and $\mathcal{L} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of links, and let $N = |\mathcal{V}|$ and $M = |\mathcal{L}|$, where $|\cdot|$ denotes the cardinality of a set. We denote by H the maximal possible number of hops (links) in a path. For any link $l = (u, v) \in \mathcal{L}$, define $S(l) = u$ and $T(l) = v$. Considering a node $u \in \mathcal{V}$, let $In(v) = \{l : T(l) = u\}$ denote the set of its ingoing links, and $Out(v) = \{l : S(l) = u\}$ the set of its outgoing links. Each link $l \in L$ is characterized by the following parameters:

- A *maximal rate* (capacity) R_l, which the link can offer to a new connection. When a new connection with a rate $r < R_l$ is established through link l, the value of R_l becomes $R_l - r$.[2]
- A *constant delay* d_l, related to the link's speed, propagation delay and maximal packet size sent over the link. We assume that d_l takes integer values.

Connections belong to a set $\mathcal{I} = \{1, 2, \ldots, I\}$ of "types." A connection (of type) $i \in \mathcal{I}$ is characterized by the following parameters:

- a pair of source and destination nodes, taken from a set $\mathcal{S} \subseteq \mathcal{V} \times \mathcal{V}$,
- a bias σ^i, related to the connection's maximal burst,
- a maximal packet size c^i,
- an end-to-end delay QoS requirement D^i, which, without loss of generality, is assumed to take integer values,
- a bandwidth QoS requirement b^i.

The above parameters characterize the traffic profile of the connection. They form part of the traffic contract between the connection and the network, and affect the actual QoS that the connection obtains, as explained in the following.

A connection should be routed through some path $\mathbf{p}$ between the corresponding source and destination nodes. We shall denote by $n(\mathbf{p})$ the number of hops of a path $\mathbf{p}$, and by $r(\mathbf{p})$ its maximal available rate, that is, $r(\mathbf{p}) = \min_{l \in \mathbf{p}} R_l$. In order to offer QoS guarantees, a scheduler is employed at each network link. We assume that the scheduling policy in the network belongs to the *rate-based* class (Georgiadis et al. (1996), Parekh and Gallager (1994), Zhang (1995)), as in the specification of the Guaranteed Service for IP (Shenker et al. (1997)). Such schedulers divide the available link rate among the active connections in a manner that is proportional to their guaranteed rates. This way, it is possible to guarantee a minimum rate and a maximum delay to each connection at each link independently of the behavior of other connections. Obviously, the connection's bandwidth requirement should not exceed its guaranteed rate. Moreover, in order to bound the worst case delay, the connection should bound its burstiness, i.e., momentary deviation from the (average) rate. In addition, the worst case delay is affected by packet size. This is because we consider nonpreemptive schedulers, i.e., that do not stop the transmission of a packet even if it belongs to a connection that has momentarily exceeded its allocated link rate; hence, larger packet sizes increase the potential deviation from the behavior of an

[2]If the maximal available rate is a nodal property, then we associate it with all its outgoing links.

"ideal" scheduler. Clearly, the worst case delay increases with the connection's burstiness and maximal packet size and with the constant delays along its path, and decreases with the guaranteed rate. These observations are formally stated in the following (for more details see, e.g., Parekh and Gallager (1994)).

When a connection i is routed over a path $\mathbf{p}$ with a reserved rate $r \leq r(\mathbf{p})$, its end-to-end delay is upper bounded by:

$$D^i(\mathbf{p}, r) = \frac{\sigma^i + n(\mathbf{p})c^i}{r} + \sum_{l \in \mathbf{p}} d_l. \tag{1}$$

Let $D^i(\mathbf{p}) = D^i(\mathbf{p}, r(\mathbf{p}))$ denote the minimal possible value of $D^i(\mathbf{p}, r)$, which shall be referred to as *the guaranteed delay* of $\mathbf{p}$. A path $\mathbf{p}$ between the source and destination nodes of a connection of type i is said to be *feasible* for connection i if $D^i(\mathbf{p}) \leq D^i$ and $r(\mathbf{p}) \geq b^i$. Paths that cannot accommodate the bandwidth requirement of the connection can be eliminated, thus the bandwidth constraint can be treated as absent. Therefore, denoting by $\alpha^i(\mathbf{p}) = \frac{\sigma^i + n(\mathbf{p})c^i}{D^i - \sum_{l \in \mathbf{p}} d_l}$ the minimal rate that satisfies the delay constraint of connection i on path $\mathbf{p}$, the feasibility of a path can be defined as follows.

Definition 1.1 *A path* $\mathbf{p}$ *between the source and destination nodes of a connection of type* i *is said to be* feasible *for a connection (of type)* i*, if* $D^i > \sum_{l \in \mathbf{p}} d_l$ *and* $r(\mathbf{p}) \geq \alpha^i(\mathbf{p})$*. Let* $\mathcal{P}^i$ *denote the set of all feasible paths for that connection.*

We assume that the number of connections is very large and the rate required by each is very small compared to link capacities. Hence, a continuous approximation is appropriate, and the (normalized) number of connections can be assumed to take continuous values. Denote by $\gamma_u^i(j)$ the expected number of type i connections with source node u and destination node j. There is a *usage cost* w_l associated with reserving a unit of bandwidth on link $l \in \mathcal{L}$ that will be referred to as the *price* of the link. The total cost for reserving bandwidth r over path $\mathbf{p}$ is, then, $r \cdot \sum_{l \in \mathbf{p}} w_l$. Each user makes its routing decisions (choice of path(s) between source and destination and corresponding reserved rate) independently, according to its individual QoS requirements and cost considerations. More specifically, each connection is established so as to minimize its total cost, while satisfying its delay constraint.

The price for usage of a link is the same for all connections and depends only on the *aggregate* rate reserved on the link. More specifically, if f_l denotes the total rate reserved by all connections on link l, then the cost for reserving a unit of bandwidth (price) on the link is $w_l(f_l)$. f_l will be referred to as the "flow" on link l. Here, it is important to note the underlying dynamics of the network; users choose feasible paths according to link prices, which, in turn, are functions of the aggregate rates reserved on the links, thus depend on the choices made by the users. This dynamic behavior can be modeled as a noncooperative game. An *equilibrium* of the network, that is, a link flow distribution $\mathbf{f} = \{f_l : l \in \mathcal{L}\}$ where prices and aggregate rates remain unchanged, is a *Nash equilibrium* of the underlying QoS-routing game.

The set of network equilibria depends on the link price functions that are determined by the network manager and/or provider. Depending on the way the network is financed and the market structure, the manager/provider might choose prices according to various objectives that combine, in general, the aim to operate the network efficiently and, at the same time, generate revenue. in determining prices, the structure of the user community, as well as competition by other providers should also be taken into account. We do not attempt to address all these issues here. Rather, we aim at a rigorous investigation of pricing strategies that lead to efficient utilization of network resources.

Network efficiency is typically defined as achieving an operating point that minimizes a global (social) cost function $J(\mathbf{f})$ that quantifies the overall network performance: the higher $J(\mathbf{f})$, the lower the network performance under $\mathbf{f}$ is. We focus on social cost functions $J(\mathbf{f}) = \sum_{l \in \mathcal{L}} J_l(f_l)$ that are the sum of link cost functions, where the cost of link l depends only on the total flow f_l on the link and satisfies the following.

Assumptions G:

G1. $J_l : [0, R_l] \to [0, \infty]$ is continuous.

G2. J_l is increasing and strictly convex.

G3. Wherever finite, J_l is continuously differentiable.

It follows from standard convex programming results that, under Assumptions G, there exists a unique operating point that minimizes J, which shall be referred to as the *network optimum*. The goal of the network manager is to determine a pricing strategy (collection of price functions) $\{w_l(f_l) : l \in \mathcal{L}\}$ that provides incentives to the users to make choices that lead to a *unique* network equilibrium that coincides with the network optimum.[3] As mentioned, any such pricing strategy is called *incentive compatible*. Here, we concentrate on price functions $\{w_l(f_l)\}$ that satisfy the following.

Assumptions P:

P1. $w_l(f_l)$ is nonnegative.

P2. $w_l : [0, R_l] \to [0, \infty]$ is continuous.

P3. $w_l(f_l)$ is increasing.

Existence of incentive compatible pricing strategies is investigated in the following section. It is shown that an incentive compatible pricing strategy of the manager always exists and its structure is specified explicitly.

As already explained, each connection is established so as to minimize its total cost while satisfying its QoS constraint. Till now, we have only considered connections with end-to-end delay constraints. In general, a connection might have multiple QoS constraints, for example, end-to-end delay and jitter. The underlying QoS-routing game and the related pricing problem, under multiple QoS constraints, are

[3] Depending on the pricing scheme, the network might have multiple equilibria, a unique equilibrium, or no equilibrium at all. Evidently, the manager would avoid pricing schemes that prevent the network from reaching an operating point. Similarly, schemes that lead to multiple equilibria are, in general, undesirable, since the operating point of the network depends on its initial state and there is no guarantee that it is the equilibrium that corresponds to the desired point.

investigated in Section 3.

Throughout the paper, we assume that the network can accommodate the total offered load, i.e., that there is a *stable* network operating point $\mathbf{f}$, in the sense that $f_l < R_l$ for all links $l \in \mathcal{L}$.

2 Incentive Compatible Pricing Strategies

In this section we investigate the existence of incentive compatible pricing strategies, that is, price functions $\{w_l(f_l) : l \in \mathcal{L}\}$ that induce a unique network equilibrium that coincides with the network optimum. To do so, we first need to analyze the structure of the QoS-routing game for given link price functions.

Consider the framework of a communication network and connections with QoS requirements as defined in the previous section. Let $\alpha^{n,d,i} = \frac{\sigma^i + nc^i}{D^i - d}$, where $1 \leq n \leq H$, $i \in \mathcal{I}$, and $1 \leq d < D^i$, i.e., $\alpha^{n,d,i}$ is the minimal rate required by a type-i connection established over a path of n hops and total constant delay d.

Let us start by studying the routing problem faced by a user. Given a set of nonnegative link prices $\{w_l\}$, the routing problem of a type i user can be defined as follows.

User Problem (U):

Find a path $\mathbf{p}^* \in \mathcal{P}^i$ *that achieves the minimum:*

$$\min_{\mathbf{p} \in \mathcal{P}^i} \frac{\sigma^i + n(\mathbf{p})c^i}{D^i - \sum_{l \in \mathbf{p}} d_l} \sum_{l \in \mathbf{p}} w_l.$$

The above problem belongs to a class of combinatorial path optimization problems with *rational objective functions.* That class has been investigated in Guerin and Orda (1999), where a polynomial solution was established. The complexity of that solution, when applied to our case, is $O(N^2 M \log(N \frac{\max_{l \in \mathcal{L}} w_l}{\min_{l \in \mathcal{L}} w_l} D^i))$. As the solution is of reasonable complexity, users can be expected to identify their optimal paths, given a set of link prices. To summarize:

Proposition 2.1 Problem U *is solvable through a polynomial,* $O(N^2 M(\log(N \frac{\max_{l \in \mathcal{L}} w_l}{\min_{l \in \mathcal{L}} w_l} D^i))$ *algorithm.*

We now proceed to characterize the structure of the QoS-routing game. Considering the network *at equilibrium* (assuming that one exists), denote by $\gamma_u^{n,d,i}(j)$ the expected number of type-i connections from source node u to destination node j, which are established over paths with n hops and total constant delay d. Similarly, let $f_l^{n,d,i}(j)$ be the aggregate flow over a link l, corresponding to type-i connections destined to node j, which are routed over paths, for which link l is the n-th hop away from the destination and the total constant delay on the remainder of the path is d. By a slight abuse of notation, define the link flow vector $\mathbf{f} = \{f_l^{n,d,i}(j) : 1 \leq n < H, i \in \mathcal{I}, 0 \leq d \leq D^i, l \in \mathcal{L}, j \in \mathcal{V}\}$.

Similarly, define $\gamma = \{\gamma_u^{n,d,i}(j) : 1 \le n < H, i \in \mathcal{I}, 0 \le d \le D^i, u, j \in \mathcal{V}, u \neq j\}$. Flow vectors $\mathbf{f}$ and γ must satisfy the following feasibility constraints.

Constraints $\mathcal{F}$:

1. $\forall\, i, \forall\, j, u \in \mathcal{V}\ u \neq j, \forall\, 1 \le n \le H, \forall\, 0 \le d < D^i$:

$$\sum_{v \in \mathcal{V}} f_{uv}^{n,d,i}(j) = \sum_{v \in \mathcal{V}} f_{vu}^{n+1,d+d_{vu},i}(j) + \gamma_u^{n,d,i}(j)\alpha^{n,d,i};$$

2. $\forall\, i, \forall\, j, u \in \mathcal{V}\ u \neq j$: $\sum_{n=1}^{H} \sum_{d=0}^{D^i-1} \gamma_u^{n,d,i}(j) = \gamma_u^i(j)$;

3. $\forall\, i, \forall\, j \in \mathcal{V}, \forall\, (u,v) \in \mathcal{L}, \forall\, 1 \le n \le H, \forall\, 0 \le d < D^i$: $f_{uv}^{n,d,i}(j) \ge 0$;

4. $\forall\, j, u \in \mathcal{V}\ u \neq j, \forall\, 1 \le n \le H, \forall\, 0 \le d < D^i$: $\gamma_u^{n,d,i}(j) \ge 0$.

The first constraint is a flow conservation constraint: it states that $\sum_{v \in \mathcal{V}} f_{uv}^{n,d,i}(j)$, i.e., the aggregate flow of a connection of type i to the destination j for which a node u is at the starting-point of the n-th hop and with d remaining units of constant delay, is equal to the sum of $\sum_{v \in \mathcal{V}} f_{vu}^{n+1,d+d_{vu},i}(j)$, i.e., the aggregate flow to j for which node u is at the end-point of the $(n+1)$–st hop and with d remaining units of constant delay (after reaching node u), plus $\gamma_u^{n,d,i}(j)\alpha^{n,d,i}$, i.e., the aggregate rate of connections of type i established between u and j over paths that have n hops paths and d units of constant delay.[4] The second constraint simply states that each connection should be routed over some path. The third and forth are obvious nonnegativity constraints. Let $\mathcal{F}$ denote the set of feasible flow vectors $(\mathbf{f}, \gamma)$, i.e., those that obey the above four constraints.

Let $(\mathbf{f}, \gamma)$ be an equilibrium flow vector. Then, the total flow on any link $l \in \mathcal{L}$ is

$$f_l = \sum_{n=1}^{H} \sum_{i \in \mathcal{I}} \sum_{d=0}^{D^i-1} \sum_{j \in \mathcal{V}} f_l^{n,d,i}(j).$$

We observe that a path chosen by a user (i.e., that solves *Problem U*) has a minimal price among paths between the same source and destination with the same number of hops and constant delay. This implies that chosen paths have the following *local optimality* property: if $\mathbf{p_1}$ and $\mathbf{p_2}$ are two paths between a node u and a destination node j, each of n hops and having the same propagation delay d (that is, $\sum_{l \in \mathbf{p_1}} d_l = \sum_{l \in \mathbf{p_2}} d_l = d$), and if $\mathbf{p_1}$ is a subpath of a path chosen by some user whose destination is j, then $\sum_{l \in \mathbf{p_1}} w_l \le \sum_{l \in \mathbf{p_2}} w_l$. This observation leads to the following equilibrium conditions.

$$\forall i, \quad \forall u, j \in \mathcal{V} \quad, u \neq j, \forall v \in Out(u), \forall 1 \le n \le H, \forall 0 \le d < D^i :$$

$$\lambda_u^{n,d}(j) \begin{cases} = w_{uv}(f_{uv}) + \lambda_v^{n-1,d-d_{uv}}(j) & , \quad f_{uv}^{n,d,i}(j) > 0 \\ \le w_{uv}(f_{uv}) + \lambda_v^{n-1,d-d_{uv}}(j) & , \quad f_{uv}^{n,d,i}(j) = 0 \end{cases} . \tag{2}$$

[4] Indeed, the expected number of such connections is $\gamma_u^{n,d,i}(j)$, and each requires a reservation of $\alpha^{n,d,i}$ units of rate.

$$\forall i, \quad \forall u, j \in \mathcal{V} \quad, u \neq j, \forall 1 \leq n \leq H, \forall 0 \leq d < D^i :$$

(3)
$$\tilde{\lambda}_u^i(j) \begin{cases} = \lambda_u^{n,d}(j)\alpha^{n,d,i} & , \quad \gamma_u^{n,d,i}(j) > 0 \\ \leq \lambda_u^{n,d}(j)\alpha^{n,d,i} & , \quad \gamma_u^{n,d,i}(j) = 0 \end{cases} .$$

(4)
$$\lambda_j^{n,d}(j) \equiv \tilde{\lambda}_j^i(j) \equiv 0.$$

The equilibrium conditions can be explained as follows. $\lambda_u^{n,d}(j)$ is the minimal equilibrium price for reserving one unit of rate between a source u and a destination j, over a path with n hops and constant delay d; expression (2) simply states that, among paths with the same number of hops and constant delay, users will prefer those with minimal per-unit rate prices. $\tilde{\lambda}_u^i(j)$, in turn, is the minimal equilibrium cost for establishing a connection of type i between a source u and a destination j; expression (3) is just a restatement of the observation made above.

The structure of the QoS-routing game is given by the following theorem.

Theorem 2.1 *For any set of link price functions $\{w_l(f_l)\}$ that satisfy Assumptions P:*

1. *There exists a unique equilibrium flow vector $(\mathbf{f}, \gamma) \in \mathcal{F}$.*

2. *Conditions (2)–(4) are necessary and sufficient for equilibrium, that is: a flow vector $(\mathbf{f}, \gamma) \in \mathcal{F}$ is an equilibrium if and only if there exist finite numbers $\{\lambda_u^{n,d}(j) : 1 \leq n < H, 0 \leq d \leq \max_{i \in \mathcal{I}} D^i, u, j \in \mathcal{V}, u \neq j\}$ and $\{\tilde{\lambda}_u^i(j) : i \in \mathcal{I}, u, j \in \mathcal{V}, u \neq j\}$, such that (2)–(4) hold.*

We now turn our attention to the problem of the manager. The manager aims at selecting link price functions $\{w_l(f_l)\}$ that enforce the network optimum, i.e., a flow vector $(\mathbf{f}, \gamma)$ that minimizes the social function $J(\cdot)$ over the set of all feasible flow vectors. According to Theorem 2.1, each choice of a set of link price functions that comply with Assumptions P induces a unique network equilibrium. Therefore, the problem faced by the network manager can be stated as follows.

Manager Problem (M):
Find a set of link price functions $\{w_l(f_l)\}$ that satisfy Assumptions P and lead to a network equilibrium, which coincides with the network optimum, i.e., a flow vector $(\mathbf{f}, \gamma)$ that solves the following problem.
Minimize:

$$J(\mathbf{f}) = \sum_{l \in \mathcal{L}} J_l(f_l) = \sum_{l \in \mathcal{L}} J_l(\sum_i \sum_{n=1}^{H} \sum_{d=0}^{D^i - 1} \sum_{j \in \mathcal{V}} f_l^{n,d,i}(j))$$

subject to Constraints $\mathcal{F}$.

The following proposition characterizes the structure of the network optimum through a set of necessary and sufficient conditions.

Proposition 2.2 *Given a social function $J(\cdot)$ that conforms with Assumptions G, a flow vector $(\mathbf{f}, \gamma)$ is the network optimum if and only if there exist (Lagrange multipliers) $\{\lambda_u^{n,d}(j) : j, u \in \mathcal{V}, u \neq j, 1 \leq$*

$n \leq H, 0 \leq d < \max_{i \in \mathcal{I}} D^i\}$ and $\{\tilde{\lambda}_u(j) : j, u \in \mathcal{V}, u \neq j\}$, such that the following (Kuhn-Tucker) conditions are satisfied:

$$\forall i, \quad \forall u, \quad j \in \mathcal{V}, u \neq j, \forall v \in Out(u), \forall 1 \leq n \leq H, \forall 0 \leq d < D^i :$$

(5)
$$\lambda_u^{n,d}(j) \begin{cases} = J'_{uv}(f_{uv}) + \lambda_v^{n-1,d-d_{uv}}(j) & , \quad f_{uv}^{n,d,i}(j) > 0 \\ \leq J'_{uv}(f_{uv}) + \lambda_v^{n-1,d-d_{uv}}(j) & , \quad f_{uv}^{n,d,i}(j) = 0 \end{cases} .$$

$$\forall i, \quad \forall u, \quad j \in \mathcal{V}, u \neq j, \forall 1 \leq n \leq H, \forall 0 \leq d < D^i \quad \lambda_u^{n,d}(j) :$$

(6)
$$\tilde{\lambda}_u^i(j) \begin{cases} = \lambda_u^{n,d}(j) \alpha^{n,d,i} & , \quad \gamma_u^{n,d,i}(j) > 0 \\ \leq \lambda_u^{n,d}(j) \alpha^{n,d,i} & , \quad \gamma_u^{n,d,i}(j) = 0 \end{cases} .$$

(7)
$$\lambda_j^{n,d}(j) \equiv \tilde{\lambda}_j^i(j) \equiv 0.$$

We have thus established the main result of this section.

Theorem 2.2 *For link price functions*

(8)
$$w_l(f_l) \stackrel{\triangle}{=} J'_l(f_l), \quad l \in \mathcal{L},$$

the flow vector $(\mathbf{f}^*, \gamma^*)$ that minimizes the social cost function J over $\mathcal{F}$ is the unique equilibrium.

The implication of Theorem 2.2 is that, by setting link price functions that are equal to the derivatives of the social cost functions, the manager can drive a network to the required social optimum.

Remark 2.1 *It is important to note that, while the equilibrium conditions (2)–(4) are stated with a considerably large number of variables, namely $O(\max_{i \in \mathcal{I}} D^i \cdot H \cdot N \cdot M)$ Lagrange multipliers, the optimality problem faced by each user is of polynomial size, as established by Proposition 2.1. Similarly, the large number of variables with which the optimality conditions (5)–(7) are stated has no implication on the manager: as established by Theorem 2.2, the size of its strategy,* i.e., *optimal price vector, is just $O(M)$.*

3 Incentive Compatible Pricing Strategies for Multi-Constrained Path Optimization

We turn to consider the case where a connection needs to accommodate additional constraints, besides end-to-end delay. In order to concretize the discussion, we focus on jitter constraints, and denote by τ^i the *maximal end-to-end jitter constraint* of a type i connection.

As before, we assume that the scheduling policy in the network belongs to the "rate-based" class (Shenker et al. (1997)). Accordingly, when a connection i is routed over a path $\mathbf{p}$ with a reserved rate $r \leq r(\mathbf{p})$, the following upper bound $\tau^i(\mathbf{p}, r)$ on the end-to-end jitter applies:

$$\tau^i(\mathbf{p}, r) = \frac{\sigma^i + n(\mathbf{p})c^i}{r} + \sum_{l \in \mathbf{p}} \delta_l, \tag{9}$$

where $\delta_l \leq d_l$ is a constant component that depends on the maximal packet size sent over the link and on the link's speed (but not on its propagation delay).

Let $\tau^i(\mathbf{p}) = \tau^i(\mathbf{p}, r(\mathbf{p}))$ denote the minimal possible value of $\tau^i(\mathbf{p}, r)$, which will be referred to as *the guaranteed jitter* of $\mathbf{p}$. Accordingly, a path $\mathbf{p}$ between s^i and t^i is said to be *feasible* for connection i if $D^i(\mathbf{p}) \leq D^i$, $r(\mathbf{p}) \geq b^i$, *and* $\tau^i(\mathbf{p}) \leq \tau^i$; as before, we can disregard the bandwidth constraint b^i. Let:

$$\alpha^i(\mathbf{p}) = \max\left(\frac{\sigma^i + n(\mathbf{p})c^i}{\tau^i - \sum_{l \in \mathbf{p}} \delta_l}, \frac{\sigma^i + n(\mathbf{p})c^i}{D^i - \sum_{l \in \mathbf{p}} d_l}\right) \tag{10}$$

denote the minimal rate that satisfies both the delay and the jitter constraints of connection i over path $\mathbf{p}$. The set $\mathcal{P}^i$ of paths between source s^i and destination t^i that are feasible for a connection (of type) i is defined as in Definition 1.1, where $\alpha^i(\mathbf{p})$ is given by eq. (10). Therefore, the QoS-routing problem faced by a user i can be stated as follows.

Jitter–Delay Problem (JD):

Find a path $\mathbf{p}^* \in \mathcal{P}^i$ *that achieves the minimum:*

$$\min_{\mathbf{p} \in \mathcal{P}^i} (\sigma^i + n(\mathbf{p})c^i) \max\left(\frac{1}{\tau^i - \sum_{l \in \mathbf{p}} \delta_l}, \frac{1}{D^i - \sum_{l \in \mathbf{p}} d_l}\right) \sum_{l \in \mathbf{p}} w_l.$$

Proposition 3.1 Problem JD *is NP-Complete.*

The above situation is typical of (multi-)constrained path optimization. These problems are, in general, intractable, but solvable through pseudo-polynomial algorithms, based on dynamic programming principles. A pseudo-polynomial solution for *Problem JD* would incur a time complexity of $O(D^i MH)$. If users could be expected to afford such a solution, the equilibrium conditions would be the same as (2)–(4), only that $\alpha^{n,d,i}$ now becomes:

$$\alpha^{n,d,i} = \max\left(\frac{\sigma^i + nc^i}{\tau^i - \sum_{l \in \mathbf{p}} \delta_l}, \frac{\sigma^i + nc^i}{D^i - \sum_{l \in \mathbf{p}} d_l}\right). \tag{11}$$

Proceeding as before, it can be shown that theorems 2.1 and 2.2 hold here too, *i.e.*, the manager can lead the users to the (system's) optimal point by setting $w_l = J'_l$. However, pseudo-optimal solutions are often prohibitively complex for connection establishment. Hence, when facing both delay and jitter constraints, users may not be able to identify their optimal solutions. This means that we cannot establish equilibrium conditions that are based on user-optimality considerations. On the other hand, such constrained path optimization problems have efficient approximation schemes, which provide ϵ-optimal solutions within polynomial time complexity (e.g., Hassin (1992), Orda (1999) and references therein).

In the following, we investigate the possible design of efficient pricing schemes when users employ such ϵ-optimal approximations. In order to concretize the discussion, we focus on approximations that are based on *scaling* techniques (Hassin (1992)). In our context, it means that users replace the original set of link prices $\{w_l\}$ with a new set $\{\widehat{w}_l\}$ of scaled values. Compared to prices $\{w_l\}$, scaled prices $\{\widehat{w}_l\}$ take values in a smaller set, thus reducing the size of the problem. To simplify the presentation, we assume that there is a lower bound $w_{\min}$ on the price that can be imposed on a link (per unit of rate). By considering the scaled prices $\{\widehat{w}_l\}$ and running an appropriate (dynamic-programming based) algorithm, a user can identify a feasible path whose price is at most $1+\epsilon$ times larger than the optimal value. Time complexity varies among the various algorithmic versions, however it is *polynomial* in the input size $(N, \log D^i)$ and in $\frac{1}{\epsilon}$.

The structure of the QoS-routing game is similar to the simpler case of *Problem U*. Equilibrium flow vectors $\mathbf{f}$ and $\boldsymbol{\gamma}$ must satisfy the feasibility constraints $\mathcal{F}$. The equilibrium conditions are identical to (2)–(4), where $\{w_{uv}(f_{uv}\})$ are replaced with $\{\widehat{w}_{uv}(f_{uv})\}$.

Let $\widehat{J}_l(f_l) \triangleq \int \widehat{w}_l(f_l) df_l$ and $\widehat{J}(\mathbf{f}) \triangleq \sum_{l\in\mathcal{L}} \widehat{J}_l(f_l)$. Following the same steps as in the analysis of *Problems U* and *M*, one can prove that Proposition 2.2, and Theorems 2.1 and 2.2, hold in this case too, but for link price functions $\{\widehat{w}_l(f_l)\}$ and a social function $\widehat{J}(\mathbf{f})$. In particular, this means that the network has a unique equilibrium $(\widehat{\mathbf{f}}, \widehat{\gamma})$, which minimizes the cost function $\widehat{J}$ over $\mathcal{F}$.

The relation between $\widehat{w}_l$ and w_l depends on the precise scaling method used by the user. A general (and quite conservative) bound is the following:

$$w_l(f_l) \leq \widehat{w}_l(f_l) \leq w_l(f_l) + \epsilon w_{\max} \tag{12}$$

where $w_{\max}$ is an upper-bound on the price of a link under a network flow that is fairly close to the optimum $\mathbf{f}^*$; in other words, it is a bound on the price at a link under reasonable congestion conditions. Consequently, we obtain the following result.

Theorem 3.1 *For link price functions $w_l(f_l) \triangleq J_l'(f_l)$, $l \in \mathcal{L}$, the network has a unique equilibrium $(\widehat{\mathbf{f}}, \widehat{\gamma})$, for which the following relation holds:*

$$J(\widehat{\mathbf{f}}) \leq J(\mathbf{f}^*) \cdot \left(1 + \epsilon \frac{w_{\max}}{w_{\min}}\right). \tag{13}$$

As prices can be expected to take relatively small values at operating points that are in the proximity of $\mathbf{f}^*$, we conclude that the ratio $\frac{w_{\max}}{w_{\min}}$ can be expected to be of manageable size. Hence, the above theorem establishes that, by setting link price functions that are equal to the derivatives of the social cost functions, the manager can drive the network into an efficient equilibrium.

Remark: Consider a general constrained path optimization problem, in which links are characterized by flow-sensitive costs $\{c_l(f_l)\}$ and constant delays $\{d_l\}$, and users attempt to identify paths of minimal cost while obeying an end-to-end delay constraint D. That is, the user problem is to identify a path $\mathbf{p}$ such that $\sum_{l\in\mathbf{p}} c_l$ is minimal while $\sum_{l\in\mathbf{p}} d_l < D$. Under this setting there is no direct dependence of path costs and delays – as opposed to Problem U where the two values depend on the rate – making the problem NP-complete. Hence, users are expected to employ suboptimal approximations. The analysis in this section can accommodate this general scenario, by simply dropping the dependence on the hop count, setting $w_l(f_l) \leftarrow c_l(f_l)$ on all $l \in \mathcal{L}$, and setting the $\alpha^{n,d,i}$ values to 1.

4 Conclusion

QoS network architectures that employ rate-based scheduling mechanisms have been widely investigated, and eventually consolidated into concrete proposals, e.g., Shenker et al. 1997. However, the efficient consumption of rates throughout the network remained an open problem, whose successful resolution is a requisite for the successful deployment of such architectures. In this study we established a novel solution methodology, based on a load-sensitive pricing mechanism.

Congestion pricing, as a network management tool, has been proposed in the past, both in the context of *best-effort* communication networks (e.g., Korilis et al. (1998), MacKie-Mason and Varian (1995)), as well as in the (rather similar) context of transportation networks (Dafermos and Sparrow (1971)). A novelty of the present study is to consider congestion at the *rate reservation* level, and price rates accordingly. A major complication in such a framework is that the path optimization problems faced by the users are considerably more complex than in a best-effort (or transportation) network. Indeed, in the latter, the user's routing scheme consists of a simple shortest-path computation; while with QoS routing, the user's scheme has been shown to be much more complex, in particular when rate-based schedulers are involved (Ma and Steenkiste (1997), Orda (1999), Pornavalai et al. (1997)). Moreover, many QoS routing problems are intractable, hence only sub-optimal solutions can be sought. Yet, we have shown that, in spite of these difficulties, simple and efficient pricing schemes can be constructed.

More specifically, we considered the employment of a pricing scheme that is based on rate consumption and depends on the aggregate reservation at the various links. Investigation of the users' response to such prices revealed the existence of a unique network equilibrium. We then showed that the network manager can drive the system into social optimum, by employing a simple pricing scheme that is based on marginal costs. The fact that such a price function is incentive compatible is not surprising and has been observed in the past for various (convex) network routing games, e.g. Dafermos and Sparrow (1971). However, in our context, the path optimization problem faced by the user is much more complex than in standard scenarios, and involves the optimization of rational objective functions; yet, interestingly, the structure of the sought incentive prices remains as simple as in the elementary setting of a single-class transportation network (Dafermos and Sparrow (1971)), where users just choose "shortest paths". Indeed, here too, the incentive compatible price function simply comprises of link prices, each being the derivative of the social link cost function, and is independent of the particular parameters of the various connections.

Our study is also the first to consider the relation between pricing schemes and multi-constrained path optimization. Specifically, we established the intractability of the related QoS routing problem when both pricing optimization as well as delay and jitter requirements are considered. The implication of this result is that, in general, users cannot identify their optimal paths, however we indicated that ϵ-optimal solutions can be expected. Accordingly, we established that, under certain conditions, the manager can still drive the network to an efficient operating point. Moreover, the corresponding pricing scheme is the same as in the standard setting, i.e. when users do identify their optimal paths. We generalized our analysis to constrained path optimization problems that are not necessarily related to rate-based schedulers.

An important implication of our results is that the manager does not need to know the exact characteristics or behavior of the network users. Specifically, the manager does not need to be aware of the specific values of the connection parameters and QoS requirements, neither it should know the exact structure of their QoS routing schemes. Moreover, the manager can choose to neglect some of the link properties, such as the constant delays. Indeed, the pricing solution is insensitive to all these, and depends solely on the corresponding social function.

Some important issues remain for future investigation. One is the incorporation of such pricing schemes in actual protocols. Another issue of practical importance is the dynamic behavior and convergence properties of these schemes. In addition, we conjecture that better performance bounds can be established for the multi-constrained (ϵ-optimal) case. While much is yet to be learned and understood, we believe that the findings of the present study provide encouraging support for the deployment of incentive compatible pricing schemes in rate-based QoS network architectures.

Notes

1. The terms "connection" and "user" are used interchangeably.

2. If the maximal available rate is a nodal property, then we associate it with all its outgoing links.

3. Depending on the pricing scheme, the network might have multiple equilibria, a unique equilibrium, or no equilibrium at all. Evidently, the manager would avoid pricing schemes that prevent the network from reaching an operating point. Similarly, schemes that give rise to multiple operating points are, in general, undesirable.

4. Indeed, the expected number of such connections is $\gamma_u^{n,d,i}(j)$, and each requires a reservation of $\alpha^{n,d,i}$ units of rate.

References

Apostolopoulos G., R. Guérin, S. Kamat, A. Orda, T. Przygienda, and D. Williams. (1999). "QoS routing mechanisms and OSPF extensions." RFC No. 2676. Internet RFC.

Crawley, E., R. Nair, B. Rajagopalan, and H. Sandick. (1998). "A Framework for QoS-based Routing in the Internet." RFC No. 2386. Internet RFC.

Dafermos, S. and F. T. Sparrow. (1971) "Optimal resource allocation and toll patterns in user-optimised transport networks." *Journal of Transport Economics and Policy*, 184–200.

Edell, R. J., N. McKeown, and P. P. Varaiya. (1995). "Billing Users and Pricing for TCP." *IEEE Journal on Selected Areas in Communications*, 13, 1162–1175.

Fudenberg, D., and J. Tirole. (1992). *Game Theory.* Cambridge, Massachusetts: MIT Press.

Garey, M. and D. Johnson. (1979). *Computers and Intractability.* San Francisco: Freeman.

Georgiadis, L., R. Guérin, V. Peris, and K. N. Sivarajan. (1996). "Efficient Network QoS Provisioning Based on per Node Traffic Shaping." *IEEE/ACM Transactions on Networking* 4(4), 482–501.

Guérin, R. and A. Orda. (1999). "QoS-based routing in networks with inaccurate information: Theory and algorithms." *IEEE/ACM Transactions on Networking* 7(3), 350–364.

Guérin, R., A. Orda, and D. Williams. (1997). "QoS Routing Mechanisms and OSPF Extensions." In *Proceedings of the 2nd IEEE Global Internet Mini-Conference*, Phoenix, AZ.

Hassin, R. (1992). "Approximation Schemes for the Restricted Shortest Path Problem." *Mathematics of Operations Research* 17(1), 36–42.

Jordan, S. and H. Jiang. (1995). "Connection Establishment in High Speed Networks." *IEEE Journal on Selected Areas in Communications* 13, 1150–1161.

Kelly, F.P. (1993). "Tarrifs and Effective Bandwidths in Multiservice Networks." *Proceedings of ITC'94* 401–410.

Korilis, Y.A., T. A. Varvarigou, and S. R. Ahuja. (1998). "Incentive-Compatible Pricing Strategies in Noncooperative Networks." In *Proceedings of the IEEE INFOCOM'98*, San Francisco, CA.

Korilis, Y.A., and A. Orda (2002). "Incentive Compatible Pricing Strategies for QoS Routing." Department of Electrical Engineering, Technion, EE Pub. 1310, may be obtained from ftp://ftp.technion.ac.il/pub/supported/ee/Network/ko.price02.ps

Low, S.H., and P. P. Varaiya. (1993). "A New Approach to Service Provisioning in ATM Networks." *IEEE/ACM Transactions on Networking* 1(3), 547–553.

Ma, Q., and P. Steenkiste. (1997). "Quality of Service Routing for Traffic with Performance Guarantees." In *Proceedings of IWQoS'97*, Columbia University, New York, NY.

MacKie-Mason, J. K., and H. R. Varian. (1995). "Pricing Congestible Network Resources." *IEEE Journal on Selected Areas in Communications* 13(7), 1141–1149.

Mendelson, H. and S. Whang. (1990). "Optimal Incentive-Compatible Priority Pricing for the M/M/1 Queue." *Operations Research* 38(5), 870–883.

Myerson, R.B., (1991). *Game Theory: Analysis of Conflict.* Cambridge, Massachusetts: Harvard University Press.

Orda, A., (1999). "Routing with end to end QoS guarantees in broadband networks." *IEEE/ACM Transactions on Networking* 7(3), 365–374.

Orda, A. and N. Shimkin. (2000). "Incentive Pricing in Multi-Class Systems." *Telecommunication Systems* 13, 241–267.

Parekh, A. K., and R. Gallager. (1994). "A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: the Multiple Node Case." *IEEE/ACM Transactions on Networking* 2, 137–150.

Pornavalai, C., G. Chakraborty, and N. Shiratori. (1997). "QoS Based Routing Algorithm in Integrated Services Packet Networks." In *Proceedings of IEEE ICNP'97*, Atlanta, GA.

PNNI (Private Network-Network Interface) Specification v1.0. (1996). ATM Forum Technical Committee.

Sairamesh, J., D. F. Ferguson, and Y. Yemini. (1995). "An Approach to Pricing, Optimal Allocation and Quality of Service Provisioning in High-Speed Packet Networks." In *Proceeding of the IEEE INFOCOM*, Boston, MA.

Shenker, S., C. Partridge, and R. Guerin. (1997). "Specification of Guaranteed Quality of Service." RFC No. 2212. Internet RFC.

Wang, Z. and J. Crowcroft. (1996). "Quality-of-Service Routing for Supporting Multimedia Applications." *IEEE JSAC* 14(7), 1288–1234.

Zhang, H. (1995). "Service Disciplines for Guaranteed Performance Service in Packet-Switching Networks." *Proceedings of the IEEE* 83(10), 1374–1399.