

Incentive Compatible Pricing Strategies for QoS Routing

Yannis A. Korilis
Bell Laboratories
Lucent Technologies
Holmdel, NJ 07733, USA
yannisk@bell-labs.com

*Ariel Orda**
Dept. of Electrical Engineering
Technion
Haifa 32000, Israel
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 the case in which users can identify only sub-optimal paths, as is often the case with multiobjective path optimization.

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

I. INTRODUCTION

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., [1], [3], [9], [18], [22], [26], [29] 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. Recent studies have proposed schedulers that map delay guarantees into rate require-

ments and have each node advertise its residual rate [8], [24], [30]. The Guaranteed Service Class proposed for the Internet [28] is based on such rate-based principles. Based on the properties of such schedulers, several recent studies [10], [11], [18], [22], [26] have analyzed algorithms for computing paths that satisfy end-to-end delay bounds. In particular, it has been shown in [10] 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 [18], [26] 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 connection¹ 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 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; see, e.g., [20]. 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 [5], [13], [14], [16], [17], [19], [23], [27]. 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 sec-

*Part of this work was done while visiting Bell Laboratories.

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

and 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 belong 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* [6], [21]. 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 [4].

In the sequel, we turn our attention to connections that con-

duct multiobjective, constraint path optimization. A typical setting is to identify a path that minimizes some target function, e.g., administrative costs, while observing one 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 multiobjective path optimization problems are, in general, NP-complete [7], 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., [12]). 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 II we present the QoS-routing model and formulate the problem of incentive compatible pricing. Focusing on end-to-end delay constraints, in Section III we investigate the structure of the QoS-routing game and study the problem of incentive compatible pricing. The multiobjective path optimization case is considered in Section IV. Conclusions are presented in Section V. Due to space limits, proofs are omitted from this version, and can be found in [15].

II. 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 $M = |\mathcal{L}|$. 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 $In(u) = \{l : T(l) = u\}$ denote the set of its ingoing links, and $Out(u) = \{l : S(l) = u\}$ the set of its outgoing links. Each link $l \in \mathcal{L}$ 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 transfer unit. Without loss of generality, we assume that d_l takes integer values.

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

- a source node s^i and a destination node t^i ,
- 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 .

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

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, $r(\mathbf{p}) = \min_{l \in \mathbf{p}} R_l$. We assume that the scheduling policy in the network belongs to the *rate-based* class [8], [24], [30], as in the specification of the Guaranteed Service for IP [28]. Accordingly, 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:

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

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 s^i and t^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 constrained 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 II.1: A path \mathbf{p} between s^i and t^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.

The arrival process of type i connections with source node u and destination node j is some ergodic process with rate $\gamma_u^i(j)$. We assume that there is an infinite number of type i connections, each with infinitesimally small rate. 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, $r \sum_{l \in \mathbf{p}} w_l$. 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 . It is important here 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. $J_l : [0, R_l] \rightarrow [0, \infty]$ is continuous.
- G2. J_l is increasing and strictly convex.
- G3. J_l is continuously differentiable.

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 which coincides with the network optimum.³ 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. $w_l : [0, R_l] \rightarrow [0, \infty]$ 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 investigated in Section IV.

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 $l \in \mathcal{L}$.

III. 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.

³Depending 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.

For ease of exposition, we consider first, in Subsection III-A, the simple case of a *homogeneous network*. The results are extended to the case of a general network in Subsection III-B.

A. Homogeneous Network

Consider a network with homogeneous connections – i.e., $D^i \equiv D$, $\sigma^i \equiv \sigma$ and $c^i \equiv c$ – and homogeneous links – i.e., $d_l \equiv d$ – that will be referred to as a *homogeneous network*. Denote by $\gamma_u(j)$ the aggregate (arrival) rate of connections between a source node u and a destination node j . Let $\alpha^n = \frac{\sigma + nc}{D - nd}$, i.e., α^n denotes the minimal rate required by a connection over an n -hop path.

Let us start by studying the routing problem faced by a user. Given a set of nonnegative link prices $\{w_l\}$,⁴ the routing problem of a type i user can be defined as follows.

Homogeneous Network – User Problem (HU):

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

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

It is easy to see that *Problem HU* is solved by the following algorithm.

Algorithm HU:

1. For all $1 \leq n < \frac{D}{d}$, run a Bellman-Ford shortest-path algorithm [2], with respect to the metric $\{w_l\}$, in order to find a shortest path \mathbf{p}^n , among those that have at most n hops. Denote: $W^n = \sum_{l \in \mathbf{p}^n} w_l$.
2. Let $n^* = \arg \min_{1 \leq n < \frac{D}{d}} \alpha^n W^n$. The required path is $\mathbf{p}^* = \mathbf{p}^{n^*}$.

We now proceed to characterize the structure of the QoS-routing game. Considering the network at equilibrium (assuming that one exists), denote by $\gamma_u^n(j)$ the aggregate (arrival) rate of connections from source node u to destination node j , which are established over n -hop paths. Similarly, let $f_l^n(j)$ be the aggregated flow over a link l , corresponding to connections destined to node j , which are routed over paths for which link l is the n -th hop away from the destination. By a slight abuse of notation, define the link flow vector $\mathbf{f} = \{f_l^n(j) : 1 \leq n < \frac{D}{d}, l \in \mathcal{L}, j \in \mathcal{V}\}$. Similarly, define $\gamma = \{\gamma_u^n(j) : 1 \leq n < \frac{D}{d}, u, j \in \mathcal{V}, u \neq j\}$. Flow vectors \mathbf{f} and γ must satisfy the following feasibility constraints.

Constraints \mathcal{FH} :

1. $\forall j, u \in \mathcal{V} \ u \neq j, \forall 1 \leq n < \frac{D}{d} : \sum_{v \in \mathcal{V}} f_{uv}^n(j) = \sum_{v \in \mathcal{V}} f_{vu}^{n+1}(j) + \gamma_u^n(j)\alpha^n$;
2. $\forall j, u \in \mathcal{V} \ u \neq j : \sum_{n=1}^{\lceil \frac{D}{d} \rceil - 1} \gamma_u^n(j) = \gamma_u(j)$;
3. $\forall j \in \mathcal{V}, \forall (u, v) \in \mathcal{L}, \forall 1 \leq n < \frac{D}{d} : f_{uv}^n(j) \geq 0$;
4. $\forall j, u \in \mathcal{V} \ u \neq j, \forall 1 \leq n < \frac{D}{d} : \gamma_u^n(j) \geq 0$.

The first constraint is a flow conservation constraint: it states that $\sum_{v \in \mathcal{V}} f_{uv}^n(j)$, i.e., the aggregate flow to a destination j for which a node u is at the starting-point of the n -th hop, is equal to the sum of $\sum_{v \in \mathcal{V}} f_{vu}^{n+1}(j)$, i.e., the aggregate flow to j for which node u is at the end-point of the $(n + 1)$ -st hop, plus

⁴Recall that there is an infinite number of connections each with infinitesimally small rate. Thus, the routing decisions of a single user do not affect link prices, which can be treated as constants with respect to the user optimization problem.

$\gamma_u^n(j)\alpha^n$, i.e., the aggregate rate of connections established between u and j over n -hop paths.⁵ The second constraint simply states that each connection should be routed over some path. The third and forth are obvious nonnegativity constraints. Let \mathcal{FH} denote the set of feasible flow vectors (\mathbf{f}, γ) for the homogeneous network, i.e., those that obey the above four constraints.

Let (\mathbf{f}, γ) be an equilibrium flow vector. Then, the total flow on any link $l \in \mathcal{L}$ is $f_l = \sum_{n=1}^{\lceil \frac{D}{d} \rceil - 1} \sum_{j \in \mathcal{V}} f_l^n(j)$. By *Algorithm HU*, it is clear that a path chosen by a user is a shortest n -hop path with respect to the metric $\{w_l(f_l)\}$, for some n , $1 \leq n < \frac{D}{d}$. This observation leads to the following equilibrium conditions:

$$\forall u, j \in \mathcal{V}, u \neq j, \forall v \in \text{Out}(u), \forall 1 \leq n < \frac{D}{d}:$$

$$\lambda_u^n(j) \begin{cases} = w_{uv}(f_{uv}) + \lambda_v^{n-1}(j) & , f_{uv}^n(j) > 0 \\ \leq w_{uv}(f_{uv}) + \lambda_v^{n-1}(j) & , f_{uv}^n(j) = 0 \end{cases} \quad (2)$$

$$\forall u, j \in \mathcal{V}, u \neq j, \forall 1 \leq n < \frac{D}{d}:$$

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

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

The equilibrium conditions can be explained as follows. $\lambda_u^n(j)$ is the minimal equilibrium price for reserving one unit of rate over an n -hop path between a source u and a destination j ; expression (2) simply states that, among paths with the same number of hops, users will prefer those with minimal per-unit rate prices. $\tilde{\lambda}_u(j)$, in turn, is the minimal equilibrium cost for establishing a connection 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 in a homogeneous network is given by the following theorem.

Theorem III.1: In a homogeneous network, for any set of link price functions $\{w_l(f_l)\}$ that satisfy Assumptions P:

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

We now turn our attention to the manager's problem. The manager aims at enforcing the network optimum, i.e., a flow vector (\mathbf{f}, γ) that minimizes the social function $J(\cdot)$ over the set of all feasible flow vectors. The manager can only decide upon link price functions $\{w_l(f_l)\}$. According to Theorem III.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.

Homogeneous Network – Manager Problem (HM):

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}, γ) that solves the following problem.

⁵Indeed, such connections arrive at rate $\gamma_u^n(j)$, and each requires a reservation of α^n units of rate.

Minimize:

$$J(\mathbf{f}) = \sum_{l \in \mathcal{L}} J_l(f_l) = \sum_{l \in \mathcal{L}} J_l\left(\sum_{n=1}^{\lceil \frac{D}{d} \rceil - 1} \sum_{j \in \mathcal{V}} f_l^n(j)\right),$$

subject to Constraints \mathcal{FH} .

Before we investigate the existence of incentive compatible pricing strategies, let us first prove the following proposition, which characterizes the structure of the network optimum through a set of necessary and sufficient conditions.

Proposition III.1: Given a social cost function $J(\cdot)$ that conforms with Assumptions G and a homogeneous network, a flow vector (\mathbf{f}, γ) is the network optimum if and only if there exist (Lagrange multipliers) $\{\lambda_u^n(j) : 1 \leq n < \frac{D}{d}, u, j \in \mathcal{V}, u \neq j\}$ and $\{\tilde{\lambda}_u(j) : u, j \in \mathcal{V}, u \neq j\}$, such that the following (Kuhn-Tucker) conditions are satisfied:

$$\forall u, j \in \mathcal{V}, u \neq j, \forall v \in \text{Out}(u), \forall 1 \leq n < \frac{D}{d}:$$

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

$$\forall u, j \in \mathcal{V}, u \neq j, \forall 1 \leq n < \frac{D}{d}:$$

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

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

Theorem III.1 and Proposition III.1 imply that link price functions $\{J'_l(f_l)\}$ lead to a unique equilibrium that minimizes the social cost function. We have, thus, established the main result of this subsection.

Theorem III.2: For link price functions $w_l(f_l) \triangleq J'_l(f_l)$, $l \in \mathcal{L}$, the flow vector (\mathbf{f}^*, γ^*) that minimizes the social cost function J over \mathcal{FH} is the unique equilibrium of the homogeneous network.

The implication of Theorem III.2 is that, by setting link price functions that are equal to the derivatives of the social cost functions, the manager can drive a homogeneous network to the required social optimum.

B. General Network

We consider now the general case, i.e., general connections – with values D^i, σ^i, c^i – and general links – with values d_l . Let $\alpha^{n,d,i} = \frac{\sigma^i + n c^i}{D^i - d}$, 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 .

Again, we 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 in a general network can be defined as follows:

General Network – User Problem (GU):

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 [10], where a polynomial solution was established. The complexity of that solution, when applied to our case, is $O(N^2 M \log(N \frac{\max_{l \in \mathcal{L}} w_l}{\min_{l \in \mathcal{L}} w_l} D^i))$. Although the complexity here is higher than in the homogeneous case, it is still of a reasonable polynomial rank. Thus, just as in the homogeneous case, users can be expected to identify their optimal paths, given a set of link prices. To summarize:

Proposition III.2: Problem GU 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 in a general network. Considering the network at equilibrium (assuming that one exists), denote by $\gamma_u^{n,d,i}(j)$ the aggregate (arrival) rate 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 $\gamma = \{\gamma_u^{n,d,i}(j) : 1 \leq n < H, i \in \mathcal{I}, 0 \leq d \leq D^i, u, j \in \mathcal{V}, u \neq j\}$. Flow vectors \mathbf{f} and γ must satisfy the following feasibility constraints.

Constraints \mathcal{FG} :

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

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

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. $\forall i, \forall j \in \mathcal{V}, \forall (u, v) \in \mathcal{L}, \forall 1 \leq n \leq H, \forall 0 \leq d < D^i: f_{uv}^{n,d,i}(j) \geq 0;$
4. $\forall j, u \in \mathcal{V} u \neq j, \forall 1 \leq n \leq H, \forall 0 \leq d < D^i: \gamma_u^{n,d,i}(j) \geq 0.$

The above constraints are similar to the set of constraints \mathcal{FH} , except that now we differentiate flows not only according to the destination and remaining number of hops, but also according to the connection type and remaining units of constant delay. Let \mathcal{FG} denote the set of feasible flow vectors (\mathbf{f}, γ) for a general network, i.e., those that obey the above four constraints.

Let (\mathbf{f}, γ) 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 GU*) 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 $\sum_{l \in \mathbf{p}_1} w_l \leq \sum_{l \in \mathbf{p}_2} w_l$. This observation leads to the following equilibrium conditions:

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

$$\lambda_u^{n,d}(j) \begin{cases} = w_{uv}(f_{uv}) + \lambda_v^{n-1,d-d_{uv}}(j) & , f_{uv}^{n,d,i}(j) > 0 \\ \leq w_{uv}(f_{uv}) + \lambda_v^{n-1,d-d_{uv}}(j) & , f_{uv}^{n,d,i}(j) = 0 \end{cases} \quad (8)$$

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

$$\tilde{\lambda}_u^i(j) \begin{cases} = \lambda_u^{n,d}(j) \alpha^{n,d,i} & , \gamma_u^{n,d,i}(j) > 0 \\ \leq \lambda_u^{n,d}(j) \alpha^{n,d,i} & , \gamma_u^{n,d,i}(j) = 0 \end{cases} \quad (9)$$

$$\lambda_j^{n,d}(j) \equiv \tilde{\lambda}_j^i(j) \equiv 0. \quad (10)$$

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 (8) 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 (9) is just a restatement of the observation made above.

The structure of the QoS-routing game in a general network is given by the following theorem.

Theorem III.3: In a general network, for any set of link price functions $\{w_l(f_l)\}$ that satisfy Assumptions P:

1. There exists a unique equilibrium flow vector $(\mathbf{f}, \boldsymbol{\gamma}) \in \mathcal{FG}$.
2. Conditions (8)–(10) are necessary and sufficient for equilibrium, that is: a flow vector $(\mathbf{f}, \boldsymbol{\gamma}) \in \mathcal{FG}$ 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 (8)–(10) hold.

We now turn our attention to the problem of the manager. As before, the manager aims at selecting link price functions $\{w_l(f_l)\}$ that enforce the network optimum, i.e., a flow vector $(\mathbf{f}, \boldsymbol{\gamma})$ that minimizes the social function $J(\cdot)$ over the set of all feasible flow vectors. According to Theorem III.3, 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.

General Network – Manager Problem (GM):

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 $(\mathbf{f}, \boldsymbol{\gamma})$ that solves the following problem.

Minimize:

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

subject to Constraints \mathcal{FG} .

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

Proposition III.3: Given a social function $J(\cdot)$ that conforms with Assumptions G and a general network, a flow vector $(\mathbf{f}, \boldsymbol{\gamma})$ 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 $\{\tilde{\lambda}_u^i(j) : j, u \in \mathcal{V}, u \neq j\}$, such that the following (Kuhn-Tucker) conditions are satisfied:

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

$$\lambda_u^{n,d}(j) \begin{cases} = J'_{uv}(f_{uv}) + \lambda_v^{n-1,d-d_{uv}}(j) & , f_{uv}^{n,d,i}(j) > 0 \\ \leq J'_{uv}(f_{uv}) + \lambda_v^{n-1,d-d_{uv}}(j) & , f_{uv}^{n,d,i}(j) = 0 \end{cases} \quad (11)$$

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

$$\tilde{\lambda}_u^i(j) \begin{cases} = \lambda_u^{n,d}(j) \alpha^{n,d,i} & , \gamma_u^{n,d,i}(j) > 0 \\ \leq \lambda_u^{n,d}(j) \alpha^{n,d,i} & , \gamma_u^{n,d,i}(j) = 0 \end{cases} \quad (12)$$

$$\lambda_j^{n,d}(j) \equiv \tilde{\lambda}_j^i(j) \equiv 0. \quad (13)$$

We have thus established the main result of this section.

Theorem III.4: For link price functions

$$w_l(f_l) \triangleq J'_l(f_l), \quad l \in \mathcal{L}, \quad (14)$$

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

Again, the implication of Theorem III.4 is that, by setting link price functions that are equal to the derivatives of the social cost functions, the manager can drive a *general* network to the required social optimum.

Remark III.1: It is important to note that, while the equilibrium conditions (8)–(10) 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 III.2. Similarly, the large number of variables with which the optimality conditions (11)–(13) are stated has no implication on the manager: as established by Theorem III.4, the size of its strategy, i.e., optimal price vector, is just $O(M)$.

B.1 Special Cases

The general network model, formulated above, includes the following cases of special interest.

- **Saturation:** Suppose that each link has a “capacity” c_l , and the goal of the manager is to avoid link saturation, i.e., f_l should be kept below c_l . This case can be accommodated by choosing any price function w_l , which complies with our standard assumptions, and, in addition, satisfies $\lim_{f_l \rightarrow c_l} w_l(f_l) = \infty$.

- **Exponential prices:** Taking a competitive analysis standpoint, It has been indicated that it is a good practice to route connections over paths that minimize the sum of link weights, each being some *exponential* function of the relative load on the link (see, e.g., [25] and references therein). In our framework, implementing this indication would correspond to choosing link prices that are some function $F(\cdot)$ of an exponent of the relative load $\frac{f_l}{c_l}$, where c_l stands for the capacity of a link l ;

i.e., $w_l(f_l) = F(e^{\frac{f_l}{c_l}})$, where $F(\cdot)$ should be increasing and (weakly) convex.

- **Elastic demand:** Suppose that a user may choose not to establish a connection, say of type i , when the incurred price is too high, say higher than some value w_{\max}^i . This environment can be modeled in our framework in the following way. Consider each traffic demand $\gamma_u^i(j)$, i.e., type- i traffic between a source u and a destination j : add a fictitious node s^i , which becomes the new source of $\gamma_u^i(j)$; connect s^i to u through a fictitious link (s^i, u) , and set $d_{s^i, u} = w_{s^i, u} = 0$, i.e., (s^i, u) has zero constant delay and zero price; connect s^i to j through a fictitious link (s^i, j) , and set $d_{s^i, j} = D^i - (c^i + c^j)$, $w_{s^i, j} = w_{\max}^i$. It is easy to verify that a connection of type i , to destination j , can be (fictitiously) established over the direct link (s^i, j) with a single unit of reserved rate and corresponding price w_{\max}^i ; thus, whenever

the connection cannot be established over the real network at a lower price, it is deferred to the fictitious link, hence modeling an actual cancellation.

• *Target operating point:* In some cases, the target operating point might be specified explicitly, rather than implied by some social cost function. An easy way to accommodate such cases in our framework is as follows. Given a target operation point \mathbf{f}^* ,⁶ choose the following price functions: $w_l^*(f_l) = f_l - f_l^*$.⁷ Indeed, the unique minimum of the cost function $J_l^* = \int w_l^* df_l$ is \mathbf{f}^* . By Theorem III.4, this is also the network's unique equilibrium point. Intuitively, w_l^* discourages the over-utilization of a link, by imposing (positive) tolls when $f_l > f_l^*$; whereas when $f_l < f_l^*$, an incentive is given to use that link, in the form of a "back payment" to the user. Once equilibrium is reached, the price over all links is zero.

C. Big User

Suppose that, except for the (infinitely) many (infinitely) small users, there is one "big" user, i.e., with a non-negligible amount of flow. For example, consider a network service provider that implements its infrastructure on top of the considered network. It is easy to verify that, in the presence of such a user, our regular pricing scheme will not work. Intuitively, such a user should take into account the externalities imposed by its routing decisions, i.e., the effect that each connection has on the rest of the user's flow. Hence, its optimal policy is different than that implied by *Problem GU*. In this subsection we devise a pricing scheme for the big user, which, together with the regular scheme elaborated for small users, constitute an incentive-compatible pricing strategy.

The arrival process of type i connections corresponding to the big user, with source node u and destination node j , is some ergodic process with rate $\gamma_u^{B,i}(j)$. Considering a network at equilibrium, denote by $\gamma_u^{B,n,d,i}(j)$ the aggregate arrival rate of type- i connections corresponding to the big user, with source node u and destination node j , which are established over paths with n hops and total constant delay d . Denote by f_l^B the aggregate flow of the big user on link l . Then, $f_l^{-B} = f_l - f_l^B$ is the aggregate flow of all other (small) users.

We impose on the big user a pricing scheme $\{w_l^B\}$, which depends both on the aggregate link flow f_l and on the user's flow f_l^B , i.e., $w_l^B = w_l^B(f_l^B, f_l)$. Given such price functions, and given \mathbf{f}^{-B} , the big user's problem on a general network can be stated as follows.

General Network – Big User Problem (GBU):

Minimize:

$$\begin{aligned} \sum_{l \in \mathcal{L}} w_l^B(f_l^B, f_l) &= w_l^B(f_l^B, f_l - f_l^B) = \\ &= \sum_{l \in \mathcal{L}} w_l^B \left(\sum_i \sum_{n=1}^H \sum_{d=0}^{D^i-1} \sum_{j \in \mathcal{V}} f_l^{B,n,d,i}(j), f_l^{-B} \right) \end{aligned}$$

subject to:

⁶ \mathbf{f}^* must be feasible.

⁷While these functions do not comply with Assumption P1, our analysis can accommodate a relaxation of that assumption.

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

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

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

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

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

Problems *GM* and *GU* remain the same, except that now the γ -values include the connections of both the small users and the big user.

Consider the following price functions:

$$w_l^B(f_l^B, f_l) = J_l(f_l) - J_l(f_l - f_l^B). \quad (15)$$

Intuitively, $w_l^B(f_l^B, f_l)$ charges the big user for the increment in the function J_l induced by its aggregate flow f_l^B . In [15] we establish the following result.

Theorem III.5: Let the price functions imposed on the small users be as in (14), and the price functions imposed on the big user be as in (15). Then, the flow vector \mathbf{f}^* that minimizes $J(\mathbf{f})$ over \mathcal{FG} is the network's unique equilibrium.

IV. INCENTIVE COMPATIBLE PRICING STRATEGIES FOR MULTI-OBJECTIVE 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 [28]. 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:

$$\tau^i(\mathbf{p}, r) = \frac{\sigma^i + n(\mathbf{p})c^i}{r}. \quad (16)$$

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:

$$\alpha^i(\mathbf{p}) = \max\left(\frac{\sigma^i + n(\mathbf{p})c^i}{\tau^i}, \frac{\sigma^i + n(\mathbf{p})c^i}{D^i - \sum_{l \in \mathbf{p}} d_l}\right) \quad (17)$$

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 II.1, where $\alpha^i(\mathbf{p})$ is given by eq. (17). 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:*

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

Proposition IV.1: *Problem JD* is NP-Complete.

The above situation is typical of multi-objective 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 $O(D^i M H)$. If users could be expected to afford such a solution, the equilibrium conditions would be the same as (8)–(10), only that $\alpha^{n,d,i}$ now becomes:

$$\alpha^{n,d,i} = \max\left(\frac{\sigma^i + nc^i}{\tau^i}, \frac{\sigma^i + nc^i}{D^i - \sum_{l \in \mathbf{P}} d_l}\right). \quad (18)$$

Proceeding as before, it can be shown that theorems III.3 and III.4 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 usually prohibitively complex for connection establishment. We thus conclude that, when facing both delay and jitter constraints, users cannot be expected 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 multi-objective path optimization problems have efficient approximation schemes, which provide ϵ -optimal solutions within polynomial time complexity (e.g., [12], [22] 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 [12]. In our context, that means that users replace the original set of link prices $\{w_l\}$ with a new set $\{\hat{w}_l\}$ of scaled values. Compared to prices $\{w_l\}$, scaled prices $\{\hat{w}_l\}$ 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 $\{\hat{w}_l\}$ 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 linear in $\frac{1}{\epsilon}$.

The structure of the QoS-routing game is similar to the simpler case of *Problem GU*. Equilibrium flow vectors \mathbf{f} and γ must satisfy the feasibility constraints \mathcal{FG} . The equilibrium conditions are identical to (8)–(10), where $\{w_{uv}(f_{uv})\}$ are replaced with $\{\hat{w}_{uv}(f_{uv})\}$.

Let $\hat{J}_l(f_l) \triangleq \int \hat{w}_l(f_l) df_l$ and $\hat{J}(\mathbf{f}) \triangleq \sum_{l \in \mathcal{L}} \hat{J}_l(f_l)$. Following the same steps as in the analysis of *Problems GU* and *GM*, one can prove that Proposition III.3, and Theorems III.3 and III.4, hold in this case too, but for link price functions $\{\hat{w}_l(f_l)\}$ and a social function $\hat{J}(\mathbf{f})$. In particular, this means that the network has a unique equilibrium $(\hat{\mathbf{f}}, \hat{\gamma})$, which minimizes the cost function \hat{J} over \mathcal{FG} .

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

$$w_l(f_l) \leq \hat{w}_l(f_l) \leq w_l(f_l) + \epsilon w_{\max} \quad (19)$$

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 IV.1: For link price functions $w_l(f_l) \triangleq J'_l(f_l)$, $l \in \mathcal{L}$, the network has a unique equilibrium $(\hat{\mathbf{f}}, \hat{\gamma})$, for which the following relation holds:

$$J(\hat{\mathbf{f}}) \leq J(\mathbf{f}^*) \cdot \left(1 + \epsilon \frac{w_{\max}}{w_{\min}}\right). \quad (20)$$

As prices can be expected to take small values, *i.e.*, around w_{\min} , at operating points that are in the proximity of \mathbf{f}^* , we conclude that the ratio $\frac{w_{\max}}{w_{\min}}$ is typically small. 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 GU 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.

V. CONCLUSION

QoS network architectures that employ rate-based scheduling mechanisms have been widely investigated, and eventually consolidated into concrete proposals, e.g., [28]. 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., [16], [19]), as well as in the (rather similar) context of transportation networks [4]. 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 [18], [22], [26]. 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. Consequently, we established a

fundamental result: in spite of the complexity of the QoS routing schemes employed by the users, the network manager can drive the system into social optimum, by employing a simple pricing scheme. The corresponding incentive compatible prices are the derivatives of the social function. It is remarkable that the solution is as simple as in the much simpler setting of a single-class transportation network [4]. Moreover, we indicated that other related problems of special interest can be accommodated as special cases of our general setting.

Our study is also the first to consider the relation between pricing schemes and multi-objective 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 multi-objective 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. Some related simulation work is currently being performed. In addition, we conjecture that better performance bounds can be established for the multiobjective (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.

REFERENCES

- [1] Private Network-Network Interface Specification v1.0 (PNNI). ATM Forum Technical Committee, March 1996.
- [2] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 1990.
- [3] E. Crawley, R. Nair, B. Rajagopalan, and H. Sandick. A Framework for QoS-based Routing in the Internet – RFC No. 2386. Internet RFC, August 1998.
- [4] S. Dafermos and F. T. Sparrow. Optimal resource allocation and toll patterns in user-optimised transport networks. *Journal of Transport Economics and Policy*, pages 184–200, 1971.
- [5] R. J. Edell, N. McKeown, and P. P. Varaiya. Billing Users and Pricing for TCP. *IEEE Journal on Selected Areas in Communications*, 13:1162–1175, September 1995.
- [6] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, Massachusetts, 1992.
- [7] M. Garey and D. Johnson. *Computers and Intractability*. Freeman, San Francisco, 1979.

- [8] L. Georgiadis, R. Guérin, V. Peris, and K. N. Sivarajan. Efficient Network QoS Provisioning Based on per Node Traffic Shaping. *IEEE/ACM Transactions on Networking*, 4(4):482–501, August 1996.
- [9] R. Guérin, S. Kamat, A. Orda, T. Przygienda, and D. Williams. QoS Routing Mechanisms and OSPF Extensions. Internet Draft, January 1998. Work in progress.
- [10] R. Guérin and A. Orda. QoS-based routing in networks with inaccurate state and metrics information. Research report, IBM, T. J. Watson Research Center, 1996. Conference version appeared in the Proceedings of Infocom'97, Kobe, Japan.
- [11] R. Guérin, A. Orda, and D. Williams. QoS Routing Mechanisms and OSPF Extensions. In *Proceedings of the 2nd IEEE Global Internet Mini-Conference*, Phoenix, AZ, November 1997.
- [12] R. Hassin. Approximation Schemes for the Restricted Shortest Path Problem. *Mathematics of Operations Research*, 17(1):36–42, February 1992.
- [13] S. Jordan and H. Jiang. Connection Establishment in High Speed Networks. *IEEE Journal on Selected Areas in Communications*, 13:1150–1161, September 1995.
- [14] F. P. Kelly. Tariffs and Effective Bandwidths in Multiservice Networks. *Proceedings of ITC'94*, pages 401–410, 1993.
- [15] Y. A. Korilis and A. Orda. Incentive Compatible Pricing Strategies for QoS Routing. Research Report No. 1182, Department of Electrical Engineering, Technion, Haifa, Israel, December 1998. May be obtained from "ftp://ftp.technion.ac.il/pub/supported/ee/Network/ko.price98.ps".
- [16] Y. A. Korilis, T. A. Varvarigou, and S. R. Ahuja. Incentive-Compatible Pricing Strategies in Noncooperative Networks. In *Proceedings of the IEEE INFOCOM'98*, San Francisco, CA, April 1998.
- [17] S. H. Low and P. P. Varaiya. A New Approach to Service Provisioning in ATM Networks. *IEEE/ACM Transactions on Networking*, 1(3):547–553, 1993.
- [18] Q. Ma and P. Steenkiste. Quality of Service Routing for Traffic with Performance Guarantees. In *Proceedings of IWQoS'97*, Columbia University, New York, NY, May 1997.
- [19] J. K. MacKie-Mason and H. R. Varian. Pricing Cognestible Network Resources. *IEEE Journal on Selected Areas in Communications*, 13(7):1141–1149, September 1995.
- [20] H. Mendelson and S. Whang. Optimal Incentive-Compatible Priority Pricing for the M/M/1 Queue. *Operations Research*, 38(5):870–883, 1990.
- [21] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, Massachusetts, 1991.
- [22] A. Orda. Routing with End to End QoS Guarantees in Broadband Networks. In *IEEE INFOCOM'98*, San Francisco, CA, March 1998.
- [23] A. Orda and N. Shimkin. Incentive Pricing in Multi-Class Communication Networks. In *Proceedings of the IEEE INFOCOM'97*, Kobe, Japan, April 1997.
- [24] A. K. Parekh and R. Gallager. A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: the Multiple Node Case. *IEEE/ACM Transactions on Networking*, 2:137–150, 1994.
- [25] S. Plotkin. Competitive routing in ATM networks. *IEEE J. Selected Areas in Comm.*, pages 1128–1136, August 1995. Special issue on Advances in the Fundamentals of Networking.
- [26] C. Pornavalai, G. Chakraborty, and N. Shiratori. QoS Based Routing Algorithm in Integrated Services Packet Networks. In *Proceedings of IEEE ICNP'97*, Atlanta, GA, October 1997.
- [27] J. Sairamesh, D. F. Ferguson, and Y. Yemini. An Approach to Pricing, Optimal Allocation and Quality of Service Provisioning in High-Speed Packet Networks. In *Proceeding of the IEEE INFOCOM*, Boston, MA, April 1995.
- [28] S. Shenker, C. Partridge, and R. Guerin. Specification of Guaranteed Quality of Service – RFC No. 2212. Internet RFC, November 1997.
- [29] Z. Wang and J. Crowcroft. Quality-of-Service Routing for Supporting Multimedia Applications. *IEEE JSAC*, 14(7):1288–1234, September 1996.
- [30] H. Zhang. Service Disciplines for Guaranteed Performance Service in Packet-Switching Networks. *Proceedings of the IEEE*, 83(10):1374–1399, October 1995.