

Proportional QoS in Differentiated Services Networks: Capacity Management, Equilibrium Analysis and Elastic Demands

Ishai Menache and Nahum Shimkin

Department of Electrical Engineering, Technion, Israel Institute of Technology,
Haifa 32000, Israel

{imenache@tx, shimkin@ee}.technion.ac.il

Abstract. The Differentiated Services (Diffserv) architecture is a scalable solution for providing Quality of Service (QoS) over packet switched networks. By its very definition, Diffserv is not intended to provide strict performance guarantees to its subscribers. We purpose in this paper a particular form of *relative* performance guarantees. Specifically, the network manager's goal is to maintain pre-defined ratios between common congestion measures over the different service classes. We assume that each service class is advertised with a constant price. Thus, in order to induce its goal, the manager dynamically allocates available capacity between the service classes. This scheme is studied within a network flow model, with self-optimizing users, where each user can choose the amount of flow to ship on each service class according to its service utility and QoS requirements. We pose the entire problem as a *non-cooperative game*. Concentrating on a simplified single-link model with multiple service classes, we establish the existence and uniqueness of the Nash equilibrium where the relative performance goal is obtained. Accordingly, we show how to compute and sustain the required capacity assignment. The extension to a general network topology is briefly outlined.

1 Introduction

Background and Motivation. The need for providing service differentiation over the Internet has been an ongoing concern in the networking community. The Differentiated Services (Diffserv) architecture [4] has been proposed by the IETF as a scalable solution for QoS provisioning. Instead of reserving resources per session (e.g., as in the Integrated Services (IntServ) model [22]), packets are marked to create a smaller number of packet classes, which offer different service qualities. The Diffserv proposal suggests to combine simple priority mechanisms at the network core with admission control mechanisms at the network edges only, in order to create diverse end-to-end services.

The two principal Diffserv classes that have been formalized are the Expedited Forwarding (EF) [7] and the Assured Forwarding (AF) [10] services. The premise of the EF is to provide no-loss and delay reduction to its subscribers. AF is intended for users who need reliable forwarding even in times of network congestion. Ongoing IETF work concentrates on defining the engineering

and architectural aspects of Diffserv-enabled routers (e.g., [2]). However, current technical specifications deliberately do not quantify the actual *service characteristics*, which users will obtain by using the above mentioned classes. Apparently, service characteristics would have to be defined and publicly declared in order to make the distinction between the service classes meaningful to the user (and possibly worth paying for).

The Diffserv network cannot offer strict quality guarantees, as resources are allocated to the service classes based on some average network conditions [6]. Instead, the provider may declare upper-bounds on QoS measures, or alternatively provide looser guarantees, such as probabilistic or time-dependent guarantees. Another option is to offer *relative* quality guarantees. This option can be easily quantified and advertised, as illustrated by the following exemplifying rule: “service class Q will offer an end-to-end average delay, which is at least two times less than any other class, independent of the level of congestion”. When a user buys class Q it is aware of what it gets, and expects the provider to uphold the agreement conditions. In this paper, we focus on the *proportional* QoS model, whereby QoS ratios between the classes are announced. Specifically, we concentrate on *delay* ratios, although analogous definitions may be suggested when considering other QoS measures. The proportional QoS model benefits from implementation-related pros. Ratios are easier to maintain in comparison with absolute end-to-end guarantees, primarily because they may hold for different levels of congestion, and secondly because keeping the ratios locally (on a node basis) leads to fulfilling this objective on the network level.

We shall examine capacity allocation as the main network management tool for achieving the proportional QoS design objective. The goal of capacity allocation is to keep the announced delay ratios, irrespectively of complementary means for network traffic control, such as pricing and admission control. Our focus in this paper is on a simplified single link network, where the network manager owns a fixed amount of capacity to be divided among the link’s offered service classes in order to maintain the QoS ratios objective. Generally, it is easy to calculate the appropriate capacity allocation when the traffic in each service class is fixed. In this paper, however, we consider the interaction of the user behavior and network conditions. Within a standard flow model (see [3], Sec. 5.4) we represent the user population as a finite set of self-optimizing decision makers. The users are heterogenous with respect to their cost functions, which reflect their *price-quantity-quality* tradeoffs. Furthermore, users may modify their flow quantities (i.e., *elastic* users [11]), and may also shift their traffic from one service class to the other in response to current congestion conditions. We pose the overall problem as a non-cooperative game between the manager and the (selfish) users, and explore the associated capacity management policies and equilibrium conditions.

In a recent paper [14], we have considered a similar problem with static (fixed) user demand in the general network context. The present paper extends this work (for a single-link network) to the case of elastic demand, as well as adding a price term to the overall user’s cost.

Related Literature. *Service differentiation approaches:* Several research papers have addressed the model of finitely many classes, in which no strict performance guarantees are given. The simplest approach for providing differentiated services is Odlyzko’s Paris Metro Pricing (PMP) proposal [17]. The idea of PMP is to create service differentiation by giving a different price to each service class. Other papers [9, 13, 16] explicitly consider elements such as the user model, the scheduling mechanism and the network objective (e.g., a social or an economic objective). The major concern is usually in calculating the prices that would lead to the network objective. An additional common ground is that the network does not declare any kind of QoS guarantees. Users are assumed to acquire the best deal there is with respect to their quality-price tradeoff. We deviate from the last assumption, by considering the upholding of the service characteristic as a primary management priority.

Selfish routing: Since our model assumes that users are allowed to split traffic between the service classes, our work is related to selfish routing models. Game-theoretic analysis is widely used to study the working conditions of these models. The involved issues include the existence and uniqueness of an equilibrium point, its calculation, and its properties (such as the degradation of performance due to selfish behavior, known as the “price of anarchy” [21]). A common routing model, originated from the field of transportation networks, has considered networks shared by infinitesimal users (see [1] for a survey). The case of finitely many users, each carrying substantial flow has been introduced to the networking literature more recently (see [12, 18, 20]). In [12], the equilibrium properties are applied for network design (namely, link capacity assignment), where the objective is to obtain the socially optimal working point. We use a similar routing model to represent the user’s choice of service class.

Proportional QoS: Dovrolis et al. [8] proposed a class of schedulers, based on the Proportional Delay Differentiation (PDD) model. This model aims at providing predetermined delay ratios. The schedulers are implemented by observing the history of the encountered delays (or alternatively, by measuring the delay of the packet at the head of each service class), and serving the class which most deviates from its nominal delay ratio. In the present work we do not rely on PDD schedulers, but rather use a capacitated links model that may be considered a proxy to existing scheduling schemes such as GPS [19].

Contribution and Organization. This paper proposes schemes for inducing proportional QoS through capacity allocation, when users can react to the allocation decisions. The precise definitions of the network and user models are given in Section 2. The analysis of this model is presented in Section 3, which establishes the existence and uniqueness of the equilibrium point for the network-users game, and presents an algorithm for its computation. An explicit formula is obtained for the best response map of the network, namely the capacity assignment that ensures the QoS-ratio objective for *fixed* network flows, which may be used as a basis for an adaptive capacity assignment scheme. Due to length

constraints, the proofs of some claims are omitted, the reader is referred to [15] for full details.

2 The Single Link Model

Our basic model considers a single link, which supports several service classes. As a stand-alone model, it may be viewed as an approximation of a single path in a network, where the variations in traffic due to other (intersecting) network paths are neglected. Let $\mathcal{I} = \{1, 2, \dots, I\}$ be a finite set of users, which share a link that offers a set of service classes $\mathcal{A} = \{1, 2, \dots, A\}$. Since each service class is characterized by its own price and performance measure (to be described in the sequel), it would be convenient to consider the link with its respective service classes as a two terminal (source-destination) network, which is connected by a set of parallel arcs. Each arc represents a different service class. Thus, the set of arcs is also denoted by \mathcal{A} , and the terms service class and arc are used interchangeably. We denote by f_a^i the flow which user i ships on arc a . User i is free to choose any assignment of $f_a^i \geq 0$. The total demand of each user will be denoted by $f^i \triangleq \sum_{a \in \mathcal{A}} f_a^i$. Turning our attention to an arc $a \in \mathcal{A}$, let f_a be the total flow on that arc, i.e., $f_a = \sum_{i \in \mathcal{I}} f_a^i$. Also, denote by \mathbf{f}_a the vector of all user flows on arc a , i.e., $\mathbf{f}_a = (f_a^1, \dots, f_a^I)$. The user flow configuration \mathbf{f}^i is the vector $\mathbf{f}^i = (f_1^i, \dots, f_A^i)$. The flow configuration \mathbf{f} is the vector of all user flow configurations, $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^I)$. A user flow configuration is *feasible* if its components obey the nonnegativity constraints, as described above. We denote by \mathbf{F}^i the set of all feasible user flow configurations \mathbf{f}^i , and by \mathbf{F} the set of all feasible flow configurations \mathbf{f} .

The network manager has a constant capacity C , to be divided between the service classes. This capacity cannot be statically assigned, since the manager cannot predict in advance the number of customers and their preferences. Practically, the network manager would modify the current capacity assignment at slower time scales than the user routing decisions, after periodically measuring class performance. We denote by c_a the allocated capacity at arc a . The capacity allocation of the manager is the vector $\mathbf{c} = (c_1, \dots, c_A)$. An allocation \mathbf{c} is feasible if its components obey the nonnegativity and total capacity constraint, namely (i) $c_a \geq 0$, $a \in \mathcal{A}$ and (ii) $\sum_{a \in \mathcal{A}} c_a = C$. The set of all feasible capacity allocations \mathbf{c} is denoted by Γ . Finally, a system configuration is feasible if it is composed of a feasible flow configuration and a feasible capacity allocation.

Pricing. Each service class $a \in \mathcal{A}$ has a *constant* price p_a per unit traffic. Thus, the network usage price of user i is $\sum_{a \in \mathcal{A}} f_a^i p_a$. Prices may be viewed as an indirect mean for admission control [6], and (among other things) prevent flooding of the better service classes. In this paper, however, we concentrate on capacity assignment as the management tool, assuming that prices are static (or change on a slower time scale). The issue of price setting in our context is left for future work.

The *performance measures* of both the users and the manager are specified through their respective cost functions, which they wish to minimize. We denote by $J^i(\mathbf{f}, \mathbf{c})$, $i \in \mathcal{I}$ the cost function of user i , and by $J^M(\mathbf{f}, \mathbf{c})$ the cost function of

the manager. The costs of the users and the manager are related to the *congestion* level at each of the service classes. We shall use the well known *M/M/1 latency function*

$$D_a(f_a, c_a) = \begin{cases} \frac{1}{c_a - f_a} & f_a < c_a \\ \infty & \text{otherwise} \end{cases} \tag{1}$$

as a congestion measure of each service class. Here c_a is the transmission capacity measured in the same units as the flow f_a . This latency function is often used to model delay costs in communication networks [3], and provides a clear sense of link capacity.

User cost functions. Users are distinguished by their utility function $U^i(f^i)$, which quantifies their utility for shipping a total flow f^i . We make the following assumption on U^i .

Assumption 1. *For every user $i \in \mathcal{I}$, the utility function $U^i : [0, C] \rightarrow \mathbb{R}$ is increasing, bounded, concave and continuously differentiable.*

We note that utility functions with the above characteristics are commonly used within the networking pricing literature [6, 11]. We define U^i in the set $[0, C]$ only, since a total flow of $f^i > C$ cannot be accommodated by the network. We note that a user may split its flow among different service classes in order to minimize the total cost. The total cost J^i of each user i is comprised from its delay cost, its network usage price, minus its utility, namely

$$J^i(\mathbf{f}, \mathbf{c}) = \beta^i \sum_{a=1}^A f_a^i D_a(f_a, c_a) + \sum_{a=1}^A f_a^i p_a - U^i(f^i), \tag{2}$$

where $\beta^i > 0$ represents the delay sensitivity of user i , and D_a is defined in (1).

Manager cost function. The objective of the manager is to impose certain predetermined ratios between the average delays of the service classes. Taking the delay of class 1 as a reference, the ratios are described by a vector $\rho = (\rho_2, \dots, \rho_A)$, $0 < \rho_a < \infty$, where the manager’s objective is to have the delays D_1, \dots, D_A obey

$$D_a(f_a, c_a) = \rho_a D_1(f_1, c_1), \tag{3}$$

where $\rho_1 \triangleq 1$. We refer to that relation as the *fixed ratio objective*. For concreteness, we may assume that $\rho_1 \leq \rho_2 \leq \dots \leq \rho_A$, so that service classes are ordered from best to worst. Similarly, prices are expected to satisfy $p_1 \geq p_2 \geq \dots \geq p_A$, although this is not essential for our results. In functional terms, the manager’s cost function may thus be defined as

$$J^M(\mathbf{f}, \mathbf{c}) = \begin{cases} 0 & \text{if (3) holds,} \\ \infty & \text{otherwise.} \end{cases} \tag{4}$$

An alternative objective of the manager that will be considered, is to minimize a weighted sum of the delay functions, that is

$$\bar{J}^M(\mathbf{f}, \mathbf{c}) = \sum_{a \in \mathcal{A}} w_a D_a(f_a, c_a), \tag{5}$$

where $w_a > 0$, $a \in \mathcal{A}$. As we shall see, there is a close relation between the fixed ratios objective and the weighted sum objective (5).

The interaction between the manager and the users will be referred to as the *users-manager game*. A Nash Equilibrium Point (NEP) of that game is a feasible system configuration $(\tilde{\mathbf{f}}, \tilde{\mathbf{c}})$ such that

$$\begin{aligned}
 J^M(\tilde{\mathbf{f}}, \tilde{\mathbf{c}}) &= \min_{\mathbf{c} \in \Gamma} J^M(\tilde{\mathbf{f}}, \mathbf{c}), \\
 J^i(\tilde{\mathbf{f}}^i, \tilde{\mathbf{f}}^{-i}, \tilde{\mathbf{c}}) &= \min_{\mathbf{f}^i \in \mathbf{F}^i} J^i(\mathbf{f}^i, \tilde{\mathbf{f}}^{-i}, \tilde{\mathbf{c}}) \quad \forall i \in \mathcal{I},
 \end{aligned}
 \tag{6}$$

where $\tilde{\mathbf{f}}^{-i}$ stands for the flow configurations of all users, but the i th user. Namely, the NEP is a network working point, where no user, nor the manager, finds it beneficial to change its flow or capacity allocation. Our users-manager game is characterized by the finiteness of the NEP costs, since users can always ship a flow of zero to encounter a finite cost. We shall formally prove this attribute in the next section.

3 Capacity Assignment and Equilibrium Analysis

In this section we analyze the equilibrium point and suggest capacity assignment schemes that induce the ratios objective. First, we show that the manager has a unique best response with respect to the ratio objective (3). This response is a simple solution to a set of linear equations. Then we prove the existence and uniqueness of an equilibrium point, in which the desired ratios are met. Accordingly, we show the equivalence between the best response with respect to (3) and the best response with respect to the modified manager objective function (5) and discuss its implications.

Theorem 1 considers the best response capacity assignment of the manager, given any (fixed) flow configuration.

Theorem 1. *Consider a fixed flow configuration (f_1, \dots, f_A) and a desired ratio vector ρ . If $\sum_{a \in \mathcal{A}} f_a < C$, there exists a unique capacity allocation $\mathbf{c} \in \Gamma$ such that (3) is met. This allocation is explicitly given by*

$$c_a - f_a = (C - \sum_{\alpha \in \mathcal{A}} f_\alpha) \frac{\rho_a^{-1}}{\sum_{\alpha \in \mathcal{A}} \rho_\alpha^{-1}}.
 \tag{7}$$

Proof. The allocation (7) is derived from solving the following set of *linear* equations: $\rho_a(c_a - f_a) = (c_1 - f_1)$, $a = 2, \dots, A$; and $\sum_{a=1}^A c_a = C$. □

The above result allows the manager to explicitly calculate its *best response* assignment, namely the capacity assignment that will satisfy the fixed ratio objective given the *current* network flows. This calculation requires just the total flows in each service class, which are easy to measure. The next theorem establishes the existence and uniqueness of the equilibrium point.

Theorem 2. *For every delay ratios vector ρ , there exists a unique Nash equilibrium point. This NEP has finite costs for the manager and for the users. In particular, the ratios objective of the manager is satisfied.*

Proof. The proof follows from the next four lemmas.

Lemma 1. *For every delay ratio vector ρ , there exists a NEP. Furthermore, in every NEP the costs of the users and the manager are finite and the ratio objective is met.*

Proof. See [15].

Lemma 2. *Let D_1, \dots, D_A be the class delays at some NEP. Then the following equations are met at the equilibrium for every $i \in \mathcal{I}$ and every $a \in \mathcal{A}$*

$$\begin{aligned} \beta^i (D_a + f_a^i D_a^2) + p_a &= U^{i'}(f^i) \quad \text{if } f_a^i > 0, \\ \beta^i D_a + p_a &\geq U^{i'}(f^i) \quad \text{if } f_a^i = 0. \end{aligned} \tag{8}$$

Proof. Observe that

$$\begin{aligned} \frac{dJ^i(\mathbf{f}, \mathbf{c})}{df_a^i} &= \beta^i \left(\frac{1}{(c_a - f_a)} + \frac{f_a^i}{(c_a - f_a)^2} \right) + p_a - U^{i'}(f^i) \\ &= \beta^i (D_a(f_a, c_a) + f_a^i D_a^2(f_a, c_a)) + p_a - U^{i'}(f^i). \end{aligned} \tag{9}$$

Then (8) may be readily seen to be the KKT optimality conditions [5] for minimizing the cost function (2) of user i subject to the flow constraint $f_a^i \geq 0$. \square

Lemma 3. *Consider a NEP with given class delays D_1, \dots, D_A . Then the respective equilibrium flows f_a^i are uniquely determined.*

Proof. Consider the following optimization problem in (f_1^i, \dots, f_A^i) :

$$\begin{cases} \min \sum_{a=1}^A \frac{1}{2} \beta^i D_a^2 f_a^i{}^2 + f_a^i (\beta^i D_a + p_a) - U^i(\sum_a f_a^i) \\ \text{s.t.} \quad f_a^i \geq 0 \end{cases}, \tag{10}$$

where we assume that the delays D_a are fixed. Note that (10) is a strictly convex optimization problem, since the objective function is the sum of a diagonal quadratic term (with $\beta^i D_a^2 > 0$ for every a) and the negation of U^i , where U^i is concave by Assumption 1. Thus, this problem has a unique minimum, which is characterized by the KKT optimality conditions. It is now readily seen that the KKT conditions for (10) coincide with the conditions in (8). Thus, by Lemma 2, any set of equilibrium flows $(f_a^i)_{a \in \mathcal{A}}$ is a solution of (10). But since this solution is unique, the claim is established. \square

Lemma 4. *Consider two Nash equilibrium points (\mathbf{f}, \mathbf{c}) and $(\tilde{\mathbf{f}}, \tilde{\mathbf{c}})$. Then $D_a(f_a) = D_a(\tilde{f}_a)$ for every $a \in \mathcal{A}$.*

Proof. Define $D_a \triangleq D_a(f_a)$ and $\tilde{D}_a \triangleq D_a(\tilde{f}_a)$. Assume that

$$\tilde{D}_\alpha > D_\alpha \text{ for some } \alpha \in \mathcal{A}. \tag{11}$$

Then $\tilde{D}_a > D_a$ for every $a \in \mathcal{A}$ since the ratios are met in both equilibria (Lemma 1). Noting (4) and (7) it follows from (11) that $\sum_{a \in \mathcal{A}} \tilde{f}_a > \sum_{a \in \mathcal{A}} f_a$, which implies that there exists some user j for which

$$\tilde{f}^j = \sum_{a \in \mathcal{A}} \tilde{f}_a^j > \sum_{a \in \mathcal{A}} f_a^j = f^j. \quad (12)$$

We next contradict (12) by invoking the next two implications:

$$f_a^j = 0 \Rightarrow \tilde{f}_a^j = 0 \quad (13)$$

$$f_a^j > 0 \Rightarrow f_a^j > \tilde{f}_a^j. \quad (14)$$

Their proof is based on the KKT conditions (8). Since the utility U^j is concave, then by (12) we have $\lambda^j \triangleq U^{j'}(f^j) \geq U^{j'}(\tilde{f}^j) \triangleq \tilde{\lambda}^j$. If $f_a^j = 0$, then $\beta^j D_a + p_a \geq \lambda^j \geq \tilde{\lambda}^j$. Since $\tilde{D}_a > D_a$, then $\beta^j \tilde{D}_a + p_a > \tilde{\lambda}^j$, hence $\tilde{f}_a^j = 0$. To prove (14) note first that it holds trivially if $\tilde{f}_a^j = 0$. Next assume $\tilde{f}_a^j > 0$. Then by (8)

$$\beta^j D_a + \beta^j D_a^2 f_a^j + p_a = \lambda^j \geq \tilde{\lambda}^j = \beta^j \tilde{D}_a + \beta^j \tilde{D}_a^2 \tilde{f}_a^j + p_a. \quad (15)$$

Since $\tilde{D}_a > D_a$ (hence $\tilde{D}_a^2 > D_a^2$), and β^j are positive, we must have $f_a^j > \tilde{f}_a^j$ in order for (15) to hold, which establishes (14). Summing user j 's flows according to (13)-(14) yields $\sum_{a \in \mathcal{A}} \tilde{f}_a^j \leq \sum_{a \in \mathcal{A}} f_a^j$, which contradicts (12). Thus $\tilde{D}_\alpha \leq D_\alpha$. Symmetrical arguments will lead to $\tilde{D}_\alpha \geq D_\alpha$, i.e., $\tilde{D}_\alpha = D_\alpha$, hence $\tilde{D}_a = D_a$ for every $a \in \mathcal{A}$. \square

The last two lemmas imply that the user flows and the class delays in equilibrium are unique. The capacities in the equilibrium must also be unique since $c_a = f_a + \frac{1}{D_a}$, where $f_a = \sum_{i \in I} f_a^i$. This establishes the uniqueness of the NEP, and completes the proof of Theorem 2. \square

A possible criticism of the fixed ratio objective, as defined in (3), is that it does not account at all for absolute congestion measures, namely the delays themselves rather than their ratios. However, the next result shows that by achieving the ratio objective, the manager in fact minimizes the cost function \bar{J}^M (defined in (5)), which is just an appropriately weighted sum of the delays over the different service classes.

Theorem 3. *Consider the users-manager game, whose NEP is defined in (6), and an additional game, which is similar except that J^M is replaced by \bar{J}^M . If the parameters are such that $w_a = \frac{1}{\rho_a^2}$ for every $a \in \mathcal{A}$, then the two games are equivalent in the sense that their (unique) equilibrium points coincide.*

Proof. See [15].

Note that the weights w_a are inversely proportional to ρ_a^2 , which assigns higher weight to better (lower relative delay) service classes, as might be expected. From an algorithmic point of view, it should be noted that \bar{J}^M is a convex and continuous function, and therefore may provide a sound basis for an iterative (e.g., gradient-based) algorithm that may be used by the manager to minimize

this cost, with the goal of eventually reaching the desired fixed-ratio equilibrium. Yet, a specific consideration of such an algorithm is beyond the scope of this work.

We conclude this section by considering the computation of the NEP. Recall that the uniqueness of the user flows was established in Lemma 3 via the definition of strictly convex optimization problems. Hence, by solving the same optimization problems, the NEP can be efficiently calculated. The only unknown variable which is required for the computation is D_1 . In our case, an iterative search for D_1 may be easily performed, by comparing the total flow $\sum f_a$ (for a given D_1), which is obtained from both the best response formula (7) and also from the solutions to (10). More details on the search method are given in [15].

Remark 1. General Networks. It is obviously of interest to extend the results of this section to a general network topology. In a recent paper [14], we have made such an extension for a model with fixed (plastic) user demands. Since the same extension can be similarly used here, we briefly outline its key features. The following setup is applied: (i) each user has a unique fixed path from its source to its destination; (ii) the QoS ratio objective is maintained on a link basis, i.e., management is performed via a *distributed* approach, where the capacity adaptation is performed locally, at the link level. Observe that if the above two features are maintained, then the QoS ratios are met *end-to-end* for every user. The game framework now includes network *managers*, one for every link. Due to the locality of the capacity management, the formula for the best response map (Theorem 1) and the use of \bar{J}^M instead of J^M (Theorem 3) are trivially extended to the general network case; so is the proof for the existence of the equilibrium. The issue of uniqueness of the equilibrium (as in Lemma 4) is however more complicated in the general network case and is currently an open problem.

4 Conclusion

The proposed approach to QoS provisioning in Diffserv networks focuses on maintaining relative congestion measures in the different service classes. Our analysis demonstrates the feasibility of this approach, and in particular establishes the existence and uniqueness of a working point, which satisfies this QoS objective in an elastic, reactive and heterogenous user environment. Our results provide an efficient algorithm for computing the Nash equilibrium. However, from the manager's viewpoint, this computation requires complete knowledge of the users' preferences, which may not be available. An alternative scheme is to use adaptive capacity assignment, for example by utilizing the best-response map (7), which requires only the total flows in each service class, which are easy to measure. The analysis of such a scheme is an important issue for future research. Additional research topics include the price setting issue, proportional QoS with other cost functions and QoS measures, and the equilibrium dynamics of capacity allocation algorithms.

References

1. E. Altman and L. Wynter. Equilibrium, games, and pricing in transportation and telecommunications networks. *Networks and Spatial Economics*, 4(1):7–21, 2004.
2. F. Baker, K. Chan, and A. Smith. Management information base for the differentiated services architecture. RFC 3289, 2002.
3. D. P. Bertsekas and R. G. Gallager. *Data Networks*. Prentice-Hall, 1992.
4. S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An architecture for differentiated services. RFC 2475, 1998.
5. S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2003.
6. C. Courcoubetis and R. Weber. *Pricing Communication Networks: Economics, Technology and Modelling*. Wiley, 2003.
7. B. Davie, A. Charny, J. Bennett, K. Benson, J. L. Boudec, W. Courtney, S. Davari, V. Firoiu, and D. Stiliadis. An expedited forwarding PHB (per-hop behavior). RFC 3246, 2001.
8. C. Dovrolis, D. Stiliadis, and P. Ramanathan. Proportional differentiated services: Delay differentiation and packet scheduling. *IEEE/ACM Transactions on Networking*, 2002.
9. Y. Hayel, D. Ros, and B. Tuffin. Less-than-best-effort services: Pricing and scheduling. In *Proceedings of IEEE INFOCOM*, 2004.
10. J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski. Assured forwarding PHB group. RFC 2597, 1999.
11. F. P. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
12. Y. A. Korilis, A. A. Lazar, and A. Orda. Architecting noncooperative networks. *IEEE Journal on Selected Areas in Communication*, 13(7):1241–1251, 1995.
13. M. Mandjes. Pricing strategies under heterogeneous service requirements. In *Proceedings of IEEE INFOCOM*, pages 1210–1220, 2003.
14. I. Menache and N. Shimkin. Capacity assignment in Diffserv networks for proportional QoS. Technical Report CCIT-544, Technion, July 2005.
15. I. Menache and N. Shimkin. Proportional QoS in differentiated services networks: Capacity management, equilibrium analysis and elastic demands. Available from <http://www.ee.technion.ac.il/people/shimkin/preprints/MS05x.pdf>, July 2005. Extended version.
16. H. Mendelson and S. Whang. Optimal incentive-compatible priority pricing for the M/M/1 queue. *Operations Research*, 38:870–883, 1990.
17. A. M. Odlyzko. Paris metro pricing for the internet. In *Proceedings of the ACM Conference on Electronic Commerce*, pages 140–147, 1999.
18. A. Orda, R. Rom, and N. Shimkin. Competitive routing in multi-user environments. *IEEE/ACM Transactions on Networking*, 1:510–521, 1993.
19. A. K. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: the single-node case. *IEEE/ACM Transactions on Networking*, 1(3):344–357, 1993.
20. O. Richman and N. Shimkin. Topological uniqueness of the Nash equilibrium for non-cooperative routing. *Mathematics of Operations Research*, submitted for publication, October 2004.
21. T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
22. X. Xiao and L. M. Ni. Internet QoS: A big picture. *IEEE Network*, 36(11):8–18, 1999.