

Pricing, Provisioning and Peering: Dynamic Markets for Differentiated Internet Services and Implications for Network Interconnections

Nemo Semret, Raymond R.-F. Liao, Andrew T. Campbell, and Aurel A. Lazar, *Fellow, IEEE*

Abstract—This paper presents a decentralized auction-based approach to pricing of edge-allocated bandwidth in a differentiated services Internet. The players in our network economy model are one raw-capacity seller per network, one broker per service per network, and users, to play the roles of whole-sellers, retailers, and end-buyers, respectively, in a two-tier wholeseller/retailer market, which is best interpreted as a “sender-pay” model. With the progressive second price auction mechanism as the basic building block, we conduct a game theoretic analysis, deriving optimal strategies for buyers and brokers, and show the existence of networkwide market equilibria.

In addition to pricing, another key consideration in building differentiated network services is the feasibility of maintaining stable and consistent service level agreements across multiple networks where demand-driven dynamic allocations are made only at the edges. Based on the proposed game-theoretic model, we are able to construct an explicit necessary and sufficient condition for the stability of the game, which determines the sustainability of any set of service level agreement configurations between Internet service providers.

These analytical results are validated with simulations of user and broker dynamics, using the distributed progressive second price auction as the spot market mechanism in a scenario with three interconnected networks, and two services based on the proposed standard expedited forwarding and assured forwarding per-hop behaviors.

Index Terms—Capacity provisioning, differentiated service, network interconnection, peering stability, second price auction.

I. INTRODUCTION

THE RECENT development of the differentiated service (DiffServ) Internet model is aimed at supporting service differentiation for aggregated traffic in a scalable manner [1], [2]. The tenet of DiffServ is to relax the traditional hard-QoS model (e.g., end-to-end per-flow guarantee of IntServ [3], and ATM) in two dimensions: slower time-scale network mechanisms and coarser-grained traffic provisioning.

The focus of the proposed differentiated services framework has been mainly on packet level behavior, with the purpose of defining building blocks for scalable differentiated services. Substantial progress has been made in the development and

standardization of packet forwarding behaviors [4], [5]. However, two issues have been lacking systematic study in the development of differentiated services:

- 1) dynamic market-pricing of edge-allocated bandwidth; and
- 2) the feasibility of maintaining consistent service level agreements (SLAs)—or DiffServ profiles—across interconnected networks where demand-driven dynamic allocations are made only on the edges.

While the role of prices as an essential resource allocation “control signal” has been established from the outset of DiffServ [6], [7], the precise development of pricing mechanisms is still at its early stages. In the simple integrated media access model [8], the service charge for a user is proportional to the nominal subscribed bit rate and the price differentiation between different service classes remains fixed. Similarly, in the user-share differentiation proposal [9], pricing is based on the user share that is allocated over long time scales. These schemes fall within the category of capacity-based pricing. Just as DiffServ aims to provide a range of “better than best-effort” services without the complexity and per-flow state of hard-QoS, capacity-based pricing schemes can be thought of as “better than flat-rates” (more rational and sustainable from the economic point of view), without the continuous measurement and accounting required by usage-based pricing. Flat-rate pricing is the extreme of capacity pricing where the capacity equals the access line speed, while usage pricing can be thought of as the extreme where capacities are continuously adapted to fit the actual transmission rate of each flow at each moment in time. A pricing scheme which explicitly covers the range between these two, as well as the service-type dimension is discussed in [10].

One consequence of resource allocation at network edges is a natural proclivity toward a “sender-pay” model. Indeed, a “receiver-pay” model would require explicit price signaling back to the source in order to allocate the corresponding resources, since prices have to relate to the resources consumed (i.e., service quality). Such signaling, if done in real time within the network, would reintroduce the same type of complexity and scalability problems as those that afflict end-to-end per-flow QoS, and that the edge-allocation model is meant to avoid.¹

A sender-pay model is a departure from the Internet tradition of receiver-paid flat rates. However, while the pricing mechanisms presented in this paper can equally apply to a receiverpay

¹Of course, the receiver may pay the sender through some off-line means, e.g., through subscription, “pay-per-view,” or indirectly in the case of advertising-supported content.

Manuscript received October 15, 1999; revised April 15, 2000.
N. Semret is with Invisible Hand Networks, Inc., New York, NY 10012 USA (e-mail: nemo@invisiblehand.net).

R. R.-F. Liao, A. T. Campbell, and A. A. Lazar are with the Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: liao@comet.columbia.edu; campbell@comet.columbia.edu; aurel@comet.columbia.edu).

Publisher Item Identifier S 0733-8716(00)09223-4.

model; there is a strong case to be made that the Internet has reached a stage in its evolution where the change is due. Indeed, consider the history of postal service: in ancient times, it was generally run on a receiver-pay model. In a system with unreliable delivery, it is more natural to require payment on the receiving side. Just like the best-effort Internet, the unreliability was compensated for by the fact that the system was lightly loaded, and messages were such that retransmissions were acceptable. As the number of users grew, the postal system went through a phase of complex bilateral agreements between countries (this occurred in Europe from about 1600–1900), much like inter-ISP peering today. In the later stage, where differentiated services are offered (e.g., air-mail, overnight express, bulk-mail), the default is for the sender to pay,² since the quality must be selected on the sending side. Thus, by analogy, the move from best-effort to differentiated services should lead to a sender-pay model.

The space of network resource pricing schemes has many dimensions (for a complete taxonomy of network pricing, see [11, Chapter 1]). One is “where” the capacity abstraction takes place: at each hop inside the network or at the edges [12] (as discussed above). Another is how much *a priori* information on demand is required. At one extreme, the seller assumes perfect *a priori* knowledge of demand and does an offline calculation of optimal prices (e.g., time-of-day pricing based on historical traffic patterns). In more sophisticated approaches, the seller assumes the functional form of demand and adjusts prices by on-line optimizations [13]–[17]. These pricing schemes are “model-based,” in that the relationship between demand and price (and possibly time) is assumed in an *a priori* formula. Knowledge of this model and its parameters is precisely the *a priori* information requirement described above.

Auctioning is the pricing approach with minimal information requirement. The more difficult it is for the seller to obtain demand information (or valuations), the stronger the case is for using auctions. In today’s Internet, because of the diverse and rapidly evolving nature of the applications, services, and population, the case is particularly compelling. With suitably designed rules, auctions can achieve efficient (i.e., value maximizing) allocations with minimal *a priori* information.

An important aspect of the problem that has not been systematically addressed is the feasibility of maintaining consistent SLAs across interconnected networks with dynamic, market driven, edge capacity allocation. Inconsistent SLAs would result in frequent reconfiguration of traffic conditioners at the edges, and/or significant violations of the service quality in the core of the networks.

In this paper, we investigate two closely coupled problems. First, on the “demand side,” we study the feasibility of auctioning capacity in real-time on a DiffServ internet. We then consider the “supply side,” focusing on the feasibility of provisioning *stable* and *consistent* SLAs across multiple networks, where allocations are dynamically driven by demand and made only on the edges.

We begin in Section II by constructing the two-tier whole-seller/retailer market model, giving the wide-area model for

pricing, provisioning, and differentiation of the services, and introduce the demand model.

Following this, in Section III, we show through game-theoretic analysis and simulation that the progressive second price (PSP) auction of [18] can provide stable and efficient pricing in a DiffServ bandwidth market. The results of this section extend those of the single sharable resource auction of [18] to the case of multiple networked resources, in an edge-capacity allocation framework. The PSP mechanism achieves the economic objectives of incentive compatibility and efficiency, while being realistic in the engineering sense (small signaling load and computationally simple allocation rule). As such, it provides a useful baseline for understanding the conditions for the economic feasibility of wide-area differentiated services.

In Section IV, we derive a necessary and sufficient condition for the stability of dynamic SLA provisioning. Then, in Section V, all the analytical results are validated by simulations, which illustrate not only conditions for stable and unstable markets, but also stable conditions which lead to certain classes of service not being offered on an internetwork basis. Finally, in Section VI, we present some concluding remarks and future work.

II. THE MODEL

A. Distributed Market Framework

Our network model assumes that each network can be abstracted into a single bottleneck capacity (e.g., as a “Norton-equivalent” [19]). The capacity may be represented by an absolute amount of bandwidth, or some relative metrics like user share in the user-share differentiation proposal [9] or resource token in location independent resource accounting [20]. Large networks can be modeled by subdivision into a set of interconnected networks, each of which can be abstracted into a bottleneck capacity. The degree of subdivision that is necessary depends on traffic, topology, and size constraints as well as the desired level of accuracy. Within each network, the routing of aggregated traffic to each peer³ is stable over the resource allocation time scale (e.g., in the order of hours).

Fig. 1 presents the model of our proposed auction pricing framework for a set of interconnected networks as described above. A two-tier whole-seller/retailer market model is used to accommodate a network of goods (i.e., bandwidth) with multiple differentiated service classes. We define three kinds of players: users, service bandwidth brokers (SBBs), and raw bandwidth sellers (RBSs), to play the roles of end-users, retailers, and whole-sellers, respectively. Each network has a single RBS and a separate SBB for each class of service being offered. The RBS can be thought of as the bearer, and the SBBs as service providers [21]. If the RBS and multiple SBBs on the same network are not owned by the same entity, a noncooperative game formulation is the best way to model the problem. Even if they are owned by the same entity, a competitive framework is valuable, the idea being that competition among SBBs results in a dynamic and efficient partition of the physical network resources among the services being offered,

²At least for the part that relates to service quality differentiation. In general, all parties pay for basic connectivity to the system.

³In this paper, we use the term “peer” in the most general sense, i.e., any network which interconnects with a given network, and not just those that choose to exchange all traffic free of charge.

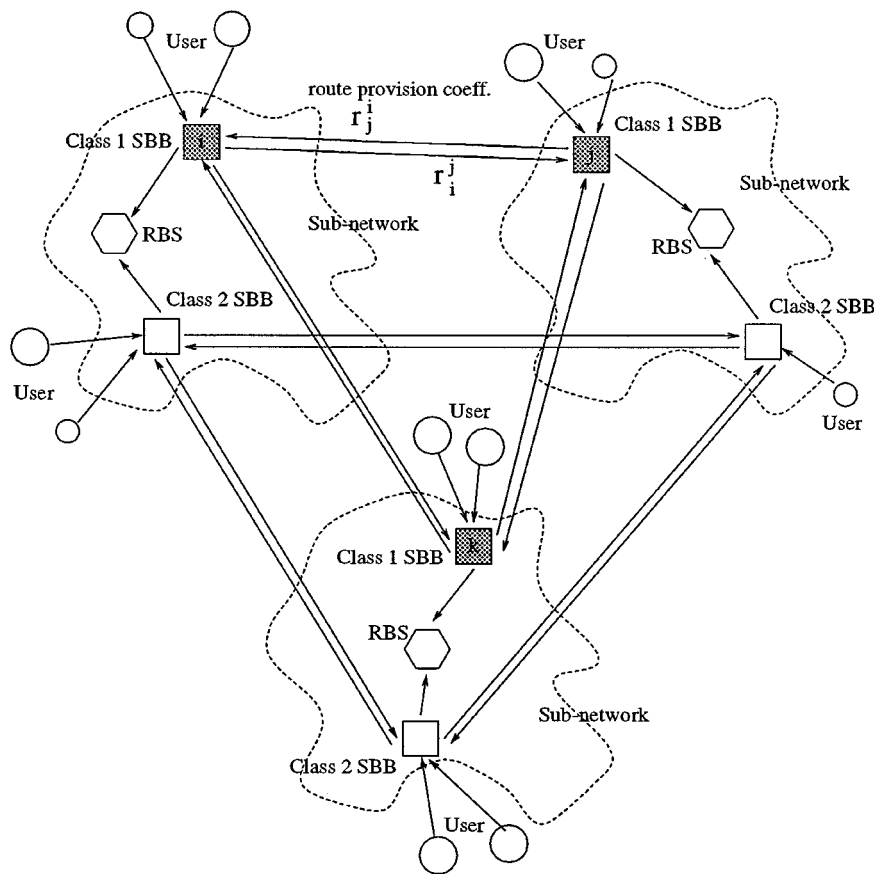


Fig. 1. The 2-tier auction pricing framework for DiffServ internet.

based on the demands from users. The users, or retail buyers, are subscribers to a particular service offered by a particular provider. In the DiffServ context, these will likely be large subscribers (e.g., web sites, various content or application server farms, intra/extranets, virtual private networks), rather than individual end users.

B. Game Theoretic Model: Message Process and Notation

Let the set of all players, including buyers, sellers, and brokers (brokers are both buyers and sellers), be denoted by $\mathcal{I} = \{1, \dots, I\}$. A player's identity $i \in \mathcal{I}$ as a subscript indicates that the player is a buyer, and as a superscript indicates the seller.

Suppose player i is buying from player j . Then he/she places a **bid** $s_i^j = (q_i^j, p_i^j)$, meaning he/she would like to buy from j a quantity q_i^j and is willing to pay a *unit* price p_i^j . Without loss of generality, we assume that all players bid in all auctions, with the understanding that if a player i does not need to buy from j , we simply set $s_i^j = (0, 0)$.

A seller j places an **ask** $s_j^j = (q_j^j, p_j^j)$, meaning he/she is offering a quantity q_j^j , with a reserve (or floor) price of p_j^j per unit. In other words, when the subscript and superscript are the same, the bid is understood as an ask.

Unless otherwise indicated, when sub/superscripts are omitted, the notation refers to the vector obtained by letting it range over all values. For example, q_i is the $1 \times I$ vector (q_i^1, \dots, q_i^I) , and q is the $I \times I$ matrix. A subscript with a minus sign indicates a vector with that component deleted

$s_{-i} \equiv (s_1, \dots, s_{i-1}, s_{i+1}, \dots, s_I)$, and $(x_i; s_{-i})$ denotes the profile obtained by replacing s_i with x_i .

Based on the profile of bids $s^j = (s_1^j, \dots, s_I^j)$, seller j computes an allocation $(a^j, c^j) = A^j(s^j)$, where a_i^j is the quantity given to player i and c_i^j is the *total cost* charged to the player i . A^j is the **allocation rule** of seller j . It is feasible if $a_i^j \leq q_i^j$, and $c_i^j \leq p_i^j q_i^j$. One possible allocation rule is the progressive second price auction as discussed in Section III.

C. Sellers' Provisioning and Peering Constraints

Suppose player $k \in \mathcal{I}$ is an RBS. Then its strategy consists of always *asking* $s_k^k = (q_k^k, p_k^k)$, with q_k^k equal to the physical bottleneck capacity of its network, and p_k^k equal to the unit cost of operation. Since it is a passive seller of physical bandwidth, k does not buy from anyone, i.e., $s_k^j = 0, \forall j \neq k$.

Suppose $j \in \mathcal{I}$ is an SBB. It offers a capacity q_j^j for sale to its users. In order to honor its contracts, the quantity offered must be constrained by the capacities that j can actually obtain. First, it must get enough bandwidth from k , the RBS in its own network, to carry the total capacity it allocates to its customers, i.e.,

$$\sum_i a_i^j \leq a_j^k. \quad (1)$$

Second, since it is selling interconnection service, j must get enough capacity from the SBBs offering the same service in each peer network. Let l denote one such peer SBB, and r_l^j be

the “fraction of traffic” generated by j 's customers that is routed to the network where player l is the peer SBB (see Remark in Section II-D for interpretations of r_j). Then, j must satisfy

$$r_j^l \sum_{i \neq l} a_i^j \leq a_j^l \quad (2)$$

for all peers l .⁴ For notational convenience, fix $r_j^k = 1$, when k is j 's RBS. Since $a_k^j = 0$, (2) includes (1) as the special case $l = k$. If l is neither a peer of j , nor its RBS, then we set $r_j^l = 0$.

Define, for any allocation a

$$e_j^l(a) \triangleq \frac{a_j^l}{r_j^l} + a_j^j.$$

We call

$$e_j \triangleq \min_{l \neq j} e_j^l(a) \quad (3)$$

the *expected bottleneck capacity* for the service offered by j .

Proposition 1 (Broker's sell-side constraints): Let $j \in \mathcal{I}$ be a SBB, and fix its buy-side allocation (a_j, c_j) . Then, on the sell-side, the quantity offered must satisfy

$$q_j^j \leq \min_{l \neq j} e_j^l(a).$$

For a broker who does not sell at a loss, the reserve price must satisfy

$$p_j^j \geq \frac{1}{q_j^j} \sum_l c_j^l.$$

Proof: Suppose $\exists l \neq j$ such that $q_j^j > e_j^l$. Then when all the offered quantity is bought, we have $\sum_i a_i^j = q_j^j > e_j^l = (a_j^l/r_j^l) + a_j^j \Leftrightarrow \sum_{i \neq l} a_i^j > a_j^l/r_j^l$, and condition (2) is violated. This proves the first assertion.

Since $\sum_l c_j^l$ is the total cost of the capacity that j is buying, the second assertion follows immediately from our assumption that the broker will not sell at a loss. \square

Remark: The obvious way for a broker to satisfy Proposition 1 is simply setting $q_j^j = \min_{l \neq j} e_j^l(a)$. Alternately, the seller can leave q_j^j equal to the maximum physical capacity, and place in its own market an artificial “buy-back” bid equal to $s_0^j = (q_0^j, p_0^j)$, where $q_0^j = (q_j^j - e_j)^+$ and p_0^j is larger than any user is willing to bid. Note that this artificial player $0 \notin \mathcal{I}$. This buy-back bid effectively limits j 's users to precisely the capacity that j can honor in forward to its peers. In other words, the buy-back bid ensures that the quantity constraint of Proposition 1 is automatically satisfied. If there is demand (bids) at prices greater than the marginal cost to j of expanding capacity, then naturally broker j will want to satisfy it, so p_0^j should be set at the marginal cost of increasing the offered quantity e_j . As we will become apparent through Proposition 4 below, p_0^j should be set to equal $\theta_j^j(e)$, which is the price at which j could obtain more capacity at its bottleneck to a peer network.

⁴We assume that service providers block “loop-back” traffic, i.e., traffic going from l through j and back to l . If that is not the case, then the summation in (2) would be over all i .




Fig. 2. Internetwork provisioning coefficients for Olympic Gold, Silver, and Bronze services, and the virtual leased line service.

D. Differentiating Services

We do not explicitly consider the per-hop behaviors (PHB's) per se, which of course are essential in assuring the service quality on the packet time scale. On our level of abstraction, only the vector of provisioning coefficients r_i differentiates broker i and the service it offers. A broker is characterized by the type of SLA that it offers.

- *Expected capacity SLA.* On average, users will get the capacity they pay for, even when the traffic enters peer networks. This could include for example services built on the DiffServ assured forwarding (AF) per-hop behaviors [5]. In this case, r_i^j is the expected fraction of the total traffic entering i that is routed to j . r_i^i the fraction of traffic that terminates with one of i 's own customers, and $\sum_{j \neq i} r_i^j = 1$, where l is the RBS in i 's network.⁵
- *Worst-case capacity SLA.* Another type of SBB may offer service agreements for worst-case bandwidth, i.e., each user always gets the amount of bandwidth they pay for, even if all of the traffic is routed to the same peer. This could include for example services built on the DiffServ expedited forwarding (EF) per-hop behavior [4]. In this case $r_i^j = 1$ for all peers j .
- *Local SLA.* For an SBB which offers SLAs valid only within its own network, $r_i^i = 1$ and $r_i^j = 0, \forall j \neq i$.

Fig. 2 illustrates several service scenarios for an SBB i with two peers j and k . In all the cases, the steady-state aggregate traffic pattern is such that 2/3 of i 's traffic flows to j 's network, and 1/3 flows to k 's network (to visualize in only two dimensions, we assume $r_i^i = 0$, i.e., i provides only “transit” service, so no traffic terminates within i 's own network). Thus, if i is offering an expected capacity service, r_i will lie along the line with slope 1/2. Here we show how the SBB would have to provision the three classes in the “Olympic service” based on AF [5], and the “virtual leased line” (VLL) service based on EF [4]. Degrees of overprovisioning must be used to differentiate among AF classes. A Bronze service class SBB would provision just enough capacity to carry the traffic on average (circle marked

⁵Note that for expected capacity, a user m whose traffic is entirely within the allocated profile a_m^i when it enters its broker i 's network could temporarily be out of profile in the peer network j , if i miscalculated r_i^j , or if there is a sudden surge of traffic from many of i 's customers to j .

“B” in the figure). If the SBB is providing Silver class service, then it must provision more generously to ensure that they are less loaded, and thus experience better service, and even more generously if the service is Gold class (circles marked “S” and “G” in the figure). For the VLL service, more conservative provisioning can be achieved by providing for worst-case flows, i.e., all the traffic can flow to any one peer and still be satisfied, as illustrated by “V” in Fig. 2.

Depending on the scheduling and buffer management algorithms used to provide the PHBs, some amount of overprovisioning may be required [4]. These engineering needs can be represented in this model by simply factoring overprovisioning into each coefficient of r , e.g., if i is offering a virtual leased line with 5% overprovisioning then $r_i = (1.05, 1.05, \dots)$.

Note that for our purposes, the provisioning coefficients r_i are known by broker i in advance, since they represent aggregate flow patterns. In practice, this means r would be measured over a time-scale slow enough to make quasi-static estimates which average out microflows.

E. Buyers

We model buyers as **bottleneck buyers**, i.e., each buyer $i \in \mathcal{I}$ seeks to maximize its utility

$$u_i = \theta_i \circ e_i(a) - \sum_j c_i^j \quad (4)$$

where

- e_i is as in (3),
- θ_i is the buyer’s **valuation** function, and
- \circ denotes composition of functions [i.e. $\theta_i \circ e_i(a) = \theta_i(e_i(a))$].

As the name indicates, the valuation function describes how much each possible allocated quantity is worth to the buyer, i.e., the willingness to pay, and is private information. Other players (including the seller) only see the buyer’s bid and not the valuation that lead the buyer to make that bid. Here, the valuation depends only on a scalar bottleneck $e_i(a)$ which is a function of the allocated quantities at all the resources.

If the buyer is a user i buying from SBB j , then $r_i^j = 1$ and $r_i^l = 0, \forall l \neq j$. Thus, $e_i(a) = a_i^j$, and (4) has the simpler form $u_i = \theta_i(a_i^j) - c_i^j$. The valuation is a function of the player’s own allocation only, and expresses the amount the user is willing to pay for each possible quantity of resource. It can be based on economic and/or information theoretic considerations (see [18, appendix]).

If the buyer is a broker, the natural utility is the potential profit so θ_j , the broker’s buy-side valuation, is the potential revenue from the sale (on the sell-side) of the capacities obtained on the buy-side. The potential revenue is derived from the demand on the sell-side: let $\forall y > 0$,

$$d^j(y) \triangleq \sum_{r_k^j \geq y} q_k^j$$

the demand at unit price y . Its “inverse” function is defined by

$$f^j(z) \triangleq \sup\{y \geq 0: d^j(y) \geq z\}.$$




Fig. 3. Demand curve for a broker j .

See Fig. 3. Note that we chose f^j to be continuous from the left. For a given demand function $d^j(\cdot), \forall z \geq 0$, $f^j(z)$ represents the highest unit price at which j could sell the z th unit of capacity. The actual prices charged to users depend on the specific allocation mechanism A^j used.

Proposition 2 (Broker’s buy-side valuation): Let $j \in \mathcal{I}$ be a broker with inverse demand $f^j(z)$. Its buy-side valuation is

$$\theta_j(x) = \int_0^x f^j(z) dz.$$

Thus $\theta_j \circ e_j(a) = \int_0^{e_j(a)} f^j(z) dz$.

Proof: Since the broker seeks to maximize profit, for a given allocation a , it will sell as much as possible; thus by Proposition 1, $q_j^j = e_j$. If e_j decreases by δ , then q_j^j must be reduced by δ . The value to j of the lost quantity is the revenue j could have gotten from it. By definition, this lost potential revenue is $f^j(e_j^m)\delta$. Thus, by abuse notation, writing θ_j as a function of e_j^l ,

$$\theta_j(e_j) - \theta_j(e_j - \delta) = f^j(e_j)\delta$$

and the result follows as $\delta \rightarrow 0$. \square

It is useful to conceptually decouple the game into two. On one hand is a “demand game” wherein users and brokers compete for the available bottleneck capacities. On the other hand, we have what may be called the “supply game” among brokers which results in the setting of the bottleneck capacities. Since the brokers are driven by the users’ demands, and the users are competing for the offerings of the brokers, the two games are interdependent, and may be played on the same or vastly different time scales.

The notation used in this paper is summarized in Table I.

III. DEMAND SIDE

In this section we consider the demand side, and derive the optimal (utility-maximizing) bidding strategies for users and brokers, and establish the existence of an efficient (value maximizing) equilibrium point among buyers, when sellers are static (i.e., do not change the offered quantity). We assume that each RBS imposes a nonzero asking (or “reserve”) price—which can be arbitrarily small. Thus, prices will always have a strictly positive floor.

TABLE I
SUMMARY OF NOTATIONS

q_i^j	player i 's bid quantity for bandwidth offered by seller j when ($i \neq j$)
q_i^i	quantity of i 's offered bandwidth
p_i^j	player i 's bid unit-price for bandwidth offered by seller j when ($i \neq j$)
p_i^i	reserve price of i 's offered bandwidth
$s_i^j = (q_i^j, p_i^j)$	player i 's bid for bandwidth offered by seller j when ($i \neq j$)
$s_i^i = (q_i^i, p_i^i)$	player i 's ask
s^j	profile of all bidders at seller j , $s^j = (s_1^j, \dots, s_l^j)$
(x_i^j, s_{-i}^j)	replacing the i th player's bid s_i^j with x_i^j , $(x_i^j, s_{-i}^j) = (s_1^j, \dots, s_{i-1}^j, x_i^j, s_{i+1}^j, \dots, s_l^j)$
a_i^j	allocation given to player i by seller j
c_i^j	total cost charged to player i by seller j
θ_i	player i 's valuation function
u_i	player i 's utility function, $u_i = \theta_i - \sum_j c_i^j$
e_j	expected bottleneck capacity at seller j
r_j^l	fraction of incoming traffic at j (excluding loop-back) that is routed to l

The design of our progressive second price auction (PSP) appears in [18].⁶ The mechanism is defined by: $\forall i, j \in \mathcal{I}$

$$\alpha_i^j(s) \equiv a_i^j(s^j) = q_i^j \wedge \left[q_i^j - \sum_{p_k^j \geq p_i^j, k \neq i} q_k^j \right]^+ \quad (5)$$

$$c_i^j(s) \equiv c_i^j(s^j) = \sum_{k \neq i} p_k^j \left[a_k^j(0; s_{-i}^j) - a_k^j(s_i^j; s_{-i}^j) \right] \quad (6)$$

where \wedge means taking the minimum. Note that each seller computes allocations from local information only (the bids for that resource). Define

$$P_i^j(z) \triangleq \inf \left\{ y \geq 0: q_i^j - \sum_{p_k^j > y, k \neq i} q_k^j \geq z \right\}.$$

Note that we define P_i^j to be continuous from the left. Under PSP, P_i^j is the market price function from the point of view of user i . Indeed, it can be shown that

$$c_i^j = \int_0^{\alpha_i^j} P_i^j(z) dz. \quad (7)$$

Remark: Except at points of discontinuity, we have $P_i^j(z) = f^j(q_i^j - z)$. This mechanism generalizes Vickrey (“second-price”) auctions [22] which are for nondivisible objects. PSP bears some similarity to Clarke–Groves mechanisms [23], [24]. The fundamental difference from the latter is that PSP is designed with a message (bid) space of two dimensions (price and quantity) in which each message is a single point, rather than an infinite dimensional space of valuation functions where each message is a revelation of the whole valuation

⁶PSP was first presented at the DIMACS Workshop on Economics, Game Theory, and the Internet, Rutgers, NJ, April 1997; and a generalized analysis at the 8th International Symposium on Dynamic Games and Applications., Maastricht, The Netherlands, July 1998.

curve (see [25] and [18] for an explanation of the “revelation principle”). This reduction of the message space is crucial in the context of communication networks, where limiting the size and complexity of the exchanged messages (signaling) is very important.

We define *elastic demand* as follows. $\forall i$, θ_i is continuous, concave, and smooth (θ_i' is continuous); and for some (possibly infinite) maximum capacity $\bar{a}_i \leq \infty$, θ_i' is strictly decreasing (i.e., $\theta_i'' \leq 0$ if θ_i' is well defined) on $[0, \bar{a}_i]$, and nonincreasing ($\theta_i'' < 0$) on $[\bar{a}_i, \infty)$.

Under elastic demand, analyzed as a complete information game, the PSP auction for a single arbitrarily divisible resource (e.g., bandwidth on one link in a network) has the following properties which are proven in [18].

- Incentive compatible: truth-telling (setting the bid price equal to the marginal valuation) is a dominant strategy.
- Stable: it has a “truthful” ϵ -Nash equilibrium [26], for any positive seller reserve price.
- Efficient: at equilibrium, allocations maximize total user value (social welfare) to within $O(\sqrt{\epsilon})$.
- Enables a direct tradeoff between engineering and economic efficiency (measured respectively by convergence time and total user value), by the parameter ϵ , which has a natural interpretation as a bid fee.

In the rest of this paper, we assume all the sellers in the network are using PSP as the allocation mechanism.

For users, the best strategy consists simply of bidding for the largest quantity such that the marginal valuation is higher than the market price, and setting the bid price equal to the marginal valuation (i.e., “truth-telling” is optimal).

Proposition 3 (User's strategy): Let $i \in \mathcal{I}$ be a user such that θ_i that is differentiable and θ_i' continuous from the left. Let $l \in \mathcal{I}$ be that user's broker. For a fixed profile s_{-i}^l , an ϵ -best reply for player i is $t_i^l = (v_i^l, w_i^l)$, such that

$$v_i^l = \sup \left\{ z \geq 0: \theta_i'(z) > P_i^l(z) \text{ and } \int_0^z P_i^l(\eta) d\eta \leq b_i \right\} - \epsilon/\theta_i'(0)$$

and

$$w_i^l = \theta_i'(v_i^l).$$

That is, $\forall s_i^l, u_i(t_i^l; s_{-i}^l) \leq u_i(s_i^l; s_{-i}^l) - \epsilon$.

Proof: This is a special case of Proposition 4, with $\forall l$, $a_i^l = 0$, $r_i^l = 1$, and $\forall l \neq i$, $r_i^l = 0$. This was derived separately in [18]. \square

Consider now a broker, participating in many auctions simultaneously. By the nature of its valuation (Proposition 2), capacity allocations are valuable to the broker only insofar as they increase its expected bottleneck capacity $\min_{l \neq i} e_i^l$. Thus, a broker must coordinate its buy-side bids (one submitted to each of its peers and its RBS) to maximize its overall utility.

Note that for Proposition 3, we do not require that θ_i be smooth. Concavity and nonincreasingness suffice, along with the purely technical condition of continuity from the left. These are satisfied by the broker's valuation (Proposition 2). Thus, we can expect that the same principle (optimality of truth-telling) should hold.

Indeed, as we will now show, it turns out that the optimal strategy is very similar to that of a single user. But instead of searching directly for the optimal capacity, the broker finds the optimal expected bottleneck e , which is the largest one such that the marginal value is just greater than the market price. The role of the market price is played by the average of the market prices at the different auctions, weighted by the route provisioning factors. The actual bids are obtained by transforming the desired optimal expected bottleneck e back into the corresponding quantities v_i^l to bid at each buy-side market. As with a user, truth-telling is optimal for the broker, so at each buy-side market, the broker sets the bid price to the marginal value.

Proposition 4 (Broker's buy-side strategy): Let $i \in \mathcal{I}$ be a broker, and fix all the other players' bids s_{-i} , as well as the broker's sell-side s_i^j (thus a_i^j is fixed). Let

$$e = \sup \left\{ h \geq 0: f^i(h) > \sum_{l \neq i} P_i^l((h - a_i^l)r_i^l) r_i^l \right\} - \epsilon / f^i(0) \quad (8)$$

and for each $l \neq i$

$$v_i^l = (e - a_i^l)r_i^l,$$

and

$$w_i^l = \frac{1}{r_i^l} f^i(e).$$

Then a (coordinated) ϵ -best reply for the broker is $t_i = (v_i, w_i)$, i.e., $\forall s_i, u_i(t_i; s_{-i}) \geq u_i(s_i; s_{-i}) - \epsilon$.

Proof: Since f_i is nonincreasing and $\forall l, P_i^l$ is nondecreasing, (8) implies $f^i(e) > \sum_{l \neq i} P_i^l(v_i^l)r_i^l$, and therefore $\forall l \neq i$,

$$\begin{aligned} & w_i^l > P_i^l(v_i^l) = f^l(q_i^l - v_i^l) \\ \Rightarrow & v_i^l \leq q_i^l - d^l(w_i^l) \\ \Rightarrow & a_i^l(t_i; s_{-i}) = v_i^l \\ \Rightarrow & e_i^l \circ a(t_i; s_{-i}) = e. \end{aligned}$$

Therefore,

$$\begin{aligned} u_i(t_i; s_{-i}) &= \int_0^e f_i(\eta) d\eta - \sum_{l \neq i} \int_0^{v_i^l} P_i^l dz \\ &= \int_0^e f^i(\eta) d\eta - \sum_{l \neq i} \int_{a_i^l}^e r_i^l P_i^l((\eta - a_i^l)r_i^l) d\eta. \end{aligned}$$

Now suppose $\exists s_i = (q_i, p_i)$ such that $u_i(s_i; s_{-i}) > u_i(t_i; s_{-i}) + \epsilon$. Let $\xi = \min_{k \neq i} e_i^k \circ a(s)$, and $\forall l \neq i$, $\zeta_i^l = [\xi - a_i^l]r_i^l$ and $\sigma_i = (\zeta_i, p_i)$. From (17) in Lemma 1, $a_i^j(\sigma_i; s_{-i}) = \zeta_i^j$; therefore,

$$u_i(\sigma_i; s_{-i}) = \int_0^\xi f^i(\eta) d\eta - \sum_{l \neq i} \int_{a_i^l}^\xi r_i^l P_i^l((\eta - a_i^l)r_i^l) d\eta.$$

By Lemma 1 (given in the Appendix), $u_i(\sigma_i; s_{-i}) \geq u_i(s_i; s_{-i})$. Therefore, $u_i(\sigma_i; s_{-i}) > u_i(t_i; s_{-i}) + \epsilon$, which by Proposition 2, is equivalent to

$$\int_e^\xi f^i(\eta) d\eta - \sum_{j \neq i} \int_e^\xi r_i^j P_i^j((\eta - a_i^j)r_i^j) d\eta > \epsilon. \quad (9)$$

Let $\bar{e} = e + \epsilon / f_i(0)$. Since f^i is nonincreasing $\int_e^{\bar{e}} f^i(\eta) \leq f^i(0)(\bar{e} - e) = \epsilon$. That, along with the fact that P_i^j is nonnegative, and (9), implies

$$\int_{\bar{e}}^\xi f^i(\eta) d\eta - \sum_{j \neq i} \int_{\bar{e}}^\xi r_i^j P_i^j((\eta - a_i^j)r_i^j) d\eta > 0.$$

If $\xi > \bar{e}$, then for some $\delta > 0$, $f^i(\bar{e} + \delta) > \sum_{j \neq i} r_i^j P_i^j((\bar{e} + \delta - a_i^j)r_i^j)$, which contradicts (8).

If $\xi \leq \bar{e}$, then $f^i(\bar{e}) < \sum_{j \neq i} r_i^j P_i^j((\bar{e} - a_i^j)r_i^j)$. But, since both f^i and P_i^j are continuous from the left, (8) implies that $f^i(\bar{e}) \geq \sum_{j \neq i} r_i^j P_i^j((\bar{e} - a_i^j)r_i^j)$, which is a contradiction. \square

As stated above, for stability of PSP, we assume that demand is elastic for all players. However, the broker does not satisfy the smoothness (continuous derivative) condition. From Proposition 2, the broker's valuation, as a function of the (scalar) expected bottleneck capacity $\min_{l \neq i} e_i^l$, is piecewise linear and concave (the derivative is the "staircase" function shown in Fig. 3). Thus, we need to assume that brokers apply some smoothing in deriving the buy-side valuation from the sell-side demand, e.g., by fitting a smooth concave curve to the piecewise linear one.

Unlike the proof of the the broker strategy, the proofs of the following results are not essential to intuitive understanding of the game and are omitted due to space constraints.

Proposition 5 (Equilibrium): In a game consisting of arbitrarily networked PSP auctions, where all buyers have utilities of the form (4), and sellers are static (i.e., with fixed s_i^j and reserve prices $p_i^j > 0$, for all sellers $i \in \mathcal{I}$), under elastic demand, for any $\epsilon > 0$, there exists a (truthful) networkwide ϵ -Nash equilibrium.

Proof: See [11, Chapter 3]. \square

At such equilibria, the allocations are efficient (i.e., arbitrarily close to the value-maximizing allocations).

Proposition 6 (Efficiency): Let a^* be the equilibrium allocations. Under elastic demand, if in addition $\forall i \in \mathcal{I}$, if θ_i^l exists and for some $\kappa > 0$, $\theta_i^l \geq -\kappa$,

$$\max_{\mathcal{A}} \sum_i \theta_i \circ e_i(a) - \sum_i \theta_i \circ e_i(a^*) = O(\epsilon/\delta + \kappa\delta)$$

where $\mathcal{A} = \{a \in \prod_j [0, q_j^j]: \sum_i a_i^j \leq q_j^j\}$, for any $\delta \leq \min_i \{e_i(a^*): e_i(a^*) > 0\}$.

Proof: See [11, Chapter 3]. \square

The bound $\epsilon/\delta + \kappa\delta$ is minimized when $\delta = \sqrt{\epsilon/\kappa}$. Thus, the strongest statement that can be made here is that as long as $\min_i \{e_i(a^*): e_i(a^*) > 0\} > \sqrt{\epsilon/\kappa}$, we get an inefficiency which is $O(\sqrt{\epsilon/\kappa})$.

In a dynamic auction game, $\epsilon > 0$ can be interpreted as a *bid fee* paid by a bidder each time they submit a bid. Indeed, in Propositions 3 and 4, the user will send a best reply bid as long as it improves his/her current utility by ϵ , and the game can only end at an ϵ -Nash equilibrium.

IV. SUPPLY SIDE

The interaction between brokers has a much richer dynamic than discussed in the previous section. For example, not all configurations of provisioning coefficients in the wide area network lead to convergence and stable allocations. Depending on

the topology and degree of overprovisioning, the interaction between brokers can lead to oscillating allocations. On the other hand, stable operating points may lead to zero allocations for some brokers resulting in certain classes of service not being offered at all. These are not mere artifacts of PSP or any particular pricing mechanism but are fundamental issues of peering and provisioning under edge-capacity allocation. The former case is analytically related to classical problems such as route-flapping using decentralized routing algorithms. The latter case relates to empirical evidence in the best-effort Internet where market forces abandon traditional “free-for-all” peering between networks of unequal size.

We now consider the supply game among brokers by itself. For that purpose, the specifics of the auction mechanism and the resulting prices are not needed. Indeed, *the analytical results presented here on the stability and sustainability of peering are independent of the actual pricing mechanism used.* It suffices to know that a broker i 's strategy results in buying capacities a_i^j from each of its peers j and offering a quantity e_i for sale according to (3), where a_i^j 's are chosen to maximize its profit; for details see [27]. We will then use simulations using PSP auctions to verify that our insights are valid when the two games are coupled.

Define the vector $e = (e_1, \dots, e_i, \dots, e_N)$ for any profile of allocations a , where e_i is the bottleneck capacity of seller i as given by (3), and $\{1, \dots, N\}$ is the subset of \mathcal{I} consisting of all the sellers (RBS' and SBBs). Pure buyers (users) are assumed to be players numbered $m = N + 1, N + 2, \dots$. From (2) and (3), at the equilibrium point, the following conditions will hold for $l \leq i \leq N$:

$$e_i = \sum_{j \in \mathcal{I}, j \neq i} a_j^i \quad (10)$$

$$a_i^j = (e_i - a_j^i) r_i^j. \quad (11)$$

Together, these equations merely state that at equilibrium, seller i will not sell more than its bottleneck capacity, and that it will not buy more than necessary from any of its peers.

The left-hand side of (10), e_i , is quantity that seller i is offering to its users given what it has obtained on the buy-side, while the right-hand side is the quantity that is actually being bought from i on its sell-side. Thus, the right-hand side can never be greater. If the left-hand side is greater, then i is buying more capacity than it can sell, which means it is wasting money (since prices are always strictly positive), and therefore will reduce some of its bids on the buy-side. Thus, an equilibrium can occur only when equality holds.

The left-hand side of (11), a_i^j , is the capacity i is buying from j , while the right-hand side is the capacity it needs to buy from j to maintain a bottleneck of at least e_i . By definition—see (3)—the right-hand side can not be greater than the left-hand side. If the left-hand side is greater, the extra capacity bought from j does not increase the bottleneck capacity that i can actually offer on the sell-side, and therefore i will buy less from j . Thus, an equilibrium can occur only when equality holds.

These conditions can be rewritten in matrix form as

$$e = \Phi e + u \quad (12)$$

where, for $1 \leq i, j \leq N, j \neq i$

$$u_i = \sum_{m > N} a_m^i \left(1 + \sum_{k=1, k \neq i}^N \frac{r_k^i r_i^k}{1 - r_k^i r_i^k} \right)^{-1}$$

$$\phi_{i,i} = 0$$

$$\phi_{i,j} = \frac{r_j^i}{(1 - r_j^i r_i^j)} \left(1 + \sum_{k=1, k \neq i}^N \frac{r_k^i r_i^k}{1 - r_k^i r_i^k} \right)^{-1}.$$

The matrix $\Phi = (\phi_{i,j})_{1 \leq i, j \leq N}$, is the key to determining the stability of the game. The spectral radius of a matrix Φ , denoted $\rho(\Phi)$, is the largest of the moduli of the eigenvalues. Let $|\Phi| = (|\phi_{i,j}|)_{1 \leq i, j \leq N}$.

Consider now the brokers dynamically playing against each other. Specifically, on the buy side, each broker uses a best-reply strategy [27], and on the sell side, limits the offered capacity to the bottleneck capacity that it can obtain. Mathematically, the brokers' game is equivalent to a distributed computation to solve (12).

Proposition 7: The provisioning game, where brokers play asynchronously (i.e., each broker can act at any time, with no assumed order of turns, and variable but finite delays between turns), will converge to an equilibrium if and only if $\rho(|\Phi|) < 1$.

Proof: This follows from the above argument and the chaotic relaxation method [28], [29]. \square

Remark (Dynamical system interpretation): The users—through the demand vector u —can be viewed as external inputs driving a dynamic system, where the dynamics are governed by (10) the brokers: the system equation is then

$$e(t+1) = \Phi e(t) + u(t). \quad (13)$$

In this simplified view, all the brokers simultaneously adjust their offered quantities from $e_i(t)$ to $e_i(t+1)$, based on the demand vector $u(t)$. The convergence of the game is exactly the notion of stability of the dynamic system (13).

Remark: Brokers of different service classes do not buy from each other. But different service brokers in the same network do compete with each other to buy capacity from the RBS, and the RBS does not buy from any other player (see Fig. 1). Thus, we have the following matrix structures in, for example, a two class network:

$$\Phi = \begin{pmatrix} \Phi_{\text{class1}} & 0 & 0 \\ 0 & \Phi_{\text{class2}} & 0 \\ \text{Id} & \text{Id} & 0 \end{pmatrix} \quad (14)$$

where Id is the identity matrix, which is in the rows corresponding to the RBSs. Since the eigenvalues of Φ comprise all the eigenvalues of the diagonal blocks (i.e., Φ_{class1} , Φ_{class2} and 0), *the different service classes are independent with regard to stability.* Therefore, for any class, we need only take $(r_i^j)_{i, j \in \mathcal{I}}$ the matrix of the brokers' internetwork provisioning coefficients, derive the corresponding $|\Phi|$, and compute its eigenvalues to test whether or not the game among brokers is stable.

Remark: When all the r_i^j are equal, i.e., $r_i^j = r, \forall i, j, i \neq j$, we have:

$$\phi_{i,j} = \phi = \frac{r}{1 + (N-2)r^2}.$$

In this case, $|\Phi|$ has a single eigenvalue equal to $(N-1)\phi$ and $N-1$ eigenvalues equal to ϕ , and

$$\rho(|\Phi|) = (N-1)\phi = \frac{(N-1)r}{1 + (N-2)r^2}.$$

Specifically, when $N = 2$, $\rho(|\Phi|) = r$, so the convergence condition becomes $r < 1$.

When $N \geq 3$, the convergence condition $\rho(|\Phi|) < 1$ is equivalent to

$$\left(1 - \frac{N-1}{2}r\right)^2 > \left(\frac{N-3}{2}r\right)^2 \Leftrightarrow r < \frac{1}{N-2} \quad \text{or} \quad r > 1.$$

Therefore, *the equal provisioning game over more than two fully connected networks does not converge if $r \in [1/(N-2), 1]$.*

V. SIMULATIONS

The strategic game analysis in Section III establishes the optimal strategies and the existence of a stable and efficient operating point in the PSP games between dynamic buyers and static sellers. But these analyzes do not give any indication as to which particular equilibria will be reached. The provisioning matrix formulation in Section IV further reveals the stability condition of the provisioning game among dynamic sellers.

In what follows, we will use simulation to further study the DiffServ PSP framework and confirm the above analytical results under a realistic service provisioning scenario.

A. Simulation Configuration

We consider two classes of services, and hence two SBBs in each subnetwork.

- Class 2 is for reliable and high-quality service (e.g., the virtual leased line service considered by the EF PHB).
- Class 1 is for adaptive multimedia applications with less stringent quality requirements (like the Olympic Bronze service in Fig. 2).

In this scenario, best-effort service does not need any explicit capacity allocation. It is charged on flat rate and does not participate in the bandwidth auction market.

The simulation network has a mesh topology of three networks as shown in Fig. 1. Two access networks, A and B, connect to each other and to a backbone network M. Internetwork links are assumed to have a capacity equal to the capacity of the destination network.

The different degrees of provisioning for the two service classes are reflected in the routing factors r_i^j that are set according to Table II. One can observe the structural similarity between r_i^j in Table II and $\phi_{i,j}$ in (14).

The simulation parameters are given in Table III. To simulate the dynamics of subscribers switching among service providers, each user is modulated by an ON-OFF Markov process. At the beginning of an ON period, the user is connected randomly to

TABLE II
INTERNETWORK PROVISIONING COEFFICIENTS: r_i^j (EMPTY ENTRIES ARE ZERO)

seller	buyer									
	class 1 SBBs			class 2 SBBs			RBS'			
	A	B	M	A	B	M	A	B	M	
class 1 SBBs	A	0.3	0.2	0.1						
	B	0.2	0.3	0.1						
	M	0.5	0.5	0.8						
class 2 SBBs	A				1.0	0.4	0.1			
	B				0.4	1.0	0.2			
	M				1.0	1.0	1.0			
RBS'	A	1.0			1.0					
	B		1.0			1.0				
	M			1.0			1.0			

TABLE III
SIMULATION PARAMETERS

available bandwidth (Mbps)		
net A	net B	net M
40	40	150
user distribution:		
uniform across classes and networks		
20 "T1" users per class		
max capacity: unif. [0.75, 2.25] Mbps		
10 "T3" users per class		
max capacity: unif. [20, 60] Mbps		
mean ON interval	mean OFF interval	
10 time units	1 time unit	

one of the three networks (a uniform load distribution). During the ON period, a user continuously bids for bandwidth based on its valuation curve and presumably sends out traffic at a rate within the allocated bandwidth. During OFF periods, the user unsubscribes from the service. ON and OFF intervals are exponentially distributed with mean of 10 and 1 time units, e.g., one second or one week. In the remainder of this paper, we use one minute in simulation time as the time unit. The users are given randomly generated valuation curves, which model them as having elastic demand. Thus, a class 1 user i with a maximum capacity $\bar{a}_i = 1.5$ Mb/s will request a quantity ranging from 0 to 1.5 Mb/s of class 1 service capacity. Both the quantity and price of a bid depend not only on the player's valuation, but also on the market conditions (the requested quantities and bid prices of the other players).

B. Valuation Function

In Section III, we assumed a very general form (i.e., elastic demand) for a user's valuation. Further specification of users' valuations requires a market study on actual Internet users (see, for example, [30]).⁷ A realistic valuation model for wholesale Internet bandwidth over the last several years can be gleaned from the following observation [31]: cutting coming communication costs in half every twelve months, the market responded by doubling the traffic every six months.

This can be written as

$$\theta_i'(a_i) = \alpha_i / \sqrt{a_i}. \quad (15)$$

⁷Recall that the difficulty in developing realistic models is one of the reasons why auctions are advantageous in the first place, since the (run-time) mechanism itself (5)–(6) does not need to know the valuations.




Fig. 4. Trace at net M—horizontal axis is time in minutes.

Thus, in the simulations, we give our users valuations of the form

$$\theta_i(a_i) = 2\alpha_i \sqrt{a_i \wedge \bar{a}_i}. \quad (16)$$

In our simulation, for each class, we generate 20 users with \bar{a}_i drawn from a uniform distribution on $[0.75, 2.25]$ (we label these “T1” users which also include users of multiple or fractional T1), and 10 “T3” users with \bar{a}_i drawn from $[20, 60]$ (Mb/s). The parameter α_i is also chosen randomly such that $\theta_i(\bar{a}_i)$ is uniform on $[0.6, 1.8]$ (c/min) for the T1 users, and on $[18.0, 54.0]$ (c/min) for the T3 users.⁸

As mentioned in Section III, the broker’s buy-side valuation must be smoothed. We select the same form as in (16). To fit the curve to the demand, the broker dynamically sets $\bar{a}_i = \sum_j q_j^i$ and chooses α_i such that $\theta_i(\bar{a}_i) = \sum_j q_j^i p_j^i$. In (15), note that as a approaches zero, the marginal valuation approaches infinity. In some circumstances, this last feature can be useful. A finite maximum marginal valuation would make it possible for the broker to be completely shut out (i.e., $a_j^l = 0$ at some peer l where enough users have very higher valuations), and when one broker is shut out, so are all its peers, and the service is no longer offered on an internetwork basis.

C. Stability of Market Pricing Mechanisms

In this subsection, we focus on the demand side, and illustrate the results of Section III.

The simulations are run with the full dynamics of both the demand and supply sides, i.e., users behave according to Propo-

⁸These numbers roughly correspond to capacities and prices in today’s Internet access market. We randomize both to reflect the wider variety of access speeds and willingness to pay that are likely with future (differentiated) services.

sition 3 and brokers according to Proposition 4 on the buy-side. On the sell-side, as required by Proposition 1, the brokers do not sell more than the expected bottleneck capacity (3), and they do so by setting a buy-back bid as explained in the remark following Proposition 1. However, we intentionally omit the floor price p_j^i that ensures the broker profitability, in order to see where profits are likely to be realized.

Simulation traces of the state of the six SBBs (two in each of the three networks) are presented in Figs. 4–6. Each figure contains four plots showing the total demand at that SBB (sum of bid quantities), the offered quantity, which is the expected bottleneck e_i (see Proposition 1), the market price, and the SBBs profit. Each quantity is shown for class 1 (solid line) as well as class 2 (dotted line).

We observe the following.

- Despite the dynamics of arrivals and departures, the two classes remain stable and the SBBs are able to maintain consistent offered capacities e_i in all three networks; price changes reflect the supply and demand, and the dynamic market successfully allocates resources, which demonstrates that the PSP distributed market mechanism can quickly converge to the equilibrium given by Proposition 5.
- In each network, as expected, the higher quality class 2 is more expensive. This is despite the fact that the demand from the users is statistically identical; thus, the difference in price arises through market dynamics, and is purely due to the provisioning coefficients (i.e., corresponds to a difference in quality).
- Relatedly, the bottleneck (or offered quantity) is smaller for class 2 in all cases. These two effects (smaller bottle-

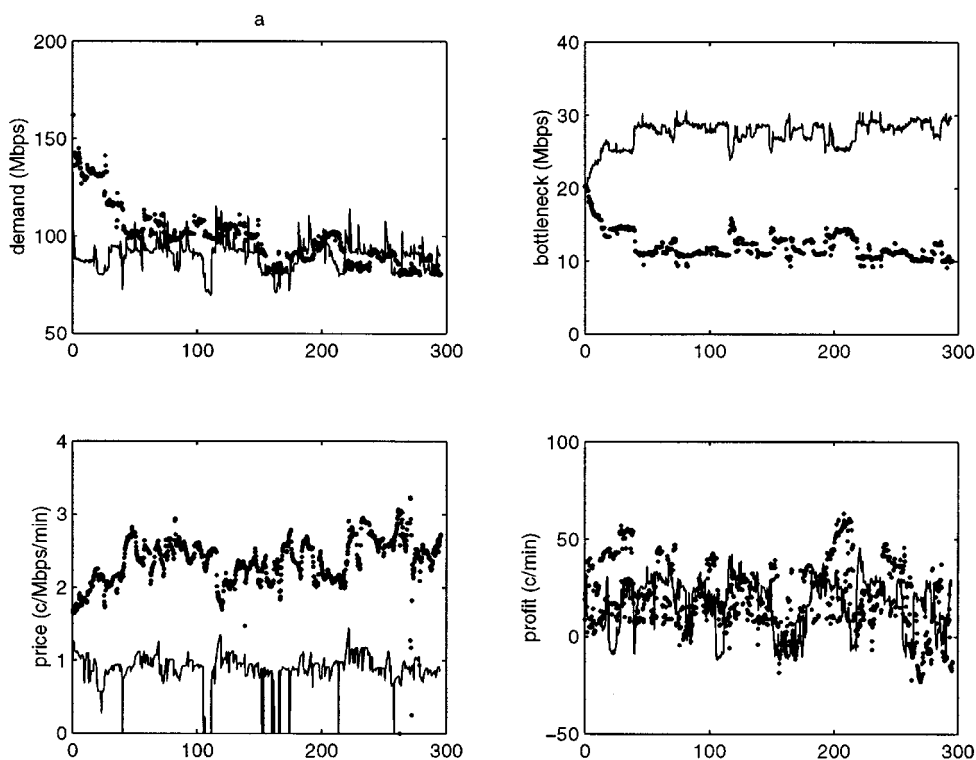


Fig. 5. Trace at net A—horizontal axis is time in minutes.

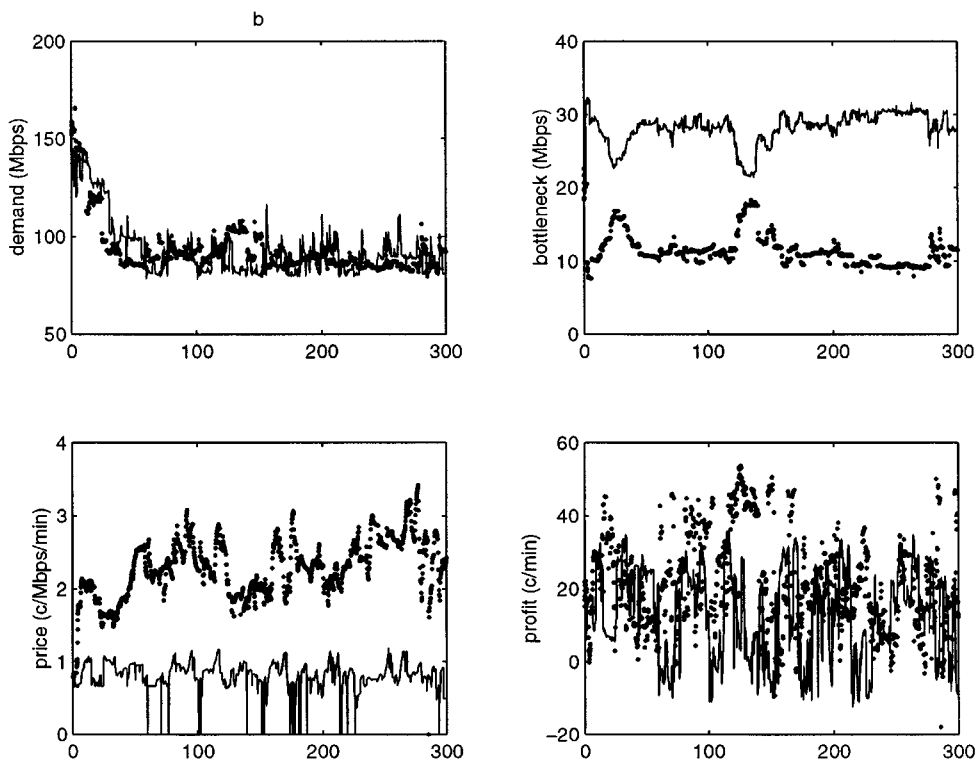


Fig. 6. Trace at net B—horizontal axis is time in minutes.

neck and higher price) balance each other out, and allow the SBBs to coexist while having differentiated quality. For example, if the market price of class 1 in network A drops “too low,” then that SBB cannot compete with the SBB of class 2 in the same network in buying from their common RBS, which causes the first SBB to reduce the

quantity of class 1 service offered in network A, which then causes more intense competition among the buyers of that service, and hence a price rise.

- The high-quality class 2 has a slightly higher share in the high-capacity network M (about 1/3 of the capacity) than it does in the smaller networks (about 1/4 of the capacity);




Fig. 7. Spectral radius as a function of internetwork provisioning coefficients, instability arises in the top left and bottom right quadrants.

this is because the demand is equally distributed across the three; therefore, M has less competition for resources; and therefore, an overprovisioned class is sustainable at a higher share of the total.

- Indeed, the large network M is consistently less expensive (in terms of unit market price) than the smaller ones.
- All SBBs remain profitable over the long run, despite not having reserve (minimum) prices, which validates the broker strategy of Proposition 4. Whenever one SBBs profit is momentarily negative, then its RBS or a peer SBB is making a corresponding extra profit. However, for the same reason outlined above, competition for the underlying resources (at the RBS level) prevents one class from being substantially more profitable than the other.

The simulation of the stable scenario provides a sanity check on the market mechanisms, and indeed results are completely in line with intuition. In the next section, we consider unstable scenarios, which as we shall see, do not always yield to intuition.

D. Stability of Internetwork Provisioning

Consider now three interconnected networks, with just one class, i.e., three brokers $\{1, 2, 3\}$. Let $r_1^2 = r_2^1 = x$, $r_1^3 = r_3^1 = y$, and $r_i^j = 0.99$ for all other pairs i, j . Fig. 7 shows $\rho(|\Phi|)$ as a function of x and y . The figure shows that when $x > 1$ and $y < 1$, or vice versa, the provisioning of this class becomes unstable. It is interesting to note that simply overprovisioning $x > 1$ and $y > 1$ does not give rise to instability. Thus, instability can be due more to asymmetry in the flows rather than to the actual degree of overprovisioning.

Neither can instability be simply attributed to the existence of “cycles” in the graph of the network. Fig. 8 shows a scenario where a single class network—with a simple topology of two access networks connected to a backbone network—can be unstable even if the graph of the network has no cycles. In Fig. 8(b), the right-hand side shows the allocations for traffic going from A to M (dotted curve), and the bottleneck capacity in A itself


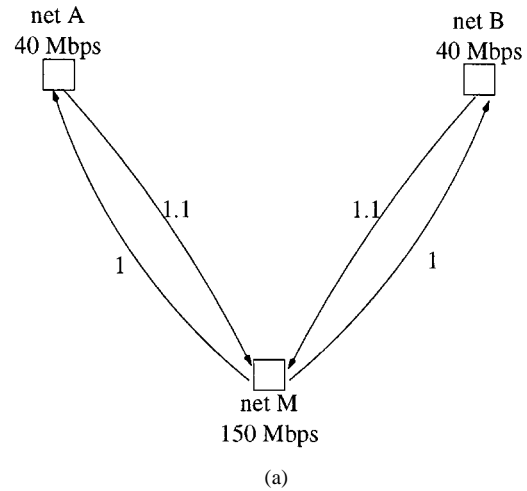


Fig. 8. Simulation of one unstable class, in the right-hand side of plot (b), the solid curve represents bottleneck bandwidth and the dotted curve represents allocated bandwidth. The horizontal axis is the number of simulation time units. The scenario is unstable as allocations do not converge. (a) Simulation topology ($\rho(\Phi) = 1.02$). (b) Trace at net A.

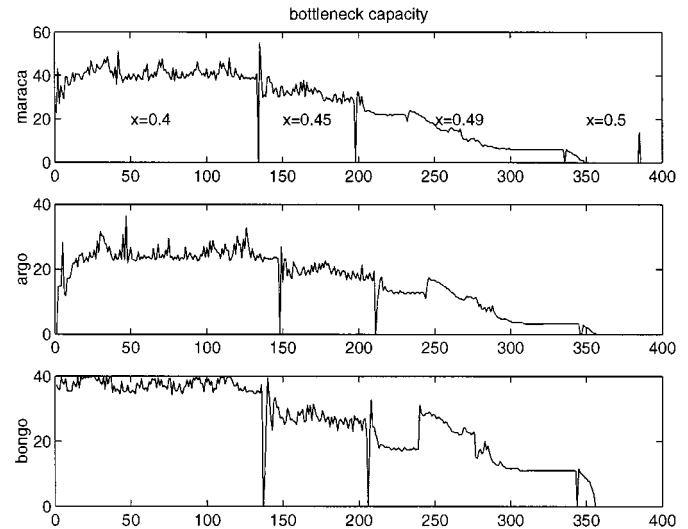


Fig. 9. “Dis-peering” effect, the legend of x axis is the number of simulation time units.

(solid curve). The instability is reflected in the volatility of the allocated capacities.

In a stable scenario, one must still worry about what kind of equilibrium is reached. Indeed, it can happen that the only equilibrium for a stable class is one where all the bottlenecks are zero. Fig. 9 illustrates this possibility, which we refer to

as “dis-peering.” Here, we simulate the network shown in Fig. 1, with a single class that is provisioned identically in all directions, i.e., $\forall i, j, i \neq j, r_i^j = x$. As x approaches 0.5, the bottleneck becomes smaller, until finally, none of the brokers has any capacity to sell. Here, there is only one class, and the physical capacity as well as the average demand from the users remains constant (even though users do come and go—see Table III). Thus, the reduction in bottlenecks is purely a result of the provisioning dynamics, and not of other traffic “squeezing out” this class. Indeed, since capacity is edge-allocated, a broker must provision for all possible routes (here there are two, one to each peer network), with a degree of assurance x . When this required assurance x reaches a critical level (which depends on the topology), it becomes impossible for the broker to satisfy any demand. This is one of the “penalties” to be incurred in exchange for the simplicity and scalability of edge-capacity allocation with stateless service differentiation. Indeed, if the broker could offer allocations tied to specific routes (e.g., with techniques such as MPLS [36]), “dis-peering” would not occur.

This effect may also be the converse of what has been observed in the current (best-effort only) Internet. In recent years, some large ISPs have decided it is not in their interest to peer free of charge with some smaller ones because they would do better by selling the bandwidth directly to their own customers [32]. Here, with differentiated services, a broker in a large network may decide to set $r_i^j = 0$ in the direction of the smaller networks (i.e., not to buy any differentiated service from the smaller network), when it is not worthwhile to get the allocations required for a high level of assurance in a congested network. Other related phenomena have been studied in the literature [33]–[35].

VI. CONCLUSION

We have presented a decentralized auction-based pricing approach for differentiated internet services. Our game-theoretic analysis identifies the best strategies for end users and bandwidth brokers. The analysis proves the existence of efficient stable operating points, and the simulations indicate that even an aggregate 50% difference in the degree of provisioning between two services does not lead to extreme differences in the market price of services, and partitioning of bandwidth between services, because of the competition among service brokers for the underlying resources (e.g., bandwidth).

In investigating the stability of provisioning differentiated internet services using a distributed game theoretic model, our results indicate that, in an internet with multiple differentiated classes competing for the same resources, even though the demand for one service affects the amount of capacity available for another, the *stability* of each class is independent of the others’. Thus, the good news is that dynamic market-driven partitioning of network capacity among services appears sustainable. The bad news is that very conservatively provisioned services can be unstable on this macro-level, even in the simplest network topologies. Even in stable cases, the only sustainable outcome may be not to peer for differentiated service traffic. These results

are not merely artifacts of PSP or of any particular pricing mechanism. They appear to be fundamental issues of market-driven peering under edge capacity allocation.

The dynamic system formulation of (13) suggests an interesting direction for future work. It may be possible to achieve certain wide-area network objectives (e.g., stability or avoiding “dis-peering”) by exercising feedback control. If such controls can be derived and are not too large in magnitude, they could be applied by injecting some service requests at multiple strategic edge points to drive the brokers of that specific class to a beneficial equilibrium. Another direction for further work is the study of the interaction between edge-allocation such (as in DiffServ) and route-pinning approaches (such as MPLS [36]), which may provide the most immediate means of addressing potentially unstable peering configurations.

APPENDIX I

BROKER’S BUY-SIDE COORDINATION

Lemma 1 (Broker coordination): Let $j \in \mathcal{I}$ be a broker. For any profile s , $s_j = (q_j, p_j)$, let $a \equiv a(s)$ be the allocations that would result, and $m = \arg \min_k e_j^k(a)$. Then, a better reply for the broker is $x_j = (z_j, p_j)$, where $\forall l \neq j$

$$z_j^l = \left[e_j^m(a) - a_l^j \right] r_j^l.$$

That is, $u_j(x_j; s_{-j}) \geq u_j(s)$. Moreover,

$$a_l^j(z_j, p_j) = z_j^l. \quad (17)$$

Proof: To avoid cluttered notation, since s_{-j} is fixed, we will omit it, writing, e.g., $u_j(\cdot, \cdot) \equiv u_j((\cdot, \cdot); s_{-j})$. Also, the argument of the function will be omitted when it is simply s , so that $u_j \equiv u_j(s_j) \equiv u_j(s_j; s_{-j})$. Note that, since we are holding all the other players fixed, and varying only the buy-side of player j , only the quantities with subscript j will change. In particular, a_l^j remains the same throughout.

We will show that

$$u_j \equiv u_j(q_j, p_j) \leq u_j(z_j, p_j). \quad (18)$$

Now, $\forall l \in \mathcal{I}$,

$$\begin{aligned} z_j^l &= \left[e_j^m(a) - a_l^j \right] r_j^l \\ &\leq \left[e_j^l(a) - a_l^j \right] r_j^l = a_l^l \\ &\leq \left[q_l^l - \sum_{r_k^l \geq r_j^l, k \neq j} q_k^l \right]^+ \end{aligned} \quad (19)$$

where the last line follows from (5). Now, using (5) again, we get

$$\begin{aligned} a_l^j(z_j, p_j) &= \left[q_l^l - \sum_{r_k^l \geq r_j^l, k \neq j} q_k^l \right]^+ \wedge z_j^l = z_j^l \\ &= \left[e_j^m(a) - a_l^j \right] r_j^l \end{aligned}$$

where the second equality follows from (19), and the last is by definition. This proves (17). Thus, we have $e_j^l(a(z_j, p_j)) = a_j^l(z_j, p_j)/r_j^l + a_j^l = e_j^m(a)$, and this holds for $l \neq j$. Therefore, by Proposition 2, $\theta_j(a(z_j, p_j)) = \theta_j(a)$, i.e., changing the bids from (q_j, p_j) to (z_j, p_j) does not change j 's bottleneck value. Therefore,

$$\begin{aligned} u_j(z_j, p_j) - u_j &= \sum_{l \neq j} c_j^l - c_j^l(z_j, p_j) \\ &= \sum_{l \neq j} \int_{a_j^l(z_j, p_j)}^{a_j^l} f^l(q_l^l - z) dz. \end{aligned}$$

Now $\forall l, e_j^m(a) \leq e_j^l(a) \Rightarrow z_j^l/r_j^l + a_j^l \leq a_j^l/r_j^l + a_j^l \Rightarrow a_j^l \geq z_j^l \geq a_j^l(z_j, p_j)$, where the last inequality follows from (5). That along with the fact that $f^l \geq 0$ implies $u_j(z_j, p_j) - u_j \geq 0$. \square

REFERENCES

- [1] Internet Engineering Task Force. Differentiated services working group. <http://www.ietf.org/html.charters/diffse,rv-charter.html>
- [2] D. Clark and W. Fang, "Explicit allocation of best-effort packet delivery service," *IEEE/ACM Trans. Networking*, vol. 6, no. 4, pp. 362–373, Aug. 1998.
- [3] S. Shenker, C. Partridge, and R. Guerin, "Specification of guaranteed quality of service," IETF RFC 2122, <ftp://ftp.isi.edu/innotes/rfc2122.txt>, Sept. 1997.
- [4] V. Jacobson, K. Nichols, and K. Poduri, "An expedited forwarding PHB," Internet RFC 2598, <ftp://ftp.isi.edu/in-notes/rfc2598.txt>, June 1999.
- [5] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, "Assured forwarding PHB group," Internet RFC 2597, <ftp://ftp.isi.edu/in-notes/rfc2597.txt>, June 1999.
- [6] D. Clark, "Internet cost allocation and pricing," in *Internet Economics*, L. W. McKnight and J. P. Bailey, Eds. Cambridge, MA: MIT Press, 1997.
- [7] K. Nichols, V. Jacobson, and L. Zhang. (1997) A two-bit differentiated services architecture for the internet. <ftp://ftp.ee.lbl.gov/papers/dsarch.pdf>
- [8] K. Kilkki, J. Ruutu, and O. Strandberg, "High quality and high utilization—Incompatible objectives for internet?," in *IEEE/IFIP 6th Int. Workshop Quality Service*, Napa Valley, CA, May 1998.
- [9] Z. Wang, "A case for proportional fair sharing," in *IEEE/IFIP 6th Int. Workshop on Quality of Service*, Napa Valley, CA, May 1998.
- [10] F. P. Kelly, "Charging and rate control for elastic traffic," *European Trans. Telecommun.*, vol. 8, 1997.
- [11] N. Semret, "Market mechanisms for network resource sharing," Ph.D. thesis, Columbia Univ., 1999.
- [12] S. Shenker, D. Clark, D. Estrin, and S. Herzog, "Pricing in computer networks: Reshaping the research agenda," *ACM Comput. Commun. Review*, vol. 26, no. 2, pp. 19–43, 1996.
- [13] A. Gupta, D. O. Stahl, and A. B. Whinston, "Priority pricing of integrated services networks," in *Internet Economics*, L. W. McKnight and J. P. Bailey, Eds. Cambridge, MA: MIT Press, 1997.
- [14] N. Anerousis and A. A. Lazar, "A framework for pricing virtual circuit and virtual path services in ATM," in *Proc. IEE Int. Teletraffic Congr.*, Washington, DC, June 1997.
- [15] Y. A. Korilis, T. A. Varvarigou, and S. R. Abuja, "Incentive-compatible pricing strategies in noncooperative networks," in *Proc. IEEE Infocommun.*, 1998.
- [16] I. C. Paschalidis and J. N. Tsitsiklis, "Congestion-dependent pricing of network services," Boston Univ., Tech. Rep., 1998.
- [17] E. Fulp, M. Ott, D. Reininger, and D. Reeves, "Paying for qos: An optimal distributed algorithm for pricing network resources," in *IEEE/IFIP 6th Int. Workshop Quality Service*, Napa Valley, CA, May 1998.
- [18] A. A. Lazar and N. Semret, "Design and analysis of the progressive second price auction for network bandwidth sharing," *Telecommunication Systems, Special Issue on Network Economics*, 2000. <http://comet.columbia.edu/~nemo/pub.html>.
- [19] M. T. T. Hsiao and A. A. Lazar, "An extension to Norton's equivalent," *Queueing Syst., Theory Appl.*, vol. 5, pp. 401–412, 1989.
- [20] I. Stoica and H. Zhang, "LIRA: An approach for service differentiation in the internet," in *Proc. of NOSSDAV'98*, Cambridge, England, July 1998.
- [21] National Research Council Renaissance Committee, *Realizing the Information Future: The Internet and Beyond*. Washington, DC: National Academy Press, 1994.
- [22] W. Vickrey, "Counterspeculation, auctions and competitive sealed tenders," *J. Finance*, vol. 16, 1961.
- [23] E. H. Clarke, "Multipart pricing of public goods," *Public Choice*, vol. 8, pp. 17–33, 1971.
- [24] T. Groves, "Incentives in teams," *Econometrica*, vol. 41, no. 3, pp. 617–631, July 1973.
- [25] R. B. Myerson, "Optimal auction design," *Mathematics Operations Research*, vol. 6, no. 1, pp. 58–73, Feb. 1981.
- [26] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1991.
- [27] N. Semret, R. R.-F. Liao, A. T. Campbell, and A. A. Lazar, "Market pricing of differentiated internet services," in *IEEE/IFIP 7th Int. Workshop Quality Service*, 1999.
- [28] D. Chazan and W. Miranker, "Chaotic relaxation," *Linear Algebra Appl.*, vol. 2, 1969.
- [29] A. D. Bovopoulos and A. A. Lazar, "Decentralized algorithms for optimal flow control," in *Proc. 25th Allerton Conf. Commun., Control Computing*, Oct. 1987.
- [30] B. Rupp, R. Eden, H. Chand, and P. Varaiya, "INDEX: A platform for determining how people value the quality of their internet access," in *IEEE/IFIP 6th Int. Workshop Quality Service*, Napa Valley, CA, May 1998.
- [31] L. G. Roberts, "Beyond Moore's law: Internet growth trends," *IEEE Computer Mag.*, Jan. 2000.
- [32] R. S. Benn. (1997) The great peering debate. <http://www.clack.net/pub/rbenn/debate.html>
- [33] P. Baake and T. Wichmann, "On the economics of internet peering," *Netnomics*, vol. 1, no. 1, 1999.
- [34] J. Gong and P. Sriganesh, "An economic analysis of network architectures," *IEEE Network*, pp. 18–21, Mar./Apr. 1996.
- [35] P. Sriganesh, "Internet cost structures and interconnection agreements," in *Internet Economics*, L. W. McKnight and J. P. Bailey, Eds. Cambridge, MA: MIT Press, 1997.
- [36] B. Davie and Y. Rekhter, *MPLS Technology and Applications*. San Mateo, CA: Morgan Kaufmann, 2000.

Nemo Semret received the Ph.D. degree at Columbia University, where he was awarded the Eliah I. Jury award for his dissertation in electrical engineering in 1999, the M.Eng. degree in electrical engineering, and the B.Eng. degree with Honors in electrical engineering and a minor in mathematics, both from McGill University.

He is a co-founder of Invisible Hand Networks, and currently serves as CEO and CTO. He is co-inventor on 3 patent applications. Prior to and during his Ph.D., he worked variously as a Researcher, Consultant, and Intern at Telcordia Technologies (Bellcore), Lucent Technologies Bell Labs, the United Nations, Intelsat, and INRS-Telecom.

Raymond R.-F. Liao received the M.A.Sc. degree in electrical and computer engineering from University of Toronto, Canada and the Bachelor degree from Huazhong University of Science and Technology, China.

He is a Ph.D. candidate and Graduate Research Assistant of the COMET Group at the Center for Telecommunications Research, Columbia University since 1996. Before that, he worked at Newbridge Networks Corporation for three years on ATM network performance analysis and traffic management, and co-authored 3 patent applications. He His current research focuses on dynamic provisioning of differentiated service Internet, and realizing adaptive quality of service in mobile multimedia networks with methodologies including programmable networks and network economics.

Andrew T. Campbell received the Ph.D. degree in computer science in 1996, the IBM Faculty Award in 1999, and the NSF CAREER Award for his research in programmable mobile networking in 1999.

He has been an Assistant Professor of Electrical Engineering and member of the COMET Group at the Center for Telecommunications Research, Columbia University since 1996. His research interests include open programmable networks, mobile networking, and quality of service. Dr. Campbell is active in the IETF and the international working group on Open Signaling (OPENSIG). His research is currently sponsored by a number of governmental agencies and industrial companies.

Andrew is a past co-chair of the 5th IFIP/IEEE International Workshop on Quality of Service (IWQOS97) and the 6th IEEE International Workshop on Mobile Multimedia Communications (MOMUC99) and is currently a co-chair of the 4th IEEE International Conference on Open Architecture and Network Programming (OPENARCH 2001). He is a Guest Editor for the IEEE Journal on Selected Areas in Communications on Active and Programmable Networks and the IEEE Personal Communications special issue on IP-based Mobile Telecommunications Networks and a Member of the Editorial Board of Computer Networks.

Aurel A. Lazar (F'93) (<http://comet.columbia.edu/~aurel/>) has been a Professor of Electrical Engineering at Columbia University since 1988.

Professor Lazar's current theoretical research interests are on networking games and pricing. His experimental work focuses on building open programmable networks. This work has led to the establishment of the IEEE standards working group on Programming Interfaces for Networks (<http://www.ieee-pin.org>).

Professor Lazar was instrumental in establishing the OPENSIG (<http://comet.columbia.edu/opensig>) international working group with the goal of exploring network programmability and next generation signaling technology. He was the Program Chair of the Fall and Spring OPENSIG'96 workshops and of the First IEEE Conference on Open Architectures and Network Programming (OPENARCH'98, <http://comet.columbia.edu/openarch>).