Peering and Provisioning of Differentiated Internet Services

N. Semret

University of California Los Angeles and Invisible Hand Networks, Inc. nemo@invisiblehand.net

R. R.-F. Liao, A. T. Campbell

Columbia University {liao,campbell}@comet.columbia.edu

A. A. Lazar

Columbia University and Xbind, Inc. aurel@comet.columbia.edu

Abstract—**A key consideration in building differentiated network services is the feasibility of maintaining stable and consistent service level agreements across multiple networks where allocations are made only at the edges. To investigate this, we consider a game theoretic model of capacity provisioning in a differentiated services internet. The players 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-users respectively in a two-tier whole-seller/retailer market. Based on this model, we are able to construct an explicit necessary and sufficient condition for the stability of the game, which determines the sustainability of a given set of service level agreement configurations among peering Internet Service Providers. The analytical results are validated with simulations of user and broker dynamics, using distributed progressive second price auctions as the spot market mechanism in a scenario with three inter-connected networks and two services.**

Keywords— **differentiated service, capacity provisioning, second price auction, peering stability**

I. Introduction

The recent development of the differentiated service (diff-serv) Internet model is aimed at supporting service differentiation for aggregated traffic in a scalable manner. The tenet of diff-serv is to relax the traditional "hard" Quality of Service (QoS) models (e.g., end-to-end per-flow guarantees of Int-serv and ATM) along two dimensions: slower time-scale network mechanisms and coarser-grained traffic provisioning.

An important aspect of the problem that has not been systematically addressed is the feasibility of maintaining consistent service level agreements (SLAs) – or diff-serv profiles – across inter-connected 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.

While the role of prices as an essential indicator for resource allocation control has been established from the outset of diff-serv [1], [2], the precise development of pricing mechanisms is still in its early stages. 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 the Internet, because of the diverse and rapidly evolving nature of the applications, services, and population, the case is more compelling. With suitably designed rules, auctions can achieve efficient (value maximizing) allocations with minimal a priori information.

In [3], we showed that when the Progressive Second Price (PSP) auction [4] is used as the diff-serv bandwidth market mechanism, it achieves economic objectives (i.e., incentive compatibility and efficiency), while being realistic in the engineering sense (i.e., small signalling and computation load). As such, it provides a useful baseline for understanding the conditions for the economic feasibility of wide-area differentiated services.

In this paper, we consider the capacity provisioning problem. Our main focus is to study the feasibility of maintaining *stable* and *consistent* SLAs across multiple networks where allocations are made only at the edges. To investigate this, in Section II, we construct a two-tier whole-seller/retailer market model, giving the wide-area pricing formulation, and the model for provisioning and differentiation of the services. Following this, in Section III, we present our analytical results identifying a necessary and sufficient condition for the stability of the SLA provisioning game. The analytical results are validated by simulations of a scenario with two types of services over three inter-connected networks. The simulations also illustrate conditions which lead to certain classes of service not being offered on an inter-network basis. Finally, in Section IV, 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" [5]). The capacity may be represented by an absolute amount of bandwidth, or some relative metrics like user share in the User-Share Differentiation proposal [6] or resource token in Location Independent Resource Accounting [7]. 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 peer[1] is stable over the resource allocation time scale (e.g., in the order of hours).

Figure 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

[1] In this paper, we use the term "peer" in the most general sense, i.e., any network which inter-connects with a given network, and not just those that choose to exchange all traffic free of charge.

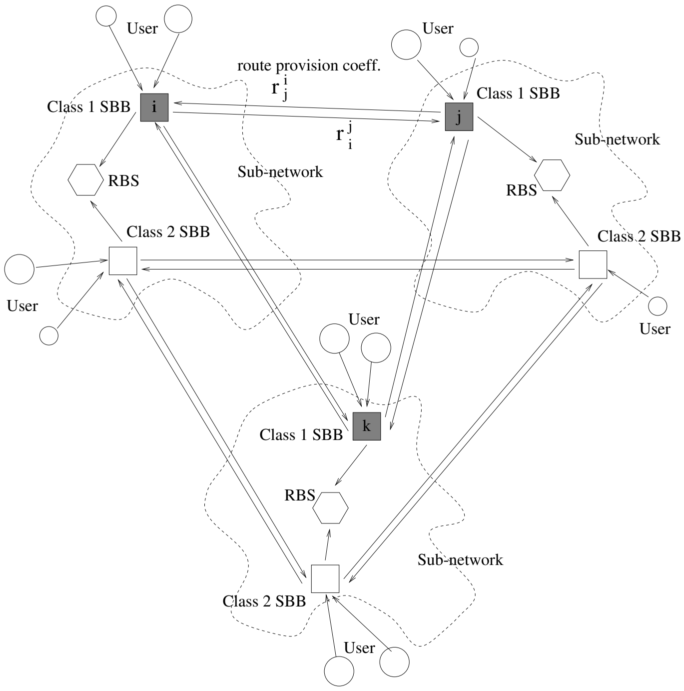

Fig. 1. 2-tier auction pricing framework for a diff-serv internet

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 [8]. If the RBS and multiple SBBs on the same network are not owned by the same entity, a non-cooperative 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, based on the demands from users. The users, or retail buyers, are subscribers to a particular service offered by a particular provider. In the diff-serv context, these will likely be large subscribers (e.g., web sites, intra/extranets, virtual private networks), rather than individual end-users.

B. Provisioning and Peering Constraints

Let the set of all players, including buyers, sellers and brokers (brokers are both buyers and sellers) be denoted $\mathcal{I}$. Following the notation in [4], 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. For seller j, its allocation to the player i is denoted a^j_i.

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

$$\sum_{m \in \mathcal{I}} a^i_m \leq a^l_i. \qquad (1)$$

Second, since it is selling wide-area service, i must get enough capacity from the SBBs offering the same service in each peer network. Let j denote one such peer SBB, and r^j_i be the "fraction of traffic" generated by i's customers that is routed to the network where player j is the peer SBB. More generally, r_i is a vector of inter-network provisioning coefficients, as discussed in Section II-C. Then, i must satisfy

$$r^j_i \sum_{m \in \mathcal{I}, m \neq j} a^i_m \leq a^j_i, \qquad (2)$$

for all peers j.[2] For notational convenience, fix $r^l_i = 1$, when l is i's RBS. Since $a^i_l = 0$, (2) includes (1) as the special case $j = l$. If j is neither a peer of i, nor its RBS, then we set $r^j_i = 0$.

Define, for any allocation a,

$$e^j_i(a) \triangleq \frac{a^j_i}{r^j_i} + a^i_j, \qquad (3)$$

$$e_i \triangleq \min_{j \neq i} e^j_i(a) \qquad (4)$$

We call e_i the *bottleneck* capacity of player i. For an SBB i, the offered quantity of bandwidth can be no larger than e_i, which follows from (2). Indeed, if SBB i offers a capacity greater than e^j_i, for some j, and that capacity is fully bought by its users, i.e., $\sum_m a^i_m > e^j_i = \frac{a^j_i}{r^j_i} + a^i_j \Rightarrow r^j_i \sum_{m \in \mathcal{I}, m \neq j} a^i_m > a^j_i$, then (2) is violated.

C. Differentiating Services

We do not explicitly consider the per-hop behaviors (PHBs) 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, e.g.:

- *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 diff-serv assured forwarding (AF) per-hop behaviors [9]. In this case, r^j_i is the expected fraction of the total traffic entering i that is routed to j. r^i_i is the fraction of traffic that terminates with one of i's own customers, and $\sum_{j \neq l} r^j_i = 1$, where l is the RBS in i's network.[3]
- *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 diff-serv expedited forwarding (EF) per-hop behavior [10]. In this case $r^j_i = 1$ for all peers j.
- *local SLA*; for an SBB which offers SLAs valid only within its own network, $r^i_i = 1$ and $r^j_i = 0, \forall j \neq i$.

[2] We assume that service providers block "loop-back" traffic, i.e., the traffic going from j through i and back to j. If that is not the case, then the summation would be over all m.

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

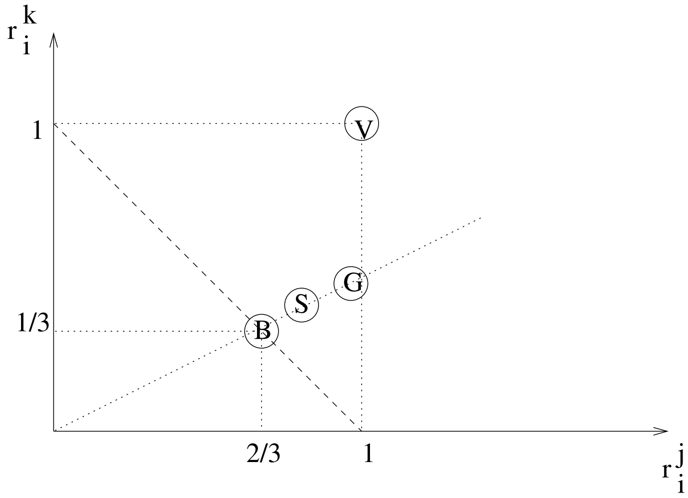

Fig. 2. Inter-network provisioning coefficients for Olympic Gold, Silver and Bronze services, and the Virtual Leased Line service

Figure 2 illustrates several service scenaria 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., 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 [9], and the "Virtual Leased Line" (VLL) service based on EF [10]. Degrees of over-provisioning 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 "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 Figure 2.

Depending on the scheduling and buffer management algorithms used to provide the PHBs, some amount of over-provisioning may be required [10]. These engineering needs can be represented in this model by simply factoring over-provisioning into each coefficient of r, e.g., if i is offering a virtual leased line with 5% over-provisioning then $r_i = (1.05, 1.05, \ldots)$.

Note that for our purposes, the provisioning coefficients r_i are known by broker i in advance, since they represents 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 micro-flows.

III. Dynamic Provisioning of Differentiated Services

A. Analysis

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 brokers, the two games are interdependent, and may be played on the same or vastly different timescales.

On the demand side, the analysis of [4], [3], [11] derives the optimal (utility-maximizing) bidding strategies for users and brokers, and establishes 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 non-zero asking (or "reserve") price – which can be arbitrarily small. Thus prices will always have a strictly positive floor.

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 Equation (4), where a_i^j's are chosen to maximize it's profit; for details see [3]. We will then use simulations using PSP auctions to verify that our insights are valid when the two games are coupled.

Define the vector $e = (e_1, \ldots, e_i, \ldots, e_N)$ for any profile of allocations a, where e_i is the bottleneck capacity of seller i as given by (4), and $\{1, \ldots, N\}$ is the subset of $\mathcal{I}$ consisting of all the sellers (RBS' and SBBs). Pure buyers (users) are assumed to be players numbered $m = N+1, N+2, \ldots$. From (2) and (4), at the equilibrium point, the following conditions will hold for $1 \leq i \leq N$:

$$e_i = \sum_{j \in \mathcal{I}, j \neq i} a_j^i \quad (5)$$

$$a_i^j = (e_i - a_j^i) r_i^j. \quad (6)$$

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

The left hand side of (5), 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 equilibirum can occur only when equality holds.

The left-hand side of (6), 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 (4) – 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 equilibirum can occur only when equality holds.

These conditions can be re-written in matrix form as

$$e = \Phi e + u, \tag{7}$$

where for $1 \le i, j \le N, j \ne i$,

$$u_i = \sum_{m>N} a^i_m \left(1 + \sum_{k=1,k\ne i}^{N} \frac{r^i_k r^k_i}{1 - r^i_k r^k_i}\right)^{-1},$$

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

$$\phi_{i,j} = \frac{r^i_j}{(1 - r^i_j r^j_i)} \left(1 + \sum_{k=1,k\ne i}^{N} \frac{r^i_k r^k_i}{1 - r^i_k r^k_i}\right)^{-1}.$$

The matrix $\Phi = (\phi_{i,j})_{1\le i,j\le N}$ 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 $|\Phi| = (|\phi_{i,j}|)_{1\le i,j\le N}$.

Consider now the brokers dynamically playing against each other. Specifically, on the buy side, each broker uses a best-reply strategy [3], 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 (7).

Proposition 1: 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 an only if $\rho(|\Phi|) < 1$.
Proof: This follows from the chaotic relaxation method [12], [13]. □

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 the brokers: the system equation is then

$$e(t+1) = \Phi e(t) + u(t). \tag{8}$$

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 (8).
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 Figure 1). Thus, we have the following matrix structures in, for example, a two class network:

$$\Phi = \begin{pmatrix} \Phi_{\text{class1}} & 0 & 0 \\ 0 & \Phi_{\text{class2}} & 0 \\ \text{Id} & \text{Id} & 0 \end{pmatrix}, \tag{9}$$

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 $\mathbf{0}$), *the different service classes are independent with regard to stability.* Therefore, for any class, we need only take $\left(r^j_i\right)_{i,j\in\mathcal{I}}$ the matrix of the brokers' inter-network 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^j_i are equal, i.e., $r^j_i = r, \forall i, j, i \ne j$, 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 $N \ge 3$, the convergence condition $\rho(|\Phi|) < 1$ is equivalent to:

$$(1 - \frac{N-1}{2}r)^2 > (\frac{N-3}{2}r)^2 \Leftrightarrow r < \frac{1}{N-2} \;\; or \;\; r > 1.$$

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

B. Simulations

In what follows, we will use simulation to confirm the above observations under realistic service provisioning scenaria. We consider two classes of services, and hence, two SBBs in each sub-network:

- class 2 is for reliable and high quality service (e.g., the virtual leased line service considered by the EF PHB), and;
- class 1 is for adaptive multimedia applications with less stringent quality requirements (like the Olympic Bronze service in Figure 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 Figure 1. Two access networks, argo and bongo, connect to each other and to a backbone network maraca. Inter-network 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 I. One can observe the structural similarity between r^i_j in Table I and $\phi_{i,j}$ in Equation 9.

The simulation parameters are given in Table II. 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 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 720 and 72 time units. The users are given randomly generated valuation curves, which model them as having elastic demand [4]. Thus, a class 1 user with a maximum capacity of 1.5Mbps will request a quantity

TABLE I

INTER-NETWORK PROVISIONING COEFFICIENTS: r^i_j (EMPTY ENTRIES ARE ZERO, A: ARGO, B: BONGO, M: MARACA)

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 II

SIMULATION PARAMETERS

available bandwidth (Mbps)		
argo	*bongo*	*maraca*
40	40	150

user distribution: uniform across classes and networks	
20 T1 users	*10 T3 users*
max capacity: 1.5 Mbps	max capacity: 40 Mbps

mean ON interval	*mean OFF interval*
720 time units	72 time units

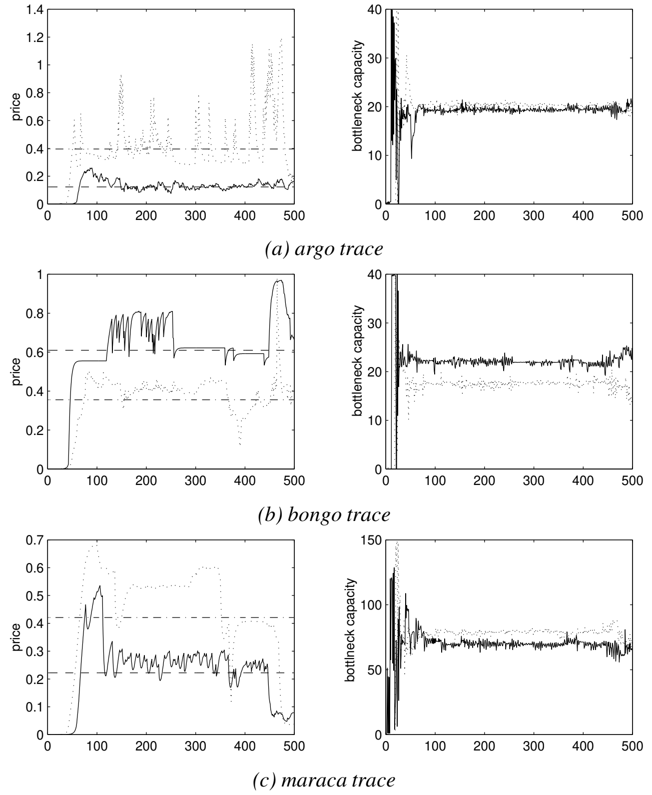

Fig. 3. Simulation of three networks with two stable classes, solid curves represent class 1 and dotted curves represent class 2. The x axis gives simulation time units, e.g., one second or one week.

ranging from 0 to 1.5 Mbps 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).

Both of the above classes are stable as shown in the simulation traces of Figure 3. The long-term average demand is the same for all three networks, and for the two classes. In the shown interval, the demand is such that capacity of argo is partitioned equally between the two classes. As expected, the higher quality class 2 service is more expensive, since it has higher provisioning coefficients – see 3(a). A similar situation occurs in maraca - see 3(c). In bongo, class 1 is more expensive because the demand is greater, which is also reflected in the larger share of capacity allocated to class 1 – see 3(b).

Consider now three inter-connected networks, with just one class, i.e., three brokers $\{1, 2, 3\}$. Let $r^2_1 = r^1_2 = x$, $r^3_1 = r^3_2 = y$, and $r^j_i = 0.99$ for all other pairs i, j. Figure 4 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 over-provisioning.

Neither can instability be simply attributed to the existence of "cycles" in the graph of the network. Figure 5 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 Figure 5(b), the right-hand side shows the allocations for traffic going from argo to maraca (dotted curve), and the bottleneck capacity in argo itself (solid curve). The instability is reflected

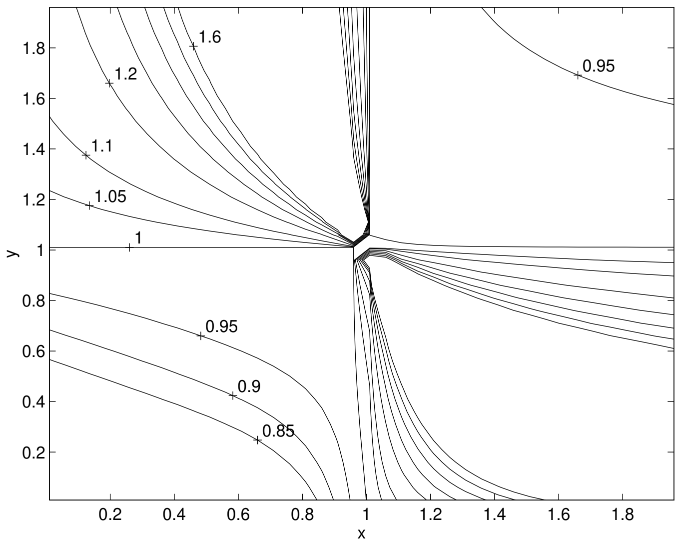

Fig. 4. Spectral radius as a function of inter-network provisioning coefficients, instability arises in the top left and bottom right quadrants.

in the volatility of the allocated capacities.

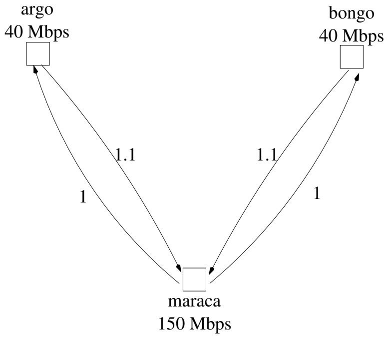

(a) Simulation topology ($\rho(\Phi) = 1.02$ *)*

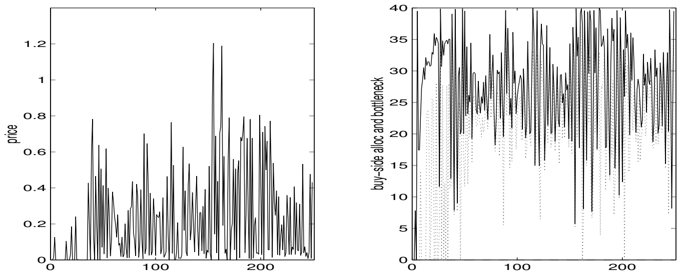

(b) Trace at argo

Fig. 5. 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 legend of x axis is the number of simulation time-units. The scenario is unstable as allocations do not converge.

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. Figure 6 illustrates this possibility, which we refer to as *"dis-peering"*. Here, we simulate the network shown in Figure 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

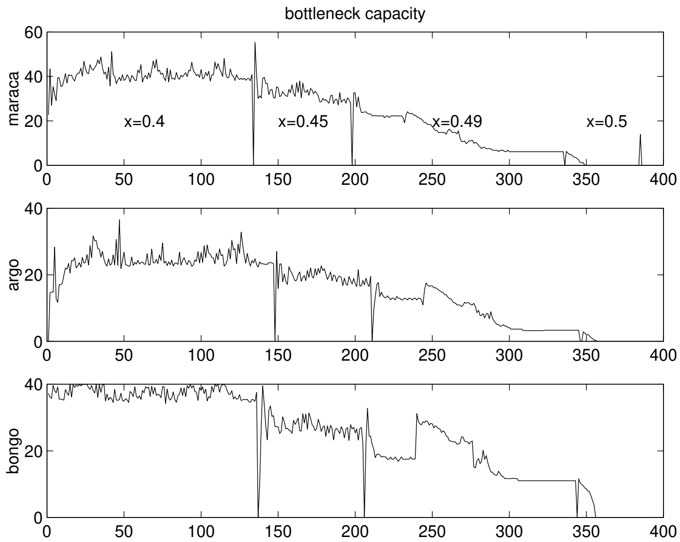

Fig. 6. "Dis-peering" effect, the legend of x axis is the number of simulation time-units.

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 II). 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, "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 [14]. 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 [15], [16], [17].

IV. Conclusion

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 to not 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 (8) 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.

References

[1] D. Clark, "Internet cost allocation and pricing," in *Internet Economics*, L. W. McKnight and J. P. Bailey, Eds. MIT Press, 1997.

[2] K. Nichols, V. Jacobson, and L. Zhang, "A two-bit differentiated services architecture for the internet," 1997, ftp://ftp.ee.lbl.gov/papers/dsarch.pdf.

[3] N. Semret, R. R.-F. Liao, A. T. Campbell, and A. A. Lazar, "Market pricing of differentiated internet services," in *IEEE/IFIP* 7^{th} *Int. Workshop on Quality of Service*, 1999.

[4] 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, 1999, http://comet.columbia.edu/~nemo/pub.html.

[5] M. T. T. Hsiao and A. A. Lazar, "An extension to Norton's equivalent," *Queueing Systems, Theory and Applications*, vol. 5, pp. 401–412, 1989.

[6] Z. Wang, "A case for proportional fair sharing," in *IEEE/IFIP* 6^{th} Int. Workshop on Quality of Service, Napa Valley, CA, May 1998.

[7] I. Stoica and H. Zhang, "LIRA: An approach for service differentiation in the internet," in *Proc. of NOSSDAV'98*, Cambridge, England, July 1998.

[8] National Research Council NRenaissance Committee, *Realizing the Information Future: The Internet and Beyond*, National Academy Press, Washington, D.C., 1994.

[9] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, "Assured forwarding PHB group," Internet RFC 2597, June 1999, ftp://ftp.isi.edu/in-notes/rfc2597.txt.

[10] V. Jacobson, K. Nichols, and K. Poduri, "An expedited forwarding PHB," Internet RFC 2598, June 1999, ftp://ftp.isi.edu/in-notes/rfc2598.txt.

[11] N. Semret, *Market Mechanisms for Network Resource Sharing*, Ph.D. thesis, Columbia University, 1999.

[12] D. Chazan and W. Miranker, "Chaotic relaxation," *Linear Algebra and its Applications*, vol. 2, 1969.

[13] A. D. Bovopoulos and A. A. Lazar, "Decentralized algorithms for optimal flow control," in *Proc. of the* 25^{th} *Allerton Conf. on Comm., Control and Computing*. University of Illinois at Urbana-Champaign, Oct. 1987.

[14] R. S. Benn, "The great peering debate," 1997, http://www.clark.net/pub/rbenn/debate.html.

[15] P. Baake and T. Wichmann, "On the economics of internet peering," *Netnomics*, vol. 1, no. 1, 1999.

[16] J. Gong and P. Sriganesh, "An economic analysis of network architectures," *IEEE Network*, pp. 18–21, March/April 1996.

[17] P. Sriganesh, "Internet cost structures and interconnection agreements," in *Internet Economics*, L. W. McKnight and J. P. Bailey, Eds. MIT Press, 1997.