Peering and Provisioning of Differentiated Internet Services

N. Semret, R. R.-F. Liao, A. T. Campbell and A. A. Lazar
Dept. of Electrical Engineering, Columbia University
New York, NY, 10027-6699, USA.
{nemo, liao, campbell, aurel}@comet.columbia.edu

March 1999

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 on 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. The model provides an explicit necessary and sufficient condition for the stability of the game, which serves as a proxy for the ability to maintain consistent SLAs. The model is 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. The simulations also illustrate conditions which lead to certain classes of service not being offered on an inter-network basis.

1 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 guarantee of Int-serv, 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 [8, 6]. A crucial aspect that has not been addressed systematically 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 on the edges, and/or significant violations of service quality in the core of the networks.

While the role of prices as essential resource allocation control signals has been established from the outset of diff-serv [4, 10], the precise development of pricing mechanisms is still at 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 compelling. With suitably designed rules, auctions can achieve efficient (value maximizing) allocations with minimal a priori information.

In [12], we showed that when the Progressive Second Price (PSP) auction introduced in [9] is used as the a diff-serv bandwidth market mechanism, it achieves economic objectives (incentive compatibility, and efficiency), while being realistic in the engineering sense (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. Section 2 introduces our game model, giving the wide-area pricing architecture, and the model for provisioning and differentiation of the services. Following this, in Section 3, we present our analytical and simulation results.

2 The Model

2.1 Distributed Market Framework

Our network model assumes that each network can be abstracted into a single bottleneck capacity (e.g. as a "Norton-equivalent" [7]). The capacity may be represented by an absolute amount of bandwidth, or some relative metrics like user share in the User-Share Differentiation proposal [15] or resource token in Location Independent Resource Accounting [14]. Large

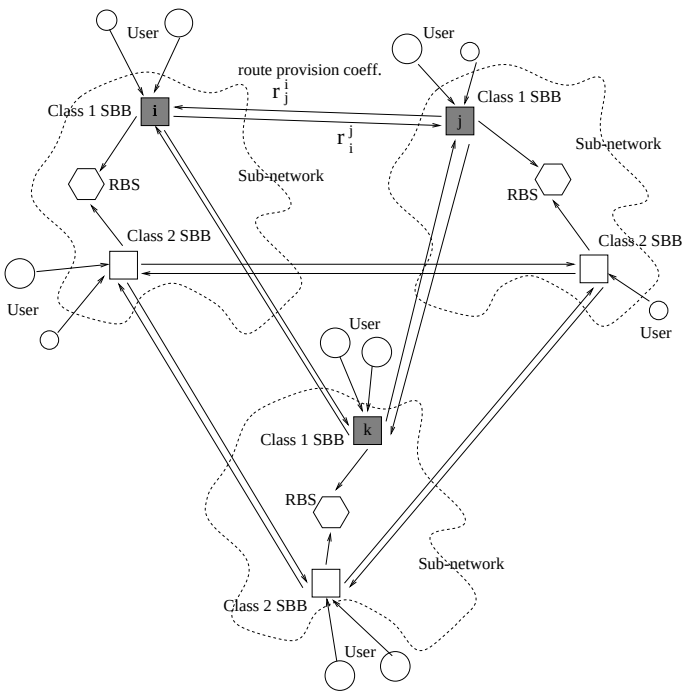

Figure 1: 2-tier auction pricing framework for diff-serv internet

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 architecture 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 [11]. If the RBS and multiple SBBs on the same network are not be 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 (web sites, intra/extranets, virtual private networks), rather than individual end-users.

2.2 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 [9], 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_i^j.

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 e constrained by the capacities that i can actually obtain. First, it must get enough bandwidth

[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.

from l, the RBS in its own network, to carry the total capacity it allocates to its customers, i.e.

$$\sum_{m \in \mathcal{I}} a_m^i \leq a_i^l. \tag{1}$$

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_i^j 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 2.3. Then, i must satisfy

$$r_i^j \sum_{m \in \mathcal{I}, m \neq j} a_m^i \leq a_i^j, \tag{2}$$

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

Define, for any allocation a,

$$e_i^j(a) \triangleq \frac{a_i^j}{r_i^j} + a_j^i, \tag{3}$$

$$e_i \triangleq \min_{j \neq i} e_i^j(a) \tag{4}$$

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 – see [12].

2.3 Differentiating Services

We do not explicitly consider the per-hop behaviors 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 service level agreement (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

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

for example services built on the diff-serv assured forwarding (AF) per-hop behaviors [6]. In this case, r_i^j is the expected[3] 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 $\sum_{j \neq l} r_i^j = 1$, where l is the RBS in i's network.[4]

- 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 [8]. 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 $r_i^j = 0, \forall j \neq i$.

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 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 [6]. 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. For the Virtual Leased Line service based on the EF PHB [8], the more conservative provisioning that is required can be achieved by providing for the worst case flows, i.e. *all* the traffic can flow to *any one* peer and still be satisfied.

Depending on the scheduling and buffer management algorithms used to provide the PHBs, some amount of over-provisioning may be required [8].

[3] Thus r represents aggregate flow patterns. r is measured over a time-scale slow enough to make quasi-static estimates which average out instantaneous micro-flows.

[4] Note 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.

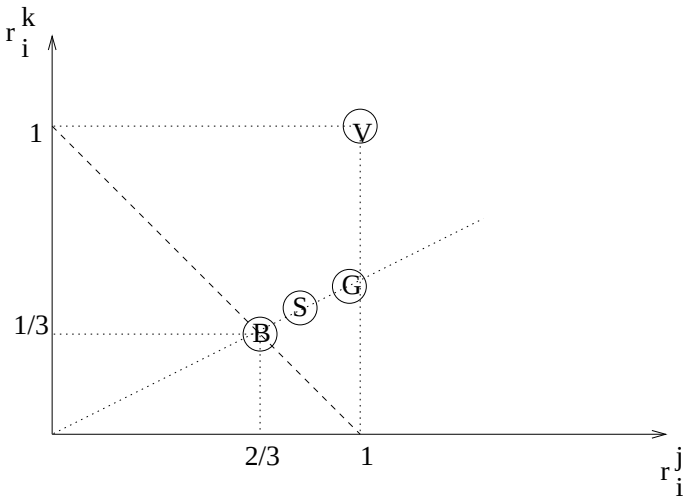

Figure 2: Inter-Network provisioning coefficients for different services: Olympic Gold, Silver and Bronze services, and the Virtual Leased Line service

These engineering needs can be represented in this model by simply factoring the over-provisioning into each coefficient of r.

3 Dynamic Provisioning of Differentiated Services

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 timescales.

On the demand side, the analysis of [9, 12] derives the optimal bidding strategies for users and brokers, and establishes the existence of an efficient (value maximizing) equilibrium point among buyers, when sellers are static. We now consider the supply game among brokers by itself. We will then use simulations to verify that the 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 $m = N+1, N+2, \ldots$. At the equilibrium point [12], the following conditions will hold $1 \leq i \leq N$:

$$\begin{aligned} e_i &= \sum_{j \in \mathcal{I}, j \neq i} a_j^i \\ a_i^j &= (e_i - a_j^i) r_i^j \end{aligned}$$

These conditions can be re-written in matrix form as

$$e = \Phi e + u, \tag{5}$$

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 [12], 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 (5).

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 [3]. □

Remark: Users can be viewed as external inputs driving a dynamic system, where the system dynamics are governed by the brokers: the system equation is then

$$e(t+1) = \Phi e(t) + u(t). \tag{6}$$

The convergence of the game is the exactly the notion of stability of the dynamic system (6).

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:

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

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 $\left(r_i^j\right)_{i,j\in\mathcal{I}}$ 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_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:

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

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

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 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_j^i that are set according to Table 1.

		buyer, class 1 SBBs			*buyer, class 2 SBBs*			*buyer, RBS'*		
		argo	bongo	maraca	argo	bongo	maraca	argo	bongo	maraca
seller class 1 SBBs	argo	0.3	0.2	0.1						
	bongo	0.2	0.3	0.1						
	maraca	0.5	0.5	0.8						
seller class 2 SBBs	argo				1.0	0.4	0.1			
	bongo				0.4	1.0	0.2			
	maraca				1.0	1.0	1.0			
seller RBS'	argo	1.0			1.0					
	bongo		1.0			1.0				
	maraca			1.0			1.0			

Table 1: Inter-Network provisioning coefficients: r_j^i (empty entries are zero)

The simulation parameters are given in Table 2. 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

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	

Table 2: Simulation parameters

distribution). 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 [9]. Thus, a class 1 user with a maximum capacity of 1.5Mbps will request a quantity 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. In each plot, the solid line represents class 1 and the dotted line represents class 2. 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 (a). A similar situation occurs in maraca - see (b). 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 (c).

Consider now three inter-connected networks, with just one class, i.e. three brokers $\{1,2,3\}$. Let $r_1^2 = r_2^1 = x$, $r_1^3 = r_2^3 = y$, and $r_i^j = 0.99$ 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.

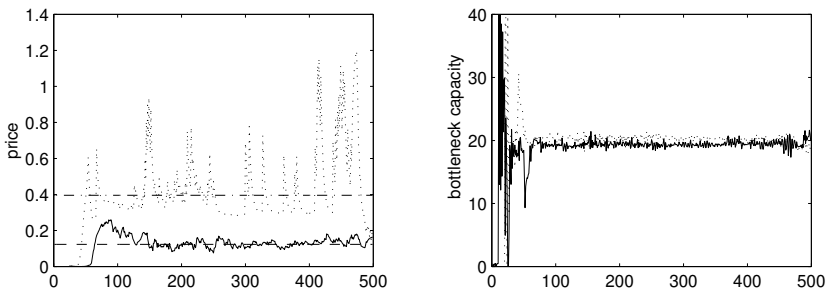

(a) argo trace

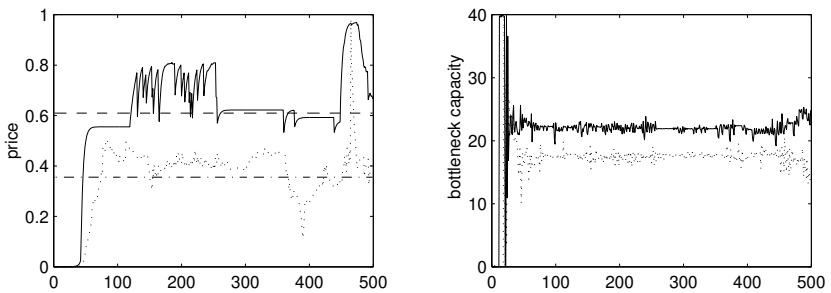

(b) bongo trace

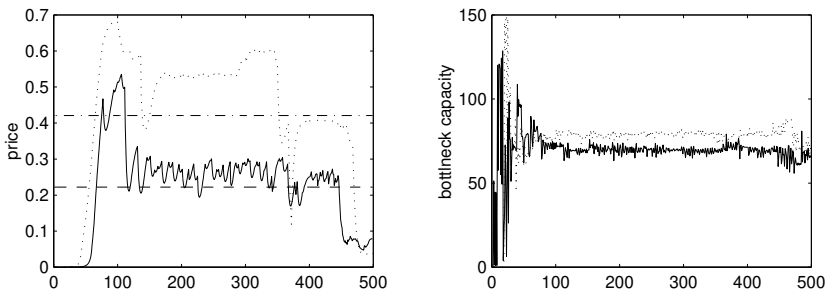

(c) maraca trace

Figure 3: Simulation of three networks with two stable classes

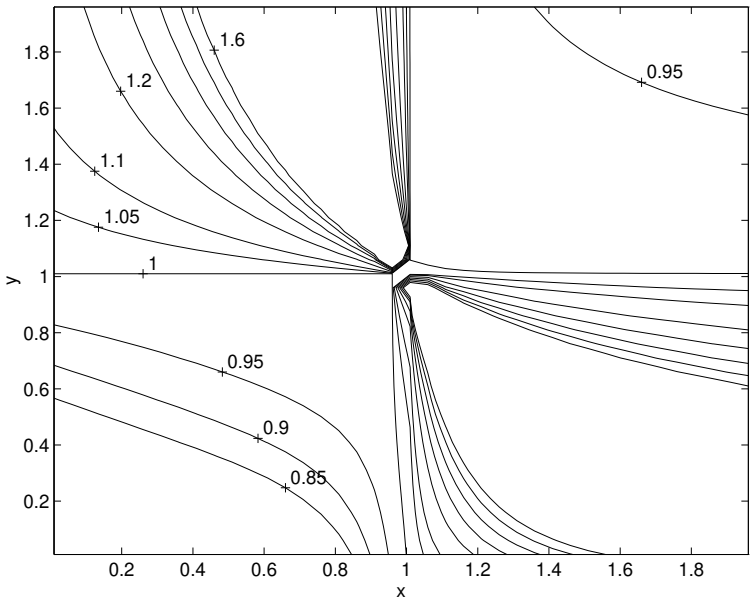

Figure 4: Spectral radius as a function of inter-network provisioning coefficients

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 (lower curve), and the bottleneck capacity in argo itself (upper 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. Figure 6 illustrates this possibility, which we refer to as *dis-peering*. Here, we simulate the network given above, 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 2). Thus, the reduction in bottlenecks is purely a result of the provisioning dy-

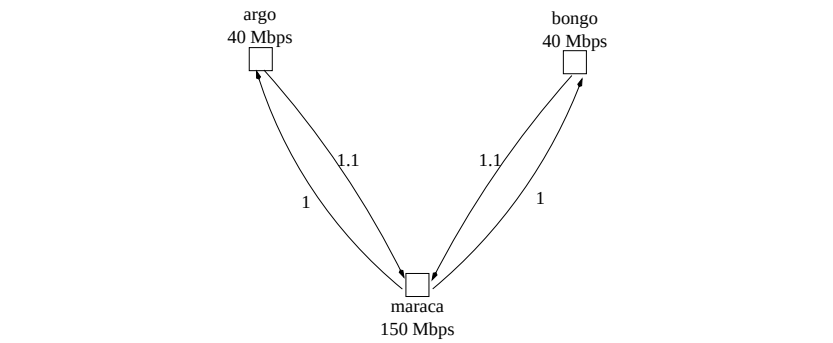

(a) Simulation topology ($\rho(\Phi) = 1.02$)

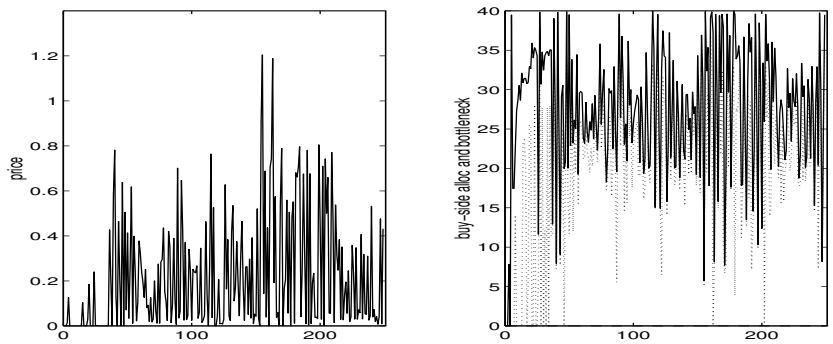

(b) Trace at argo

Figure 5: Simulation of one unstable class

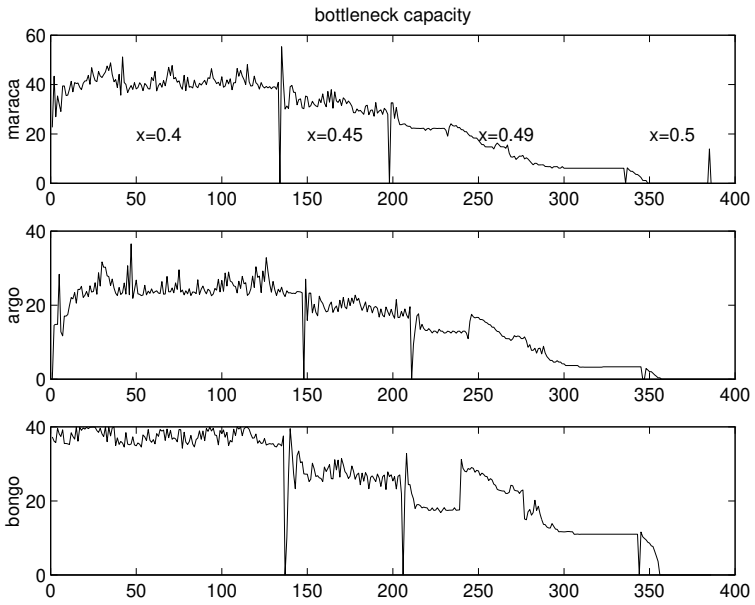

Figure 6: Dis-peering effect

namics, 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 [2]. 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 [1, 5, 13].

4 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 (6) suggests an interesting direction for future work. It may be possible to achieve certain wide-area network objectives, (e.g. stability or avoiding dispeering) 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] P. Baake and T. Wichmann. On the economics of internet peering. *Netnomics*, 1(1), 1999.

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

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

[4] D. Clark. Internet cost allocation and pricing. In L. W. McKnight and J. P. Bailey, editors, *Internet Economics*. MIT Press, 1997.

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

[6] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski. Assured forwarding PHB group. Internet Draft – work in progress, Feb 1999. http://www.ietf.org/internet-drafts/draft-ietf-diffserv-af-06.txt.

[7] M. T. T. Hsiao and A. A. Lazar. Optimal decentralized flow control of markovian queueing networks with multiple controllers. *Performance Evaluation*, 13(3), Dec 1991.

[8] V. Jacobson, K. Nichols, and K. Poduri. An expedited forwarding PHB. Internet Draft – work in progress, Nov 1998. http://www.ietf.org/internet-drafts/draft-ietf-diffserv-ef-01.txt.

[9] A. A. Lazar and N. Semret. Design, analysis and simulation of the progressive second price auction for network bandwidth sharing. Technical Report CU/CTR/TR 497-98-21, Columbia University, 1998. http://comet.columbia.edu/publications.

[10] 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.

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

[12] N. Semret, R.-F. Liao, A. T. Campbell, and A. A. Lazar. Market pricing of differentiated internet services. Technical Report CU/CTR/TR 503-98-37, Columbia University, 1998. http://comet.columbia.edu/publications.

[13] P. Sriganesh. Internet cost structures and interconnection agreements. In L. W. McKnight and J. P. Bailey, editors, *Internet Economics*. MIT Press, 1997.

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

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