On the Rise and Fall of ISPs

Ashlesh Sharma
New York University

Nathan Silberman
New York University

Lakshminarayanan Subramanian
New York University

Nicholas Economides
New York University

The Internet topology has witnessed significant changes over the years with the rise and fall of several Internet Service Providers (ISP). In this paper, we propose a new economic model that can aid in understanding the evolution of the Internet topology and provide insight into why certain ISPs fail and others succeed. Our economic model is motivated by the Cournot model for characterizing oligopolistic markets. We model the Internet topology as a conglomeration of Cournot markets across different geographic regions comprising of regional markets within each geographic region and transit markets across geographic regions. By analyzing the Nash equilibrium of the overall system, we characterize a simple relationship between the Nash price, demand, the number of ISPs and fraction of traffic exchange across regions. Our economic model is powerful enough to provide a simplified characterization of the aggregate evolution of the Internet topology within and across geographic regions without the need for capturing individual variations across each ISP. Based on this model, we show evidence that existing bandwidth pricing trends are in striking contrast with the expected Nash equilibrium behavior thereby resulting in the rise and fall of ISPs. We also corroborate the model based on analyzing Internet topology evolution from 2002 to 2008.

1. INTRODUCTION

In the past decade, we have witnessed dramatic changes in the evolution of the Internet's topology, especially at the level of top-tier ISPs. These top-tier ISPs, otherwise referred to as tier-1 or tier-2 ISPs, are the primary suppliers in the market for Internet access. During this growth period, several tier-1 ISPs have either merged or been bought by other ISPs while other new tier-1 ISPs have emerged. For example, AT&T and Genuity (formerly BBN) sold their backbone networks to SBC and Level 3 communications [1]. Similarly, MCI Communications and Worldcom merged and later were bought over by Verizon after Verizon filed for bankruptcy [2]. At the other end of the spectrum, several globally operational ISPs have emerged in Asia such as KDDI, Japan. We observe a strong geographic correspondence to the rise and fall of ISPs with many ISPs in US and Europe having failed while several others in Asia and South America have risen to prominence.

Why do certain ISPs succeed and others fail? In this paper, we propose a simple economic model that provides a mechanism for analyzing the evolution of the Internet topology. Our model is based on the Cournot model for oligopolistic markets. In Cournot competitive markets, the product being sold is homogeneous and each firm chooses a price based only on the quantity of the product they produce. While the assumptions in Cournot competition do not hold across the global bandwidth markets, we model the Internet topology as a collection of regional and transit Cournot markets. In other words, we split the Internet topology into geographic regions and apply a local Cournot competition model within each region and transit Cournot market for transit ISPs that interconnect regions. In our economic model, each ISP is a player in the bandwidth market either within a specific geographic region or a transit player across regions. Based on limited pricing data available online, we find that bandwidth prices across ISPs within a geographic region are comparable; hence, assuming separate Cournot markets within and across regions is a reasonable assumption.

Our economic model provides a simple yet powerful way of analyzing Internet topology evolution due to three important factors. First, the Cournot model enables us to analyze the Nash equilibrium of a market in the aggregate without analyzing the individual strategies and variations across each player. In the Internet topology case, this is powerful because, it can provide an insight of how regional and transit bandwidth prices should vary as a function of variations in aggregate quantity and number of ISPs. Second, while the Internet topology as a whole is tightly coupled, our economic model allows us to analyze each regional and transit market in isolation under Nash Equilibrium. Third, the final set of Nash Equilibrium equations we arrive in our system are simple and easy to analyze involving few parameters.

We draw many important insights from our model. First, for a given demand, as the fraction of traffic from a geographic region increases, regional bandwidth prices should *increase* and *not decrease as what is happening today*; hence, as Internet traffic becomes more global, overall prices should increase. Second, if local bandwidth prices decrease with increasing traffic outflow, the number of regional ISPs sustainable at Nash equilibrium decreases. We establish a Nash equilibrium relationship between the fraction of traffic outflow and the number of ISPs within a region. Third, in our model, we show that the Nash price for bandwidth connectivity is the sum total of fractional Nash price for transit connectivity and the Nash price for regional connectivity (assuming a local Cournot market for the same demand). Hence, in Nash equilibrium, Internet pricing should be similar to pricing in telephone networks where the price of an international call is the sum total of the local and the international transit cost. Finally, our model not only helps in explaining the existing economic topology of the internet, but also provides insights into predicting future market behavior.

We corroborate the theory using a detailed analysis of the evolution of the Internet topology from 2002 to 2008. For this analysis, we use the 5-tier characterization of the Internet topology along with the inter ISP relationships as characterized in prior work in Subramanian *et al* [19]. We also use a wide variety of Internet sources to estimate quantity, price, fractional traffic and number of ISPs. While the topology characterization and our parameter estimation techniques are not very precise, it does provide intuition in analyzing the aggregate evolution of the Internet topology.

2. RELATED WORK

There have been several studies which have mapped and analyzed the evolution of the Internet topology at both the AS-level and the router-level using a variety of techniques. Traceroute analysis have often been used in efforts to map router topologies [13, 3, 7, 14]). The commonly used approach to infer AS-level topologies is to analyze BGP routing tables at the Internet routers.

One of the commonly used models for Internet topology evolution is the power law model by Michalis Faloutsos et al. [9]. There have been several followup studies to this work which have presented contrasting variants of the power law model [20, 18]. While these models are useful from understanding the graph structure and its evolution, these models do not consider the financial or economic relationship between the different players (ISPs and end-users) in the Internet.

There have been few works which have analyzed the Internet evolution from an economic perspective. Freiden [11, 10] discusses Internet balkanization and the ISP peer relationships and its implications. The author argues why growing service demand, congestion at quasi-public interconnection sites and commercialization of the industry have motivated major ISPs to pursue more reliable, quasi-private network interconnection, as opposed to blind collaborative behavior between the entities. Bailey [6] surveys the economics of Internet interconnection and infers that large networks or ISPs form bilateral relationships where coordination costs are low and there is little chance for opportunism; and smaller networks or ISPs are better off forming cooperative relationships.

Recent efforts have attempted to model these complex commercial interactions between AS entities using game theory. Recently Ma et al. [15] used *Shapely value*, a concept from coalition games to model the cooperative behavior of ISPs. Their model curtails maximization of individual ISP profits, but maximizes the aggregate profit of all ISPs put together. This prevents selfish routing strategies [16] among ISPs and provides a aggregate network profit, which in turn encourages ISP connectivity and reduces balkanization. Cao et al. [8] and Shakkottai et al. [17] apply game theory to analyze Internet pricing and the economics of ISP interconnection. In particular, Cao et al. show that modeling the interactions between AS entities as a cooperative game leads to better outcomes for both the ISPs and the users [8]. They also prove the existance of a Nash-equilibrium point, where two ISPs would not move out without cooperation. Shakkottai et al. show that in the real world, interactions between ISPs are often non-cooperative. They show that these interactions can be modeled as a multi-stage game, where ISPs are linked together through transit-ISPs.

Our work builds on these insights and uses an economic model to understand and predict the large-scale behavior of ISPs across regions. Critical to our economic analysis is to leverage knowledge about commercial relationships across ISPs. Prior work by Gao et al [12] and Subramanian et al. [19] have proposed different algorithms for inferring AS relationships among three categories: provider-customer, peer-peer and sibling relationships. There have been many followup works which have improved these relationship inference algorithms. In this paper, we leverage these relationship inference mechanisms to analyze the economic rise and fall of ISPs.

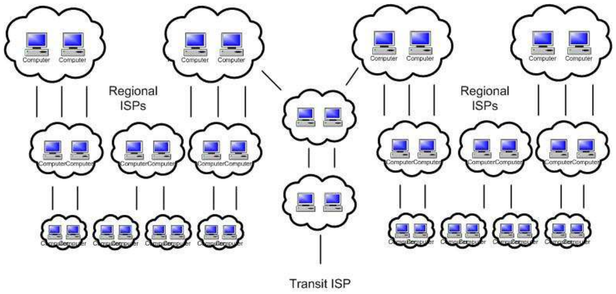

Figure 1: An illustration of Internet topology: Regional ISPs with tier structure and Transit ISPs connecting them

3. INTERNET ECONOMIC MODEL

In this section, we provide a high level overview of our Internet Economic Model. and follow up with more details on the model in Sections 4 and 5.

The Internet is a collection of Autonomous Systems (AS) where each AS is either an Internet Service Provider (ISP) or a customer stub network (university, company etc) representing the consumer such as a university or a company. Each autonomous system is associated with a AS number and AS's use the Border Gateway Protocol to determine routing paths. BGP is a policy based routing protocol where each AS uses its own set of routing policies to choose routes.

The Internet topology has three common types of relationships: customer-provider, peer-peer and transit. A transit relationship can be one of two kinds: (a) a provider-customer relationship between a large ISP and a small ISP; (b) an ISP that provides transit connectivity across geographic regions (trans-pacific, trans-Atlantic). Based on these inter-ISP relationships, the Internet topology has an inherent hierarchical structure where tier-1 ISPs represent the most important inter-continental ISPs, tier-2 represents the national ISPs, tier-3 and tier-4 ISPs represents the regional ISPs and the consumers (stub networks) form the lowest portion of this hierarchy. Prior work by Subramanian *et al.* [19] provides a 5-tier characterization of the Internet hierarchy based on inferring AS relationship information from multiple BGP routing tables. Our model and our evaluations leverage this 5-tier characterization.

Based on the existing Internet hierarchical structure, we use the following simplified model for the Internet topology illustrated in Figure 1. We characterize the Internet into different geographic regions and associate each geographic region with a set of regional ISPs and a set of transit ISPs that interconnect geographic regions. In practice, if a large ISP has both a regional and a transit presence (such as ATT, Sprint), we model these ISPs as separate ISPs for the regional and transit cases; in the existing Internet, these ISPs register the regional and transit ISPs as separate entities though they may be owned by the same parent entity. A transit relationship, in our model, is a customer-provider relationship between a transit ISP and a regional ISP; hence the regional ISP has to pay the transit ISP for connectivity. We assume that in a peer-peer relationship, peers do not pay each other for routing traffic between themselves. Hence, if two regional ISPs or two transit ISPs share a peering relationship, they do not need to pay each other for the mutual traffic exchanged. The set of customers in our model are distributed across geographic regions and they represent the "quantity of demand" in each geographic region. While customers pay their regional ISPs and the regional ISPs in turn pay the transit ISPs.

Our economic model for analyzing Internet topology evolution is based on the Cournot model for oligopolistic markets. We assume that each geographic region represents a regional Cournot market and the collection of transit ISPs that interconnect a set of regions as a transit Cournot market. The Cournot Competition model is an economic model that makes the following basic assumptions:

1. The product being sold is homogeneous across firms
2. Every firm in the market adjusts the quantity of the product they produce simultaneously
3. The price of the good is determined by the cumulative output of all the firms

The Cournot model is a natural fit for analyzing regional and transit markets. While bandwidth prices may vary significantly across regions, the bandwidth prices are roughly comparable within a geographic region. If there is a significant price difference, then the ISP charging lower prices (assuming it is profitable) has the capability to expand coverage and get a larger portion of the market. In addition, given that regional ISPs peer with each other at zero cost (assuming these ISPs are in equal standing), ISPs can route between their customers without additional cost. Hence, in the regional case, the competition is between how the demand is distributed across competing ISPs. The same case extends for the transit market where the regional ISPs represent the customers of the transit ISPs.

The Cournot model represents a powerful way to analyze Internet topology evolution since it can enable us in analyzing the Nash equilibrium of the system at an aggregate level without the individual parameters or strategies of each player. In other words, to analyze the Nash equilibrium of the system, we do not require detailed values of each ISP's demand, bandwidth and other statistics; rather we can predict the behavior of ISPs and their interaction using only the aggregated values from all the ISPs.

In formal terms, given N firms in a market, each firm's produces a quantity q_i and a price function $p_i(q_i)$. Additionally, the marginal cost of production, c is the same for all firms. When the market is at equilibrium, each firms price is equal and the equilibrium price function P is be given by $P(q_1 + ... + q_N)$. The only tool each firm has to maximize their profits is adjusting the quantity q_i being produced.

As we describe later in Section 4, in the general case of N firms, the Nash Equilibrium is characterized by the following equation:

$$p'(q)\frac{q}{N} + p(q) = c_0$$

where $p(q)$ is the variation of price with aggregate demand q and c_0 is the unit cost incurred by the ISP per unit demand. Hence, the Cournot model characterizes the Nash equilibrium of N firms purely based on the aggregated state or aggregate demand as opposed to the individual states of each firm. This aggregate characterization is the main reason why we picked the Cournot model as a base point to design an economic model for inter-ISP interaction. Note that the aggregate characterization holds even if the cost function $c(q)$ is not linear. As long as the cost function is uniform across all firms, we can replace c_0 in the equation with $c'(q)$.

4. REGIONAL MODEL

We use the simple N-party Cournot model for our basic regional case as illustrated in Figure 2. Consider the scenario where the entire network comprises of N ISPs co-located in a region and provide connectivity to M stub networks where $M >> N$. In addition, we assume that all the N ISPs peer with each other and peers exchange traffic at no additional cost.

We make some assumptions about the interaction of these ISPs in the regional setting. Let $q_i > 0$ be the quantity produced per ISP I_i or the traffic demand for the ISP I_i, $q = \sum_{i=1}^{i=N} q_i$ be the net demand across all I_i^s. In the simple case where all stub networks generate comparable amounts of traffic, q is a function of M, the number of stub networks in the region. Let $c_0 > 0$ be the unit cost required to produce q_i (we assume that c_0 is the same for each I_i, since the cost of "producing" (routing) traffic q_i for each I_i is relatively the same). Let $p(q)$ be the price function or the market price charged by each I_i. Also, we assume that $p()$ is differentiable with $p'(q) < 0$ at all $q \geq 0$.

We will briefly illustrate the Cournot model for the 2-ISP case. To find the Nash equilibrium of this model, we have to start with the profit maximization function. Consider the case of two ISP's, where each ISP tries to maximize profit, which is given by the profit maximization function,

$$\max_{q_i \geq 0}\{p(q_i + q_j)q_i - c_0 q_i\}$$

where q_i,q_j is the quantity of I_i,I_j respectively. Then the optimal quantity choice for ISP I_i has to satisfy the first order function,

$$p'(q_i + q_j)q_i + p(q_i + q_j) = c_0$$

If (q_1^*, q_2^*) is a Nash equilibrium for such a system, then it has to satisfy these equations,

$$p'(q_1^* + q_2^*)q_1^* + p(q_1^* + q_2^*) = c_0$$

and

$$p'(q_1^* + q_2^*)q_2^* + p(q_1^* + q_2^*) = c_0$$

Adding these two equations, we have,

$$p'(q_1^* + q_2^*)(\frac{q_1^* + q_2^*}{2}) + p(q_1^* + q_2^*) = c_0$$

This equilibrium equation can be extended to N ISPs that have identical cost and price functions as,

$$p'(q^*)(\frac{q^*}{N}) + p(q^*) = c_0$$

To further simplify this equation, we can write $p'(q)$ as $p'(t)/q'(t)$. In general, the price variation as a function of aggregate demand is not known. However, the price variation across time and the quantity variation across time are easier to estimate. Given a certain time period T, $p'(t)/q'(t)$ can be crudely approximated as $(p(t+T) - p(t))/(q(t+T) - q(t))$. In addition, $p(q) - c_0$ represents the profit per unit demand. Given this profit per customer, the expected variation in price at Nash equilibrium can be estimated as a function of N, $q(t)$ and $q(t+T) - q(t)$. Note that this equation is independent of how the price function varies with the aggregate demand q.

4.1 Implications

The significance of this Nash equilibrium stems from the fact that the price function is independent of individual quantities q_i of any one ISP I_i, but dependent on the aggregate quantity $q = \sum_{i=1}^{i=N} q_i$. At one extreme, if $N \rightarrow \infty$, then $p(q^*) = c_0$; the market price equals cost of the quantity and the ISPs make zero profit. On the other extreme, if N = 1, then equation reduces to the monopolistic first order function, $p'(q^*)q^* + p(q^*) = c_0$; providing a monopoly of the market for the ISP. The evolution of ISPs in the regional model has remained between these two extreme market scenarios.

The market tries to be in Nash equilibrium, since any deviation from this equilibrium would reduce the profit of some of the ISPs while increasing the profits of other ISPs. As bandwidth/data demand grows, the aggregate quantity q increases and the price $p(q)$ reduces to a marginal amount. The profit function of each ISP would depend on the production of quantity q_i. This behavior affects the smaller ISPs: ISPs that produce relatively less q_i, as the reduction in market price might not be balanced by increased demand for q_i; due to this situation, the smaller ISPs profit might be marginal and would cause it to quit the market or be taken over by larger ISPs. The Cournot oligopoly model captures this competitive market strategy of ISPs on the regional level.

4.2 Hierarchy within a region

In reality, every geographic region may have a hierarchy of regional ISPs. The above model can be directly applied at the top-tier ISPs within the region (these are not tier-1 ISPs). If all the ISPs within the region are dependent on N ISPs for connectivity within and across regions, then the above model can be extended to these N ISPs. Outside of these N ISPs, if there are other regional ISPs which depend on these N ISPs for regional connectivity, we can analyze such a hierarchical structure using a two-level hierarchy with N_1 regional ISPs at the top level and N_2 ISPs which depend on the N_1 ISPs.

The corresponding Nash equations are:

$$p_1'(q^*)(\frac{q^*}{N_1}) + p_1(q^*) = c_1$$

$$p_2'(q^*)(\frac{q^*}{N_2}) + p_2(q^*) = c_2$$

The following constraints should hold to make sense for the N_2 regional ISPs to exist. First c_2 should be significantly different from c_1; otherwise the N_1 ISPs can capture the local market (unless regulatory laws interfere). Second, p_2 should be larger than p_1 since the N_2 regional ISPs pay the price difference per unit demand to the N_1 ISPs. In our evaluations, we primarily analyze the regional case as a single hierarchy in our evaluations due to lack of fine grained pricing and demand data to analyze the lower tiers of the regional hierarchy.

5. TRANSIT ECONOMIC MODEL

In this section we characterize the internet beyond the Regional Model into a *Transit Model*, where regional ISPs are connected with each other using *transit* ISPs. The transit ISPs act as interconnect hubs between ISPs of different geographic locations.

A typical transit model is shown in figure 2. Let f_i be the fraction of traffic flow from ISPs N_i of each geographic region to the transit ISPs k. We assume that peer-peer traffic between transit ISPs are cost-free: the ISPs do not pay each other for the traffic routing between themselves. The Cournot model can be applied to each region taking into consideration the transit ISPs to which they route the traffic. (Here, we consider only the economics of the regional ISPs and not the transit ISPs.)

In the scenario, the profit maximization function of i^{th} ISP in a region is given by,

$$\max\{p_l(q_i)\alpha_i - c_0\alpha - p_t(q_i)\beta_i\} \tag{1}$$

where p_l is the price function in the region (local), p_t is the transit price given to the transit ISPs by the regional ISPs, α is the demand per ISP in a region and $\sum_i \alpha_i = q_1$, where q_1 is the total demand in a region, β is the demand per ISP in a region with fraction of traffic f_i from the transit ISP and $\sum_i \beta_i = f_1 q_1$, where f_1 is the total fraction of traffic to/from transit ISP, c_0 is the cost per unit quantity.

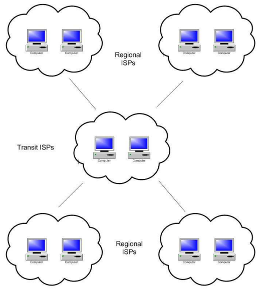

Figure 2: **Regional ISPs and Transit ISPs**

To compute the Nash equilibrium of the system, we first have to compute the first order profit maximization function, which can be obtained by differentiating (1)

$$\frac{\partial p_l}{\partial q_i}\alpha_i + p_l\frac{\partial \alpha_i}{\partial q_i} - c_0\frac{\partial \alpha_i}{\partial q_i} - \frac{\partial p_t}{\partial q_i}\beta_i - p_t\frac{\partial \beta_i}{\partial q_i} = 0$$

We know that, $\frac{\partial \alpha_i}{\partial q_i} = 1$, $\frac{\partial \beta_i}{\partial q_i} = f_1$, Adding over all i ISPs,

$$\sum_i p_l'\alpha_i + p_l - c_0 - p_t'\beta_i - p_t f = 0$$

$$p_l' q_i + N_1 p_l - N_1 c_0 - p_t' f_i q_i - N_1 p_t f = 0$$

where N_1 is the total number of ISPs in a region. Therefore the first order joint profit maximization is,

$$(p_l' - p_t' f_1)\frac{q_1}{N_1} + (p_l - p_t f_1) = c_0 \tag{2}$$

The price values in this equation is a function over q_1. So, Eq. (2) can be precisely written as,

$$(p_l'(q_1) - p_t'(q_1) f_1)\frac{q_1}{N_1} + (p_l(q_1) - p_t(q_1) f_1) = c_0$$

Let

$$\lambda(q_1) = p_l(q_1) - f_1 p_t(q_1) \tag{3}$$

so the equation reduces to,

$$\lambda'(q_1)\frac{q_1}{N_1} + \lambda(q_1) = c_0 \tag{4}$$

The Nash equilibrium of this system can be written as,

$$\lambda'(q_1^*)\frac{q_1^*}{N_1} + \lambda(q_1^*) = c_0 \tag{5}$$

This equation has the same significance as the Nash equilibrium of the Regional Model. As the number of regional ISPs $N_1 \rightarrow \infty$, $\lambda(q_1^*) = c_0 \Rightarrow p_l = c_0 + f_1 p_t$. The market price equals the cost price plus the transit price and there is no profit for the ISP. On the other hand, if $N_1 = 1$, then the equation characterizes a monopolistic market model.

To understand the effect of transit ISPs on the model, we show how transit ISP prices affect the ISPs in a geographic region. The regional price p_l is affected by transit prices and this is shown by

(3) by,

$$p_l(q_1^*) = \lambda(q_1^*) + f_1 p_t(q_1^*) \tag{6}$$

where $p_l(q_1^*)$ is the price function in the region, $\lambda(q_1^*)$ is the price paid by the customers in a region and $f_1 p_t(q_1^*)$ is the transit price based on the fraction of traffic between the regional ISP and the transit ISP.

As bandwidth demand increases, the transit ISPs charge higher prices to the regional ISPs. The regional ISPs charge less or marginal prices to the customers in the region due to competition with other regional ISPs. This reduces their profit and smaller regional ISPs go bankrupt or they are taken over by larger regional ISPs. As $f_1 p_t(q_1^*)$ increases and $\lambda(q_1^*)$ decreases, $p_l(q_1^*)$ will be comparable (or almost equal) to $f_1 p_t(q_1^*)$. Also, as demand for bandwidth increases the number of ISPs increase, which causes each ISP to have reduced profits. This is a deviation from the Nash equilibrium of the transit model. If demand or quantity q_1 increases, transit price p_t increases, customer price $\lambda(q_1^*)$ reduces due to regional ISP competition, then the only way for the market to be in Nash equilibrium is to reduce the number of ISPs N_1 in the region. By reducing the number of ISPs each regional ISP has a larger chunk of bandwith (or demand q_1), which in turn increases the profit.

5.1 k-region case

In a duopoly (k = 2) Cournot model, if we consider an inverse demand function $p(q) = a - bq$, with $a > c \geq 0$ and $b > 0$, then in Nash equilibrium, $q*$ of each firm is $\frac{a-c}{3b}$. To include regional ISPs and transit ISPs in the model, we extend the two firm Cournot model to multiple firm Cournot model. In a multiple firm setting, the inverse demand function of a firm changes to

$$p = a - b(q_2 + q_3 + \ldots + q_N) - bq_1$$

Let $q_{-1} = q_2 + q_3 + \ldots + q_N$, therefore

$$p = a - bq_{-1} - bq_1$$

The profit maximization of a firm now depends on the other firms' output. As the number of firms increases the profit of each firm comes closer to the cost function c and attains only marginal profit. The cost function can be written as,

$$c = a - bq_{-1} - 2bq_1$$

Therefore at Nash equilibrium, q_1^* of a firm is,

$$q_1^* = \frac{a-c}{2b} - \frac{q_{-1}}{2}$$

Since the firms are in equilibrium, q_{-1} can be approximated as $(N-1)q_1*$. So, q_1* can be computed as,

$$q_1^* = \frac{a-c}{N+1}b \tag{7}$$

Therefore for N firms the joint output quantity q^* at Nash equilibrium is,

$$q^* = (\frac{a-c}{b})(\frac{N}{N+1}) \tag{8}$$

The joint profit function is,

$$p^* = a - bq^* = \frac{a+Nc}{N+1} \tag{9}$$

So, the profit per firm p_1^* is given by,

$$p_1^* = (p^* - c)q_1^* = \frac{(a-c)^2}{(N+1)^2}b \tag{10}$$

We apply this Cournot model of multiple firms to regional ISPs and transit ISPs and provide a relationship among them based on the quantity produced, the fraction of traffic flow between them and the number of ISPs in each of the regions. For a two region-one transit model, a conservation equation can be written as,

$$f_1 q_1 + f_2 q_2 = f_3 q_4$$

where f_1,f_2 is the fraction of traffic flowing from the two regions to the transit ISPs, f_3 is the incoming fraction of traffic at the transit region. Generalizing it to n regions,

$$\sum_i^n f_i q_i = f_t q_t \tag{11}$$

where f_t is the aggregate fraction of incoming traffic to the transit ISPs. In Nash equilibrium condition, the equation can be written as,

$$\sum_i^n f_i(\frac{a-c}{b})(\frac{N_i}{N_i+1}) = f_t(\frac{a-c}{b})(\frac{K_i}{K_i+1}) \tag{12}$$

where K_i is the number of transit ISPs. At Nash equilibrium, we also show that the Nash price per customer p_{nash} can be written as: $p_{nash} = p_{region} + fp_{transit}$, where p_{region} is the Nash price within the geographic region assuming no transit traffic and $p_{transit}$ is the transit price at Nash equilibrium.

We draw several important insights from these results. First, at Nash equilibrium, the Nash price that a regional ISP should charge is the sum of the Nash price for local connectivity and the fractional Nash price for transit connectivity. In other words, if the fraction of external traffic is 0, then the Nash price for connectivity is the price for local connectivity as drawn from the regional model. However, as the fraction f increases, the Nash price should correspondingly include the fractional price for transit connectivity. This pricing structure is present in existing telephone networks where the cost of any call is dictated by the destination of a call.

Unfortunately, the existing Internet pricing structure does not follow the Nash equilibrium pricing structure. For a given demand, as the fraction of traffic from a geographic region increases, we require regional bandwidth prices to *increase*. However, existing bandwidth prices have steeply decreased. The only way to retain Nash equilibrium is for N to decrease; hence some regional ISPs may need to fail at Nash equilibrium. Hence, for a given demand q, as Internet traffic becomes more global, overall bandwidth prices should increase or some regional ISPs have to fail. If q significantly increases, N can correspondingly increase.

The Nash equilibrium also sets up a relationship between N_i and f_i within each region. When all other parameters are kept a constant, any fluctuations in f_i have corresponding implications on N_i. Hence, in the absence of external fluctuations, $f_i N_i/(N_i+1)$ exhibits roughly constant behavior over time. We show this in our evaluations using real world data.

Finally, we note that the Nash equilibrium in our model can be characterized by a set of equations with very few aggregate parameters and no individual parameters corresponding to specific players. The larger the number of parameters, the harder it becomes to plug in values and perform meaningful empirical analysis from such a model.

6. PRELIMINARY RESULTS

In practice, it is very hard to obtain accurate information about traffic demand, prices, price fluctuations,fraction of outgoing traffic per geographic region and number of competing ISPs in the same

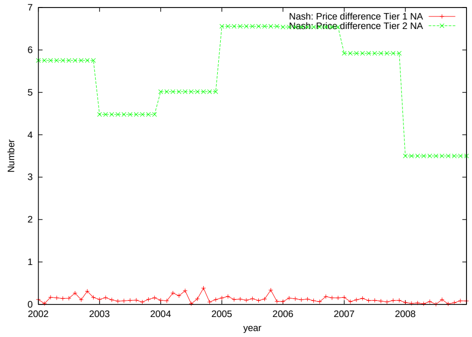

Figure 3: The Nash price difference in North America(NA) for Tier1 and Tier2&3

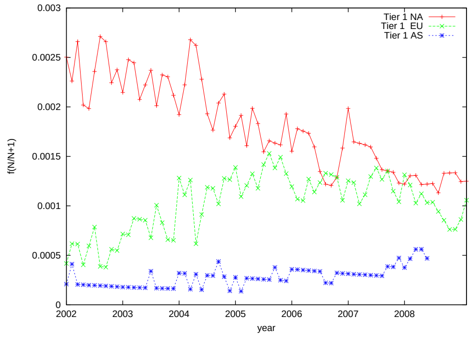

Figure 4: Variation of Transit function for Tier1 across regions

hierarchy within a region. We have tried our best to approximately infer these parameters from a variety of sources and apply them in the model. To measure how ISPs were performing over time, we used a heuristic developed by Subramanian et al. [19] to classify autonomous systems into a 5-tier hierarchy. Since the generated AS model is a heuristic, the differences between tiers are not always clear. Consequently, instead of focusing on exactly which tier a specific AS is ranked, we will be focusing on the relative movement of the different ASes.

Our data model partitions ISPs and IP prefixes by geographic region based on ASGeo Netlantis database [4]. Traffic is then categorized by whether its source and destination are in the same region, producing internal traffic, or whether the source and destination are in different regions, producing transit traffic.

We estimate the number of ASes per region i, N_i and these N_i are further grouped into various tiers. The aggregate traffic per region q_i is estimated by adding all the tier 5 ASes in that region. The fraction of traffic per region per tier f_i is estimated as $\frac{N_{ti}}{\sum_{ti} N_{ti}}$, where N_{ti} is the number of ASes per tier in that region. In the equation $p'(q^*)(\frac{q^*}{N}) + p(q^*) = c_0$, $p'(q^*)$ can also be written as $\frac{p'(t)}{q'(t)}$ and we estimate this parameter using variation in price and quantity observed in quantized intervals (on a monthly basis). The price variation data is taken from OECD's [5] broadband pricing data. Also, we use a linear price function to approximate the prices for years where no data is available.

6.1 Nash Equillibrium Analysis

We present our initial analysis of the Nash equillibrium condition for the existing Internet by using estimated values of the individual parameters.

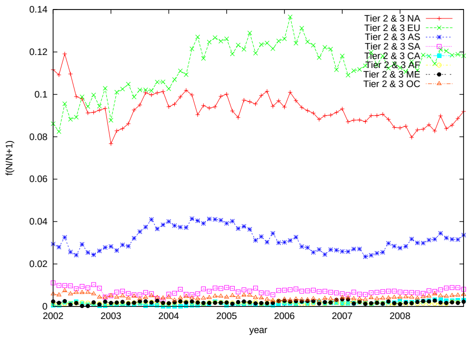

Figure 5: Variation of Transit function for Tier2&3 across regions

Variation of $p'(t)$: First, we analyze the variation of $p'(t)$ across North America for Tier 1,2,3 providers. Here, we approximate p'(t) as the expected relative difference in price over one year. Figure 3 shows the variation of price difference across time in North America. The Tier 1, the expected price difference values vary only slightly. This shows the Tier 1 ISPs tend to have same profit margin per customer and their only way to increase profit would be to increase the quantity q, thereby serving lot more lower level ISPs. Hence, at Nash equillibrium the expected price difference for Tier-1 providers is relatively small. While transit prices are not readily available online, the small set of published numbers online appear to be relatively constant. We find that the price differences for Tier 2 & 3 correlate with the OECD broadband price variations published online [5].

Transit Model: To analyze the Nash equillibrium equation for the transit model and computed the variation of $\sum_i^n f_i(\frac{a-c}{b})(\frac{N_i}{N_i+1})$ with respect to Tier 1 and Tier 2 & 3 across regions. Figure 4 shows the variation of transit function $\sum_i^n f_i(\frac{a-c}{b})(\frac{N_i}{N_i+1})$ with respect to Tier 1 across regions. The variation is not much and this result holds good with the Transit model for multiple ISPs. Hence, this validates our Nash Equillibrium equation that the quantity $f_i N_i/(N_i+1)$ is roughly constant over time for different geographic regions. Figure 5 shows the variation of transit function with respect to Tier 2 & Tier 3. The values per region seem to fluctuate only slightly and they correlate with the Transit model as stated earlier.

7. CONCLUSIONS

This paper proposes a simple Cournot market based model for analyzing the evolution of the Internet topology within and across geographic regions. The power of the model lies in the fact that the Nash equillibrium can be characterized by few aggregate parameters without the need for any individual parameters of each player in the system. By analyzing Internet topology data from 2002 to 2008, we corroborate the model and also perform detailed case studies of the rise and fall of individual ISPs across geographic regions.

8. REFERENCES

[1] `http://telephonyonline.com/access/finance/level3_att_assets_040407`.

[2] `http://newscenter.verizon.com/press-releases/verizon/2005/page.jsp?itemID=29709748`.

[3] `http://www.caida.org/tools/measurement/skitter/`.

[4] `http://asgeo.netlantis.org/`.

[5] `http://www.oecd.org/document/54/0,3343,en_2649_34225_38690102_1_1_1_37441,00.html`.

[6] J. P. Bailey. The economics of internet interconnection agreements. pages 155–168, 1997.

[7] A. Broido and k claffy. Internet topology: connectivity of ip graphs, 2001.

[8] X.-R. Cao, H.-X. Shen, R. Milito, and P. Wirth. Internet pricing with a game theoretical approach: concepts and examples. *IEEE/ACM Trans. Netw.*, 10(2):208–216, 2002.

[9] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM*, pages 251–262, 1999.

[10] R. Frieden. When internet peers become customers: The consequences of settlementbased interconnection. *Telecommunications Policy Research Conference*, 1999.

[11] R. Frieden. Without public peer: The potential regulatory and universal service consequences of internet balkanization. *Available at SSRN: http://ssrn.com/abstract=102927 or DOI: 10.2139/ssrn.102927*, June 1998.

[12] L. Gao. On inferring autonomous system relationships in the internet. In *Proc. IEEE Global Internet Symposium*, Nov. 2000.

[13] R. Govindan and H. Tangmunarunkit. Heuristics for internet map discovery. In *IEEE INFOCOM 2000*, pages 1371–1380, Tel Aviv, Israel, March 2000. IEEE.

[14] Y. Hyun, A. Broido, and k claffy. Traceroute and bgp as path incongruities.

[15] R. T. B. Ma, D. ming Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein. Internet economics: the use of shapley value for isp settlement. In *CoNEXT '07*, pages 1–12, New York, NY, USA, 2007. ACM.

[16] T. Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.

[17] S. Shakkottai and R. Srikant. Economics of network pricing with multiple isps. *IEEE/ACM Trans. Netw.*, 14(6):1233–1245, 2006.

[18] G. Siganos, M. Faloutsos, P. Faloutsos, and C. Faloutsos. Powerlaws and the as-level internet topology, 2003.

[19] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the internet hierarchy from multiple vantage points. In *INFOCOM 2002, New York, NY*.

[20] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. Network topology generators: Degree-based vs. structural. In *SIGCOMM*, 2002.