Link-Level Internet Structures

Scott Schremmer Kevin Wampler

ABSTRACT

The growth of the interconnections upon which the Internet runs is far from random, but rather follows primarily economic rules. Companies attempt to provide a high quality of service to a maximum number of customers while at the same time minimizing their costs. We examine some models of these economic tradeoffs, and the network topologies which can result. We look at how various parameters related to Internet growth and economic decisions alter the resulting network topology, particularly within a network undergoing incremental growth. We also consider how the structure of an incrementally growing network compares to that of a network which is planned in advance.

1. INTRODUCTION

Ever since its inception, the topology of the Internet has been shaped by various forces. In more recent times, these forces have been largely economic in their nature. Much of the current link structure in the Internet was laid by ISPs attempting to provide a service to their customer s and generate profits which outweigh the costs of setting up their network.

It is a natural hypothesis that these economic goals have shaped the interconnections in the Internet, but to the best of the author's knowledge no research has been done on this. Although there has been a reasonable amount of remarking over the heavy tailed router degree distributions (perhaps the easiest to measure feature of the network's topology), current models for this rely on assumptions without an explicit basis in reality. This leaves the reasons for the current topology of the Internet still somewhat of a mystery.

In response to these observations take some steps toward an economically motivated model of the topology of the Internet. In our models we consider the costs incurred by an ISP which attempts to provide service to some set of customers while minimizing the costs of actually providing this service. We focus on models of the link-level structure on the Internet – the graph of points of presence and backbone links. In addition to having qualitatively plausible structures, these models exhibit heavy tailed router degree distributions similar to those generated with existing models, but using economic incentives. Furthermore some of the aspects of these degree distributions (the very heavy tailed nature of the ROCKETFUEL router degree measurements) are also shown to arise from one of these models.

In addition we consider one of these models in the context of a system with more than one ISP. This allows for the consideration of factors such as market share, routing policy and the cost to purchase bandwidth. We show how these factor affect the networks formed and their relative goodness.

2. SIMULATING A SINGLE ISP

A single ISP is modeled as an edge set over a set of nodes representing points of network demand. Each node is associated with an amount of demand, which we will refer to as its *population*. This can be thought of as the total bandwidth demand of point in a network (such as the bandwidth usage of a city or other localized packet of population). We make the assumption that the traffic from a node is only in transfers to other nodes, and ignore any traffic within a node. It should be noted that this restriction is not important in the simulation, as a node which has some amount of traffic to itself (inter-city traffic) can be thought of as a node of lesser population which only transfers between other nodes. To aid in various aspects of the simulation, in particular the cost of laying a link between two nodes, we also give each node a position (x, y) in two dimensional space.

Each edge in this simulation represents to a link between two nodes. Just as each node is associated with a bandwidth demand, each link is associated with a maximum bandwidth capacity. If a link does not have sufficient bandwidth to accommodate the demand placed upon it we refer to the link as *overloaded*. We make the assumption that the quality of service provided by an ISP is only a function the bandwidth it provides. This the latency of a network is not considered to be of importance and any suboptimal service results from the overloading of links.

To model the economic considerations in designing an ISP backbone we consider not only the quality of service provided by a link structure but also the cost required to build such a backbone. To do this, for every pair of nodes n_1, n_2 we associate a positive real number representing the cost of laying a link directly connecting n_1 and n_2. This cost is a function, $c(n_1, n_2, b)$, of the the desired bandwidth capacity of the link and, since we assume homogeneous terrain, of the Euclidean distance between n_1 and n_2. In our simulations we simply estimate this cost function, but it can be reasoned that it should satisfy some basic properties.

Neglecting fixed costs associated with the bureaucratic transactions required to lay a link the cost function can be thought of as the sum of two terms. One of these terms represents the costs associated with transporting the cable, and with actually laying the cable out. Thus this term should depend only on the distance between n_1 and n_2 and should be linear in this distance. This assumes that the terrain in which the cable is laid is homogeneous, so that there is no advantage to laying the cable in certain areas. In practice this is not true, and there are regions and paths in which it is cheaper to lay cable. The effect is not only the result of the terrain in which the cable is laid (flat versus mountainous for example) but also of the ease and cost of obtaining the rights to actually lay the cable. Thus it is often the case that it is easiest to lay fiber along the route where another link has been previously laid, since the negotiations for laying the fiber along this route are likely to be easier. We do not fully simulate these effects in our simulations, though our incremental growth model does make some assumptions deriving from these observations.

The other term in the cost function represents the cost of the cable itself. It is clear that this term should be linear in the distance between n_1 and n_2 (neglecting discounts for buying in bulk, etc.) Furthermore, this term should also be sublinear in the bandwidth of the link. The rationale for this is essentially that twice the bandwidth can always be achieved for twice the cost by simply laying two cables instead of one, so that to be economically competitive a single cable of twice the bandwidth would have to be priced at less than double the cost of a single cable. The cost to connect n_1 and n_2 can then be written as:

$$w_1 dist(n_1, n_2) + w_2 dist(n_1, n_2) f(bandwidth) \tag{1}$$

Where w_1 and w_2 are weighting terms representing the cost per unit distance of physically laying out the cable and the cost per unit distance of the cable itself respectively. In our simulation we somewhat arbitrarily choose

$$f(bandwidth) = \sqrt{bandwidth} \tag{2}$$

though in general any sublinear function would do. The particular functions chosen is not actually particularly important, since as far as network topology is concerned what matters is that fatter links cost more and that buying a fetter link is cheaper than buying multiple links (f should be monotonically increasing and sublinear in the bandwidth). We further say that the cost required in deploying a given backbone structure is equal to the sum of the costs of its links.

It is also necessary to model how the populations which the ISP connects place demand on the links in the network. To do this we define another function giving the amount of traffic demanded between all pairs of nodes n_1 and n_2. To simplify the simulation we assume that the demand of each node is evenly distributed amongst the population of the rest of the Internet (note that this is different than being evenly distributed amongst the nodes in the Internet). Thus the demand between n_1 and n_2 in an Internet of total population P can be written as:

$$population(n_1) \frac{population(n_2)}{P - population(n_1)} \tag{3}$$

It is further necessary to describe the links in an ISP along which traffic between n_1 and n_2 is routed. To do this we make the simplifying approximation that all packets are routed along the shortest graph theoretic path between n_1 and n_2. We can thus describe the total bandwidth demand placed upon an ISP use to a single node n by placing a demand equal to the demand between n and n_2 on every link in the shortest path between n and n_2 for every other node n_2 in the Internet. The total bandwidth demand on a backbone is the sum of these demands as n ranges over every node in the graph. If the demand on a link exceeds its bandwidth we call the excess bandwidth that the link does not have the capacity to support to be the *overload* of the link.

3. INCREMENTAL GROWTH

In the real world, ISPs rarely plan their entire network in advance. Rather, an ISP grows over time by providing service to increasingly larger sets of customers. When an ISP adds new customers to it's service coverage, they make use of whatever existing backbone structure they have, and merely try to connect the new customers to this existing framework as efficiently as possible.

To model this sort of growth we consider an *incremental growth model*. In such a model an ISP starts out as a single node, and links to other nodes one at a time. Each node is added in a greedy manner, not accounting for possible future additions. Thus an ISP grows a spanning tree over the nodes in an Internet.

Such models have been investigated in [1]. In their simulation they use an incremental growth model where a new node, a in connected to a node in the current spanning tree, b, with a probability proportional to the degree of b. This is intended to account, in the case of a network topology, for the tendency of an ISP to favor points of presence which are already well connected. In this model they observe heavy tailed degree distributions, such as those observed in the real Internet [2].

There are some problems with this model. First, the probabilistic preference for nodes of higher degree is not directly related to any real world property. Even if it is the case where nodes of higher degree are preferred, it is highly unlikely that these nodes would be linked to with a probability proportional to their degree (without some other factor accounting for this). Since an ISP is trying to optimize some

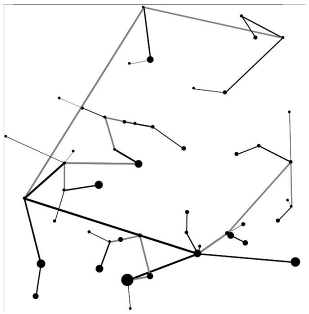

Figure 1: A network grown so that the incremental cost of each link is minimized.

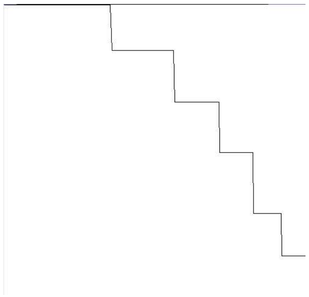

Figure 2: The degree distribution of networks which are grown to greedily minimize the incremental cost of each new link.

cost/service metric, they are highly unlikely to link to any node other than the optimal one. If a link to one node provides even a 5% advantage over another, the first node would be linked to nearly 100% of the time and not the 51% probability suggested by the other method.

Furthermore, this model has the disadvantage that it gives at best a vague explanation as to why nodes of high degree are preferred. Since router degrees in the actual Internet do exhibit a heavy tailed distribution, there is certainly some mechanism driving toward the use of high degree nodes, but the reasons for this have not been adequately answered.

To address these concerns we investigate in incremental growth model in which nodes are selected on an economic basis. In contrast to the Barabasi model, each time a new nodes is added to the network it is linked in the economically most efficient way, rather than only preferring more efficient links with some probability. We consider this growth in an Internet model consisting of randomly located nodes and in which nodes are added to an ISP in a random order.

When each node which is added to an ISP, we link that node to a node which is already in the ISP. For each such link we associate some positive real number which represents the cost of adding the link. Different cost functions yield different network topologies.

We first consider a network which is grown to attempt to minimize the cost of the links in its backbone structure. In the incremental growth model this is equivalent to minimizing the cost of each additional link, and thus the cost function for adding a given link is simply equal to the link cost function, $c(n1, n2, b)$. Here the bandwidth b is the minimum bandwidth required to provision the new link so that it has new overload. Since this bandwidth will be the same in all cases (more specifically, it will be proportional to the population of the new node) this reduces to connecting the new node to the closest existing node in the ISP.

An ISP grown in this manner is shown in figure 1, and a log-log plot of the degree distributions averaged of 10 such networks, each with a total of 50 nodes, is displayed in figure 2. As can readily be seen, such networks do not exhibit a heavy tailed degree distribution, though they do contain almost entirely short links. The reason for this is simple, and essentially reduces to the observation that a single node can only be the closest neighbor of at most 5 other nodes (6 if ties are allowed). This we should expect to see nodes of degree greater than 5 only rarely, when the order in which the nodes are added causes a globally suboptimal link to be chosen.

As is suggested by the resulting non-realistic network topologies, this model does not account for some critical aspect of network growth. In the previous model, each new link was given the minimum bandwidth needed to keep it from being overloaded. What was not accounted for, however, was the extra load placed on the other existing links of the network by the addition of the new bandwidth demand. As an ISP grows their backbone may become overloaded from the increased demand, resulting in a need to reprovision their existing backbone links in order to accommodate this demand.

We thus consider a model in which the cost of adding a new node is not just the cost of connecting it to the rest of the network, but also include the cost of all increases in the bandwidth of existing links which must be made to account for the increased demand. Whenever the bandwidth of a link is increased, we assume that the cost of this is equal to the cost of laying a new link of higher bandwidth between the same two nodes which the old link connected. In reality, an ISP might often ease the demand on existing links by laying entirely new fiber between points of presence which were not previously directly connected. We ignore this possibility to simplify the incremental growth model, and merely note that such an assumption is not entirely unreasonable due to the preference for laying links along the routes of old links that was mentioned earlier.

Initially we will consider a model in which upon the addition of new demand the ISP merely allocates the minimum bandwidth to each link required to eliminate overload along it. Thus each time a new population is added to the ISP every link in the network must be reprovisioned. It can be analytically suggested that such networks will be grown to prefer nodes of high degree.

For this analysis, we will consider the cost of reprovisioning and growing a network which grows from 1 to n nodes for networks of two different topologies: path and star. For the

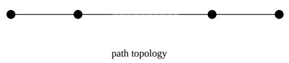

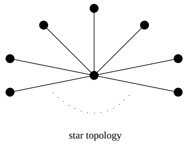

purpose of simplification we will assume that all links are of equal length, all nodes are of unit population, and that in the path topology the nodes are added in order.

First, let us consider a path of length $m-1$ which has been grown by the addition of a new node to length m. Since all of the links in the network were previously minimally provisioned, every link in the network must be reprovisioned. Since we know the network topology, we need only derive what bandwidth is required for each link to determine the cost (since, because all links are of equal length, the cost of a link is a function, $\alpha c(b)$ of only the bandwidth of the link times a normalizing constant).

Consider a single link in the path such that a nodes lie to the left of it and b nodes lie to the right of it. The bandwidth that such a link must accommodate is equal to the demand placed upon it be each node to its left communicating to each node on its right plus the demand from each node on its right communicating with each node on its left. Since there is only one route between each pair of nodes all such traffic must flow through the link.

Since the bandwidth demand of a node is evenly distributed amongst the population of the rest of the network, and since each node has equal population, the amount of traffic between any two nodes is $\frac{1}{m-1}$. The the total traffic across the link can be written as the sum of these demands for all a nodes communicating with b nodes plus the demand from all b nodes communicating with a nodes:

$$\begin{aligned}\sum_{i=1}^{a}\sum_{i=1}^{b}\frac{1}{m-1} \quad + \quad & \sum_{i=1}^{b}\sum_{i=1}^{a}\frac{1}{m-1} \\ = \quad & 2\sum_{i=1}^{a}\sum_{i=1}^{b}\frac{1}{m-1} \\ = \quad & 2\sum_{i=1}^{a}\frac{b}{m-1} \\ = \quad & \frac{2ab}{m-1} \\ = \quad & \frac{2a(m-a)}{m-1}\end{aligned}$$

making the total cost of reprovisioning the link equal to:

$$c\left(\frac{2a(m-a)}{m-1}\right)$$

The suboptimially of this configuration can be immediately seen by the fact that the links near the center of the path are burdened with quadratically more bandwidth than links near the edges. A graph of this load across the link is shown in 3.

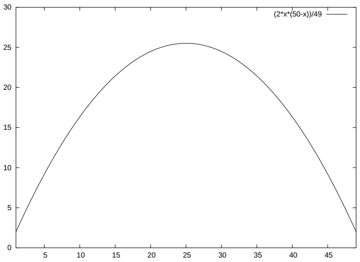

Figure 3: Load across the links of a path network of 50 nodes. The x-axis is the position of the link in the path, the y axis is the load placed upon a link at that point.

From the load on each link in the path, it is easy to express the total cost of a single reprovisioning as the sum of the costs for each of the $m-1$ links:

$$\sum_{a=1}^{m-1} c\left(\frac{2a(m-a)}{m-1}\right)$$

We now consider a network with a star topology. The high symmetry of this case makes the analysis somewhat simpler. First note that the node at the center of the star places a demand of $\frac{1}{m-1}$ on each link. Each other node places a demand of 1 on the link connecting it to the center node, and a demand of $\frac{1}{m-1}$ on each other link. Thus the total demand on each link is:

$$\frac{1}{m-1}+1+\sum_{i=1}^{m-1}\frac{1}{m-1}=2+\frac{1}{m-1}$$

making the total cost of the reprovisioning equal to

$$\sum_{a=1}^{m-1} c\left(2+\frac{1}{m-1}\right)$$

Notice that, as opposed to the path case where the demand across the center links grows linearly with the number of nodes, the demand on each link in a star network is bounded by a constant. Since c is monotonically increasing, this means that (particularly for large networks) the reprovisioning of links in a star network is cheaper that for a path network. Although though this most certainly does not constitute a full proof, it is suggestive that there will be a preference for higher degree nodes.

We have incorporated such an incremental growth model into our simulation. In it, nodes are added one at a time (as before), but instead choosing links to minimize the cost of the single link connecting th new node the the rest of the network, we minimize the cost of this new link plus the cost of all links which must be reprovisioned. An example of a resulting network on the same nodeset as in figure 1 is

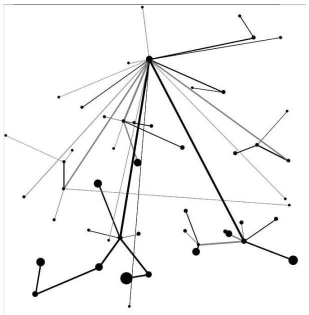

Figure 4: A network grown so that the sum of the incremental cost of each link plus any necessary reprovisionings of existing links is minimized.

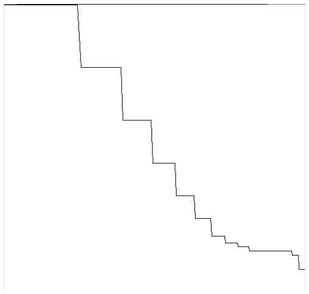

Figure 5: The degree distribution of networks grown to minimize the cost of new links and reprovisionings.

shown in 4, and the corresponding degree distribution graph in figure 5

As is illustrated in these figures, such networks are very heavy tailed – even more so than a power distribution would give, and thus suggest that minimizing the demand on existing links is a possible reason for the heavy tailed degree distribution in the actual Internet.

Although it accounts for some of the aspects of Internet topology it is clear that this model is highly unrealistic. In real life ISPs do not provision their networks minimally so that whenever there is new demand they must reprovision their network. Such a strategy would be obviously impossible to economically maintain.

Instead, ISPs overprovision their links not only to account for normal fluctuations in demand (flash crows, etc.) but also to allow for future growth, i.e. new customers being added into the existing network. We account for this in our model by specifying a fixed constant, ρ, which greater than or equal to one and specifies the factor by which new links are to be overprovisioned. Thus if $\rho = 2$ each new link is created with twice the capacity that is required.

Networks grown in this method exhibit a topology which is a blend of those which minimize the cost of each new link and those which always minimize accounting for every link being reprovisioned. What essentially happens is that the the overprovisioning of links allows for flexibility in placing links without always having to account for overloading the rest of the network. An example of the result of such a growth procedure is shown in figure 6, and a corresponding degree distribution is shown in 7. It is in particular worth noting that these networks display a degree distribution which resembles a power distribution (as well have a qualitatively plausible looking structure). Suggesting that some critical aspects of the topology of the Internet may be able to be explained with this model.

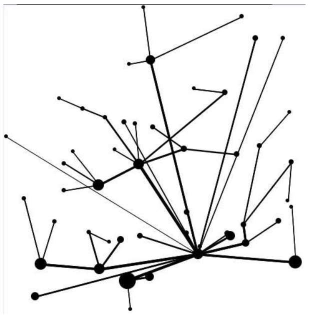

Figure 6: A network grown with overprovising of $\rho = 2$.

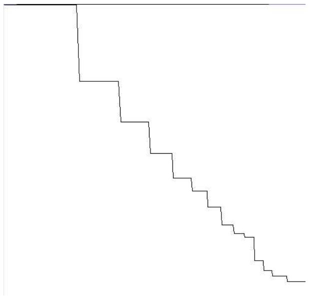

Figure 7: The degree distribution of networks grown as in figure 4, but with overprovisioning.

4. GLOBAL OPTIMIZATION

So far we have considered only networks which are grown a node at a time with greedy optimization and which have a tree topology. It is interesting for the sake of comparison to consider networks which are globally optimized to connect some set nodes as best as possible.

To define such an optimization problem it is necessary to create an objective function describing in essence the goals of an ISP. We formulate this as a simple weighted sum of the cost of the network and the sum of the overloads of the links in the network (G):

$$w_c cost(G) + w_l \sum_{l \in links(G)} overload(l) \tag{4}$$

Where w_c and w_l are weighting terms representing the emphasis an ISP places on minimizing the cost of a backbone versus providing maximal service (or alternatively how tolerant users of service defects compared to backbone cost).

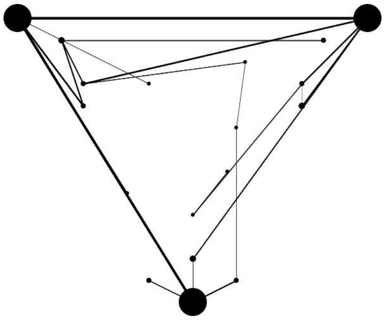

Figure 8: A balance of cost and service demands: $w_c = 1$ and $w_l = 1$

It should be noted that this is slightly different that the incremental growth objective function is that overloaded links are allowed. The reason for this is that without a concept of growth it is hard to express a tendency like overprovisioning. To partially compensate for this, we allow link overloads, but with the observation that they may also be interpreted as (up to a normalization constant) the tolerance for potential future overloads on a link. This model also has the advantage that it easily allows to view how different network topologies arise from different levels of emphasis on cost versus overload minimization.

For the single ISP backbone simulation we minimize the objective function for the ISP by performing a simulated annealing procedure. We start out with a random backbone structure which connects all of the nodes in the graph and then randomly apply four transformations:

1. add a link connecting two nodes
2. remove a random link
3. increase the bandwidth of a link
4. decrease the bandwidth of a link

The resulting backbone is accepted as a basis for the next iteration provided it is better than the previous backbone, or if it is worse with a random probability depending on how much worse and the current iteration.

We show the results of this simulation for various ratios of w_c and w_l corresponding to a high emphasis on the cost of the network, a high emphasis on overload minimization, and a balance of the two. The simulation takes place on a hypothetical network consisting of three large populations and several smaller populations. The results of this can be seen in figures 8, 9, and 10.

In these figures node size is proportional to population and link width is proportional to the log of the bandwidth of the link. Node that the more emphasis is placed on cost, the more sparse the network becomes and the more it exhibits a 'star' structure with larger cities being connected to each other and each having many links to nearby smaller cities. It is interesting to note the resemblance of these figures to the topologies of the AT&T and Sprint backbones and described in [3].

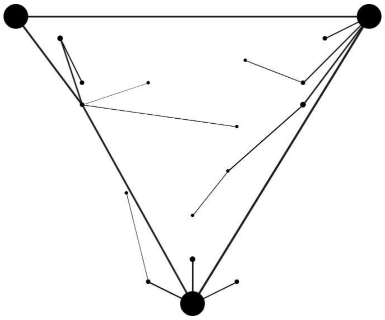

Figure 9: Emphasis on minimizing cost: $w_c = 5$ and $w_l = 1$

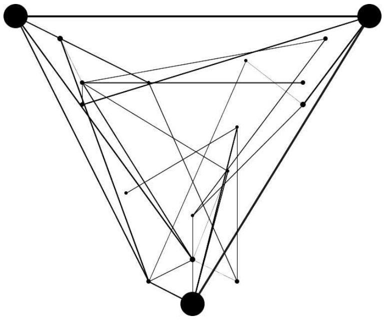

Figure 10: Emphasis on minimizing overload: $w_c = 1$ and $w_l = 5$

There are also some differences which can be noted between these topologies and those exhibited by the incremental growth model. Firstly, it can be seen that even when a large emphasis is placed on minimizing cost, it is still economical to connect the large cities with fat links, then branch service to the smaller cities off from there.

It is also notable that in the case of a large emphasis on minimizing overload, the networks exhibited a tendency to be more fully connected (rather than to simply have fatter links). This leads to a very heavy tailed distribution in which most of the nodes have high degree. This is, of course, reflective of the true Internet, but these results show a resemblance not displayed by the incremental growth models. The graphs of router degree distributions given by the ROCKETFUEL project [3] are not merely heavy tailed, but also display a downwardly concave distribution in a log-log plot. This downward concavity in a heavy tailed distribution is displayed on those networks optimized under an objective function which places a significant weight on overload optimization, see figure 11. This helps to verify our model, and indicates that an overconnection of the Internet used to minimize link overload my be responsible for many of its topological properties.

There is the possibility that this is a result of the annealing process rather than reflective of a global optimum. Even so it still suggests that similar structures could be expected

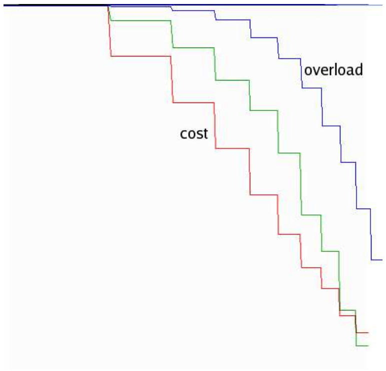

Figure 11: Node degree distribution for three ratios of cost/overload minimization weightings. red: $w_c = 5$, $w_l = 1$. green: $w_c = 1$, $w_l = 1$. blue: $w_c = 1$, $w_l = 5$.

to arise in networks constructed under a similar objectives using something resembling a local search strategy – a reasonable assumption for the growth of actual networks.

5. MULTIPLE ISPS

In this section we consider an extension of the model based on simulated annealing to consider more than 1 ISP. We consider 2 ISPs one of which serves all the nodes in our Internet while the other serves a portion of the nodes. This could be thought of as a small ISP attempting to gain a foothold which hasnt had the opportunity to expand to cover the whole network yet.

We define the market share of the smaller network to be the portion of the population in the cities in which the smaller network has a presence that subscribe to the smaller network. We assume that each network creates links to route its own packets within cities in which it has a presence. (A provider never buys bandwidth from another network in order to transfer data between 2 cities in which it has a presence). We assume that the smaller ISP purchases bandwidth from the large ISP to transfer data to cities in which it doesnt have a presence.

We use the same assumptions about population and its relationship to traffic as in the previous section. Also as in the previous section, within an ISP the data is routed along the paths with the least number of hops (and the network links are often chosen to support this). One important question when dealing with multi-ISPs is the question of how to route the data packets traveling between ISPs. We consider 2 models for this. In the first model, the smaller ISP gives its packets to the large ISP in the city in which the packets originate. In the second model, the small ISP transfers its packet to the city closest in terms of Euclidean distance to the city in which the packet is destined. (We assume that the small ISP does not know the internal details of the large ISPs structure but does know enough in order to determine geographically where to send the packet).

One difficulty in a system like this is determining how much an ISP should charge for its bandwidth. We assume (perhaps unrealistically) that the large ISP has the ability to track the usage of the small ISP and determine over exactly which links the data is being sent. The cost is based on this. Even if this is not possible, one could at least conclude that the large ISP has a general idea of the bandwidth used by the users of the small ISP and considers this when the contract is signed. Considering this information in our simulations allows us to have a reasonable charge on behalf of the large ISP as we vary the parameters of the simulation. (ie as the small ISP grows its bill will go down).

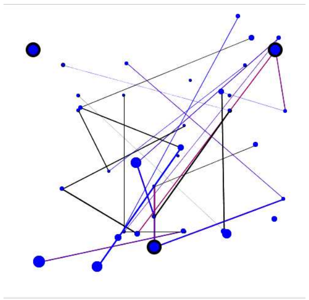

Figure 12: The larger of the two networks.

In order to calculate this cost we first consider the flow on the links over which the traffic will flow. (Using shortest paths and considering whether the ISP is being billed for the traffic from the point of origin or the node closest to destination). As in the previous section we assume that cost of the link is linearly proportional to its length and sub-linearly proportional to the bandwidth present. We calculate the cost to the large ISP for the extra traffic by subtracting the cost without the extra traffic from the cost with the extra bandwidth required to transport the data. We define the resale factor as the value by which this cost is multiplied by to determine the bill to the small ISP. Thus, if the resale factor is 1 it is being sent at cost.

We considered a range of values for this parameter. It seems in most cases the ISP would want to sell at a cost higher than what it costs for the link. However, it does seem there would be instances in which the larger ISP might be forced to sell for less. These include instances in which an ISP buys a large cable for future growth or is not able to realize growth that is expected.

In our simulations in this section we anneal as in the previous section. Each network is set up randomly and links are added, removed and have their bandwidth increase and decreased. If the network is improved by the change the change is kept. Also with some random probability which gets smaller over time changes which make the network worse are kept. In this section we anneal each network separately from the other. In most situations the owners of the network would try to maximize the profit for their own network rather than the overall goodness of the network as a whole.

Figures 12 and 13 show an example of a network created in this way. The first image shows the larger network. Note that all the cities are served by this network. The second image shows the smaller network. Only the cities connected on this picture are part of that network. For this plot we are

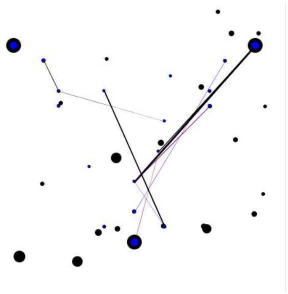

Figure 13: The smaller of the two networks.

using the balanced cost/service metric and assuming that the smaller network routes its packets to the edge of its network before they are passed to the other network.

We would like to determine under what circumstances a network would grow or recede. We assume that a network that has a lower cost (thus able to charge its customers less) and provides better service is more likely to acquire more customers. We normalize our goodness metric (consisting of the inverse of the cost added to the network overuse) by multiplying by the overall population that the ISP serves. Using this it is possible to guess which network is likely to grow and which is more likely to shrink and possibly fail. To characterize this we consider the network shown above consisting of 2 ISPs. We vary the market share of the smaller network as well as the resale factor of the larger ISP.

These results are shown in figures 14 and 15. The plots show the ratio of the small isp's goodness to the large isp's goodness as a function of marketshare and the resale factor. Examining the plots one quickly notices the many peaks and valleys. One might conclude that this is an artafact of the simulated annealing process. However, this could also indicate a certain amount of instability which makes an ISP's initial choices very important. A few significantly suboptimal choices early on could lead to a signficantly lower local maxima later. This does seem reflective of real life in which the difference between success and failure can be very slight.

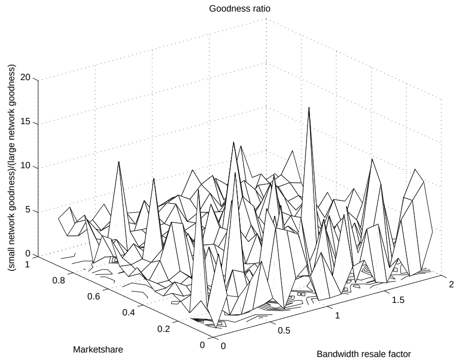

Figure 14: A plot showing the relative goodness of the 2 networks as a function of the smaller networks market share and the bandwidth resale factor. In this network the smaller ISP routed the traffic bound for the larger ISP to the point closest to its destination.

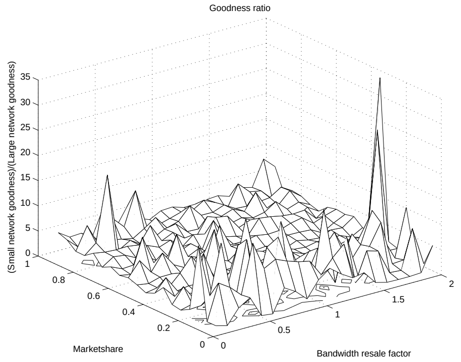

Figure 15: A plot showing the relative goodness of the 2 networks as a function of the smaller networks market share and the bandwidth resale factor. In this network the smaller ISP passed all traffic bound for the larger network immediately and was charged more as a result.

6. CONCLUSION AND FUTURE WORK

We have given two economically motivated models of Internet topology and shown that they exhibit similar properties as in the actual Internet. One of these models, the incremental growth model, dives a power degree distribution similar to that of other models, but using plausible econom ic assumptions for growth. This model suggests that a heavy tailed router degree distribution can may have arisen in the Internet due to the minimization of link cost and load on existing links when new links are added by an ISP. The other model, a global optimization model, showed a very heavy tailed distribution (downwardly concave on a log-log plot) similar to those measured by the ROCKETFUEL project. This model suggests that over connection of POPs to provide better service is a likely factor in this distribution.

There is a great deal of work which remains to be done in the direction of these models. Although the models rely on economic assumptions, these assumptions have not been checked against reality. It would be quite useful to set the parameters in the models to values close to their real-world counterparts, and check the structures similar to those in the actual Internet still arise. Furthermore, interviews with a few ISPs would help to verify (or bring into question) the assumptions for the factors impacting network structure that we make in out models. It would also be fruitful to explicitly include a customer reaction to service in our models, as this is almost certainly a large factor in the choices made by ISPs.

As the internet is more complex than one ISP providing service to its customers, we considered networks of more than one ISP. This allowed us to examine the results of the interaction between the 2 networks as they attempted to grow. There is much work to be done here to consider the effects of geography, peering and routing agreements, initial choices related to topology and other such parameters and their effect on the success of an ISP.

Although our results do shed light on plausible reasons for the structure of the Internet, these answers cannot be said to be conclusive. This is partially due the large and diverse number of factor effecting the growth of the Internet and to the difficulty accurately comparing the similarity of the topologies generated by a model to those in the Internet. Still we hope that this first step given some useful information both about the reasons behind the Internet's structure and of avenues where more work can be done in this area.

7. REFERENCES

[1] A.-L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[2] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origin of power laws in internet topologies revisited.

[3] N. Spring, R. Mahajan, and D. Wetherall. Measuring isp topologies with rocketfuel, 2002.