An Agent-based Model for the Evolution of the Internet Ecosystem

Amogh Dhamdhere and Constantine Dovrolis
{amogh, dovrolis}@cc.gatech.edu
Georgia Institute of Technology

***Abstract*—We propose an agent-based model for the evolution of the Internet ecosystem. We model networks in the Internet as selfish agents, each of which tries to maximize a certain utility function in a distributed manner. We consider a utility function that represents the monetary profit of a network. Our model accounts for various important constraints such as geography, multihoming, and various strategies for provider and peer selection by different types of networks. We implement this model in a simulator, which is then used to solve the model and determine a "steady-state" of the network. We then present a set of "what-if" questions that can be answered using the proposed model and by studying the properties of the resulting steady-state.**

I. Introduction

The Internet, at the interdomain level resembles in several ways a natural ecosystem. Autonomous networks engage in competitive transit (or customer-provider) relations, and also in symbiotic peering relations. These relations, which are represented as interdomain logical links, transfer not only traffic but also economic value between networks. The Internet ecosystem is dynamic, as networks continually attempt to maximize, in a distributed manner, their utility obtained from connecting to the Internet. The dynamics of the Internet ecosystem are determined both by external "environmental" factors (such as the state of the global economy or the popularity of new Internet applications) and by complex incentives and objectives of each network. Specifically, networks attempt to optimize their utility or financial gains by dynamically changing, directly or indirectly, the networks they interact with. In recent times, there has been a great deal of interest in understanding the evolution of the Internet ecosystem. Recent trends such as the rise of major content providers, the increasing amount of peer-to-peer traffic, and the changing peering landscape raise important questions about where the Internet is heading, in terms of topological and economic organization. This paper presents an approach that can be used to answer some of these questions about the evolution of the Internet.

We propose an agent-based, dynamic model for the evolution of the Internet ecosystem. The model considers the Internet at the interdomain level as a network of interacting, selfish agents, each of which is concerned with optimizing its utility function in a distributed manner. An agent has only local knowledge, and can change only its interactions with its neighbors by creating or removing links to them. We propose a first-principles model for these local actions of networks, taking into account the various objectives and constraints faced by networks. Our model relies on the concept of profit as the utility function that networks attempt to maximize, and takes into account practical considerations such as geography, multihoming, transit, peering and operational costs, and various provider and peer selection strategies often followed by networks. As this model is analytically intractable to solve, we rely on simulations and simply allow the model to "run". This results in an equilibrium or steady-state where no network has the incentive to make a further change in its connectivity.

Our model is not intended to be a topology generator, i.e., our goal is not to produce a topology that matches structural properties of the Internet topology such as degree distributions. Instead, the focus of our study is to be able to answer "what-if" questions about possible evolution paths for the Internet ecosystem. For example, one could ask questions of the form: What is the best strategy for provider and peer selection for transit providers when the interdomain traffic matrix consists of mostly client-server traffic? What are the properties of the steady-state when each network usess its optimal strategy? What if the interdomain traffic matrix changes such that most of the traffic is peer-to-peer traffic? How does multihoming affect the economics of the Internet ecosystem?

The scope of possible "what-if" questions we can answer with the proposed model is quite broad. In this paper, we focus on describing the model, and present a *canonical model* which we believe represents the current state of the Internet. We then present a number of specific questions that we plan to answer in future work. First, we plan to study various topological and economic properties of the steady-state networks that result from the canonical model, when networks use different strategies for provider and peer selection. Then, we plan to find the optimal provider and peer selection strategies for each network. Finally, we plan to study the effect of external factors such as the interdomain traffic matrix and the underlying cost structures for transit, peering and local costs.

II. The Internet ecosystem

We model the Internet ecosystem at the granularity of autonomous networks. Each network is independently operated and managed and is inherently *selfish*, as it tries to optimize its utility obtained from connecting to the Internet. Networks can have different utility functions depending on their requirements and business interests in connecting to

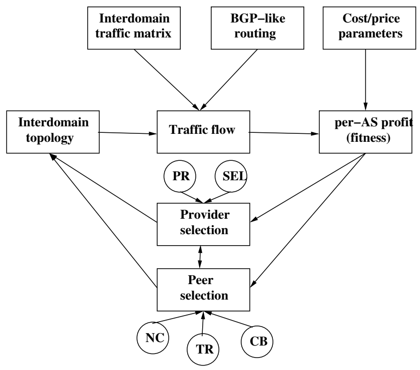

Fig. 1. The feedback loop between topology, traffic flow and fitness in the AS ecosystem

the Internet. Further, each network in the Internet operates under the condition of *limited knowledge*, meaning that it has limited ability to predict the actions of other networks, or to evaluate the long-term effects of its own actions. The focus of this paper is mainly on the networks at the core of the Internet, that are in the business of providing Internet transit. These provider networks are mainly concerned with maximizing their financial utility, or *fitness*. Providers try to achieve this goal by changing the set of networks they interact with, by intelligently choosing their own providers and peers, and attracting customers.

As a result of this optimization performed in a distributed manner by each network, the Internet topology continuously evolves. In an attempt maximize their fitness, networks create and drop links with other networks. The change in the topology affects the flow of traffic in the system. The interdomain traffic flow is a crucial factor affecting the evolution of the Internet. The fitness of a network is determined by the traffic it receives from its customers (for which it gets paid), the traffic it sends to its providers (which it pays for), and the traffic it sends to its peers (which results in transit costs saved). Consequently, a change in the topology and traffic flow can cause a change in the fitness function of a network. This change in the fitness of a network could create an incentive for that network to change its set of providers or peers, which again changes the topology. We thus have a *feedback loop* in which topology changes cause changes in fitness, which lead to more topology changes.

This effect is illustrated in Figure 1. The topology in combination with the interdomain traffic matrix and routing policies determine the traffic flow in the network. The traffic flow, together with external conditions such as pricing schemes and operational costs determine the fitness of each network. Networks attempt to optimize their fitness by changing their connectivity, which results in topology changes. Networks could differ in terms of the factors that affect their fitness, and the types of actions those networks take in order to optimize their fitness. In the next section, we describe in detail each component of this model: The different types of networks, the traffic model, the routing model, the economic model, the role of geography, and different possible provider and peer selection strategies.

III. A MODEL FOR AS INTERACTIONS

In this section we describe the key components of the model that we develop for the Internet ecosystem. interactions among networks in the Internet ecosystem.

A. Network types

We consider the following types of networks:

Enterprise Customers (EC): These are networks at the "edge" of the Internet. Such networks are normally sources of traffic, such as websites and hosting companies, or sinks, such as campus/corporate/residential access networks. A fraction s of ECs are mostly content sinks, while the remaining are mostly content sources. We do not model the economic fitness of ECs. The actions of ECs are limited to choosing the desired set of upstream providers. In our model, ECs do not engage in peering relationships.

Content providers(CP): CPs are networks that are mostly sources of traffic. These are, however, distinct from content ECs in that they also engage in peering relationships. Recent studies [1], [2] show that content providers are increasingly active in forming peering links. Google, Akamai and MSN are examples of CPs.

Small Transit Providers (STP): Transit providers are networks whose main business function is to provide transit to other networks. STPs are providers that have limited geographical presence, often just a small set of geographical regions. In our model, STPs do not themselves source or sink any traffic. STPs are motivated by economic benefit, and choose their set of providers and peers in such a way as to maximize their profit (fitness). France Telecom, Rogers Telecom, and Chinacom are examples of STPs.

Large Transit Providers (LTP): LTPs are similar to STPs in the sense that their business function is to provide Internet transit to other networks. LTPs, however, have a larger geographical presence, often spanning the entire world. LTPs, too are motivated by economic benefit, and choose their own providers and peers in such a way as to maximize their profit. LTPs are like "tier-1" networks, but in our model, an LTP may have providers if necessary for reachability. Sprint, AT&T and Level 3 Communications are examples of LTPs.

We refer to each network in our model as Autonomous Systems (ASes), in the sense that each of them independently chooses its set of providers/peers, and has complete control of its internal network. This does not refer to ASes in the BGP sense, because it also captures the presence of networks that do not have AS numbers.

B. Traffic model

The traffic model concerns the generation of the inter-AS traffic matrix that determines the amount of traffic sent from each AS to every other AS. In order to generate this traffic matrix, we assume that each AS i consumes an amount of traffic I_i. We consider two different types of traffic: Client-Server (CS) traffic and Peer-to-Peer (P2P) traffic. CS traffic flows from content sources (source ECs and content providers) to sink ECs, while P2P traffic flows between sink ECs. A *popularity* $p_{c,i}$ is associated with each content source i (source EC or CP), which determines the popularity of the traffic generated by that source. Similarly, a popularity $p_{p,i}$ is associated with each sink EC i, which determines the popularity of the p2p traffic generated by that sink EC. The incoming traffic for each AS can now be split among different sources according to the particular traffic model being used. By changing the relative fraction of CS and p2p traffic, we can model various inter-AS traffic scenarios. A parameter c determines the fraction of the total incoming traffic at a network that is CS traffic, meaning that the sources of this traffic are source ECs and CPs. Let S_c be the set of sources of this traffic and n_c be the number of these sources. The sources are ranked in decreasing order of the client-server popularity metric $p_{c,j}$. For each access stub i, we then assign the incoming CS traffic from each content source j using a Zipf distribution as follows:

$$T_{ji}^{CS} = c \frac{\frac{1}{j^{\alpha}}}{\sum\limits_{k=1 \ldots n_c} \frac{1}{k^{\alpha}}} I_i$$

Next, we assign the incoming P2P traffic for each network i as follows. We construct the set of sources of p2p traffic for i. This set S_p of size n_p is simply the set of all sink ECs. These access stubs are ranked in decreasing order of the p2p popularity metric $p_{p,i}$. For each access stub i, we then assign the incoming P2P traffic from other access stubs using a Zipf distribution as follows:

$$T_{ji}^{P2P} = (1-c) \frac{\frac{1}{j^{\alpha}}}{\sum\limits_{k=1 \ldots n_p} \frac{1}{k^{\alpha}}} I_i$$

The parameter α for the Zipf distribution controls the distribution of traffic from different sources to a particular destination. The skewness of this distribution can be controlled by varying this parameter. An important feature of the traffic matrix is that the CS popularity and P2P traffic popularity of a source is the same for all sinks. This reflects the fact that certain popular websites are the largest sources of traffic for most destination networks in the Internet. We assume that the same popularity characteristics also apply to sources of P2P traffic.

To control the overall nature of the traffic matrix, we vary the parameter c. c=1 means that the traffic matrix is completely CS in nature. c=0 means that the traffic matrix is completely P2P in nature.

C. Routing model

We capture the key properties of interdomain routing. In general, Internet routing follows the policy of "no-valley, prefer customers". The no-valley policy implies that traffic that enters a network i from one provider cannot exit through a provider of i. Also, traffic that enters network i from a peer cannot exit through another peer. For a network i, the rules for selecting the next hop network towards each destination j can be summarized as follows: i first prefers the customer that advertises the shortest path to reach j. If j cannot be reached through a customer link, then i chooses the peer that advertises the shortest path to j. If there are no customers or peers that can reach j, then i chooses the provider that advertises the shortest path to j. In each of the previous cases, if multiple neighbors of i advertise equal-cost paths to j, then i breaks ties deterministically using the neighbor's node id.

Calculating policy-compliant shortest paths between all pairs of nodes can be computationally very expensive (O(N^3), where N is the number of nodes in the graph). We optimize the computation of routing tables by using an algorithm inspired by the method proposed by Gao and Wang [3]. We simplify the routing computation by assuming that stub nodes do not form peering links. We can then calculate the shortest policy-compliant paths among providers. This can be done efficiently in $O(N_p E_p)$, where N_p is the number of providers and E_p is the number of edges among providers. Following this step, each provider p learns the best path towards each stub s via the provider p' of s for which p has the shortest path. This can be done in $O(N_p N_s d_{p,s})$, where $d_{p,s}$ is the multihoming degree of stubs, and N_s is the number of stubs. Finally, each stub s determines the best path towards stub s'. To do this, s chooses the provider p among its set of providers that gives the shortest path towards s'. The final step can be done in $O(N_s^2 d_s)$.

D. Economic model

Given the interdomain traffic matrix, the interdomain topology and the routing model, we can calculate the traffic flow in the network. The traffic flow determines the aggregate amount of traffic that flows over each link in the network. The economic fitness of each transit provider in the network is a function of this traffic flow.

For a transit provider, income is in the form of payments by its transit customers. The provider also pays transit costs to its own providers, peering costs, and also pays to maintain and operate its own network (money spent on leasing/purchasing infrastructure, staff salaries, etc.). We formulate an expression for the fitness of a network i as follows. Let $\mathcal{C}_i$ be the set of customers, $\mathcal{P}_i$ be the set of providers and $\mathcal{R}_i$ be the set of peers of a transit provider i. Its "fitness" is:

$$f_i = \sum_{c \in \mathcal{C}_i} T_i(v_{ic}) - \sum_{p \in \mathcal{P}_i} T_p(v_{ip}) - \sum_{r \in \mathcal{R}_i} R_i(v_{ir}) - L_i(v_i)$$

The function $T_i(v_{ic})$ determines the transit payment made by customer c to i when the aggregate traffic exchanged by the two networks is v_{ic}. The function $T_p(v_{ip})$ determines the transit payment made by i to provider p when the aggregate traffic exchanged by the two is v_{ip}. In practice, transit prices show *economies of scale*, meaning that the per-bit cost of purchasing Internet transit decreases as the volume of traffic increases. We model this phenomenon by using a concave increasing function to represent the cost of Internet transit as a function of the traffic volume. We use the following function to determine the transit payment made by network i to provider p.

$$T_p(v_{ip}) = m_{t,p} * v_{ip}^{e_t} \tag{1}$$

The function $R_i(v_{ir})$ determines the monetary cost of maintaining the peering link between i and r when the aggregate traffic exchanged over the peering link is v_{ir}. These costs are not paid by the peers to each other; rather, they are incurred to purchase/lease a sufficiently high-bandwidth link to the peering location. These costs are traffic dependent, and also show economies of scale. We use the following function to determine peering costs:

$$R_i(v_{ir}) = m_{r,i} * v_{ir}^{e_r} \tag{2}$$

The function $L_i(v_i)$ determines the monetary cost of maintaining the local network for network i when the aggregate traffic handled by i is v_i. We use a local cost function of the form:

$$L_i(v_i) = l_i + m_{l,i} * v_i^{e_l} \tag{3}$$

The term l_i represents the traffic-independent costs, while the other term represents the traffic-dependent costs incurred by i for operating its local network.

The transit, peering and local cost exponents (e_t, e_r and e_l) control the extent of the economies of scale associated with the various costs. A lower value of the exponent results in larger economies of scale, *i.e.*, the per-bit cost of transit decreases faster as the total volume of traffic increases. All providers are assigned the same exponents for their transit, peering and local cost functions, but differ in the multipliers $m_{t,i}, m_{p,i}$ and $m_{l,i}$. This is consistent with data collected by Norton [4] and Chang [5].

E. Geographical presence

Our model captures geographical constraints by assigning a geographical presence to each network. It is important to account for geographical presence in a model model of interdomain interactions, because it constrains both provider and peer selection for networks. A network can only choose a provider with which it has at least one region in common. We model geography by dividing the world into a number of regions. Each network is then assigned a geographical presence in a subset of those regions. Similarly, a network can only peer with a network with which it has at least one region in common. Networks of different types have significantly different geographical presence. In particular, STPs are present in a small number of regions, while LTPs typically have a presence in most of the world.

F. Provider selection methods

Here, we describe the process and methods used by a network i to choose its set of upstream providers. Network i first determines its set of potential providers. Potential providers of i have at least one region in common with i and are not in the customer tree of i or of any peers of i. Let the set of potential providers of i be $\mathcal{F}_i$. Network i must now choose its final set of providers from the set $\mathcal{F}_i$. Networks of different types can have different objectives in choosing providers. We consider the provider selection goals that are reasonable for different types of networks. For example, it is likely that many edge networks are concerned simply with minimizing the price they pay for Internet transit. Transit providers, on the other hand, may avoid choosing their competitors as providers. We consider three methods of provider selection that capture such expected behavior by edge networks and transit providers.

Price-based (PR): In this method, the goal of network i is to simply choose the cheapest providers. The metric used for comparing providers is the transit price multiplier $m_{t,j}$ associated with provider j.

Price-based Selective (SEL): This method is applicable only to transit providers. For a provider that is concerned with maximizing its profit, it makes sense to select the cheapest providers. A provider i, however, would not want to choose as provider a network that it could peer with in the future. Provider i would also avoid a provider that it views as a competitor, *i.e.*, a provider that competes for the same set of customers as i. We consider a provider selection strategy where an STP does not choose other STPs in the same region as providers. By using this strategy, a provider i simply removes all STPs from the set $\mathcal{F}_i$. As in PR provider selection, the metric used for comparing providers is the transit price multiplier $m_{t,j}$ associated with provider j.

Performance-based (PF): We consider a method of provider selection whereby edge networks (ECs and CPs) choose providers in a "performance-aware" manner. The recent popularity of performance-sensitive streaming video and peer-to-peer applications means that the sources and sinks of this content have an incentive to connect to the providers that can offer the best performance (*e.g.,* in terms of short paths) to the sources/destinations of this content. In this method, network i needs to have an estimate of the traffic sent to and received from each of its destinations. i estimates a performance metric for each of its potential providers. The metric is the path length to the major sources and destinations of traffic, weighted by the traffic volume to those destinations. For each destination j of i, let A_{ij} be the total traffic sent and received by i to/from j. Let l_{kj} be the path length from provider k to destination j. Now, the performance metric associated with provider k is given by $L_k = \sum_j A_{ij} l_{kj}$.

G. Multihoming

Multihoming refers to the connection of a network to multiple upstream providers. Networks that focus on reliability and availability have used multihoming for several years. Multihoming is increasingly popular recently, particularly by transit providers in the core of the network [2]. In our model, each network i is assigned a *maximum multihoming degree* $d_{p,i}$. A network may impose a maximum multihoming constraint due to practical or economic considerations. In practice, it may not be possible for network i to always connect to $d_{p,i}$ providers. This could be the case if i cannot find $d_{p,i}$ providers in its set of *potential providers*.

For each of the provider selection methods mentioned in Section III-F, lower values of the performance metric are better. Network i ranks its set of potential providers in increasing order of their performance metric. It then chooses the top $d_{p,i}$ providers according to this ranking. Each network i can have a different multihoming degree $d_{p,i}$. Multihoming degrees for networks are drawn from a range, depending on the type of network.

H. Peer selection methods

Here, we describe the process and methods adopted by a network i to choose its set of peers. For a network, the objective for peering is to save transit costs, and in some cases to maintain reachability to the rest of the network. We consider three different methods for peer selection by networks, based on what commonly happens in the Internet.

Peering by necessity (NC): In this mode, a network i only peers with j if that is necessary to maintain reachability. In some cases, it is not possible for i and j to reach consensus on which of the two should be the provider of the other. This happens in the case where both i and j would choose the other as their provider according to their own provider selection criterion. In this case, i and j form a peering link "due to necessity".

Peering by traffic ratios (TR): A common approach used by ISPs to make peering decisions relies on traffic exchange ratios. In this method, two ISPs i and j peer if they exchange "roughly equal" volumes of traffic. In practice, this is implemented by measuring the ratio of the traffic that flows from i to j and from j to i. If this ratio is less than r_t, then i and j agree to peer.

Peering by cost-benefit analysis (CB): In this method for peering, a network i assesses both the costs associated with a specific peering link, and the potential benefits that can be achieved by that peering. The costs associated with peering are due to the fixed and traffic dependent costs of establishing a peering link. The potential benefits are due to saved transit fees. Network i chooses to peer with j if the estimated benefits are greater than the estimated costs. This implies that the fitness of i after establishing the peering link would be greater than the fitness before. In practice, a network i needs to estimate the "peerable traffic volume" with network j, which can be used to calculate the costs of peering. Network i then subtracts the peerable traffic volume from the total traffic sent to the transit provider of i, and calculates the estimated transit payments that can be saved through peering.

I. Network actions

In our model, networks periodically reconsider their provider and peer selection to connect to the optimal set of providers and peers, *i.e.*, to maximize their own fitness. A detailed description of this process follows.

1) **Provider selection:** First, a network i identifies the set of preferred providers, according to its provider selection criteria. Let this set be P_i.
2) **Try to peer with providers:** If network i does not form peering links, skip to step 3. Else, i tries to convert each of its provider links to peering links. For this purpose, we evaluate the provider selection criteria of j, and find the set P_j. If $j \in P_i$ and $i \in P_j$, then i and j become peers "due to necessity". This condition captures the situation where i and j cannot agree on who should be the provider of whom. In this case, they need to peer to maintain global reachability for their customers. Network i then removes transit links to providers that are also in the customer tree of j. The intuition for this is as follows: When i and j form a peering link, some providers from P_i may be in the customer tree of j. i will never use such providers to reach nodes in the customer tree of j, since the direct path through the peering link is preferred. Figure 2 represents a case where i can safely remove providers k and l after forming a peering link with j. [1]

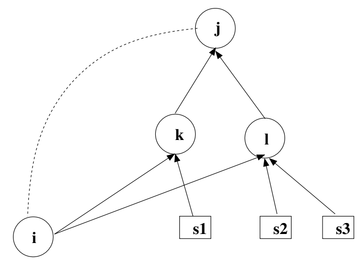

Fig. 2. Network i can remove providers k and l after forming a peering link with provider j.

3) **Check for potential peering candidates:** Analogous to the set of potential providers, network i maintains a list of possible peering candidates, R_i. As ECs do not peer in our model, the set of peering candidates of i is restricted to LTPs, STPs, and CPs that have a geographical region in common with i. For each possible peering candidate k, i performs the following actions: If k is already a peer of i, then i checks whether the peering link with k should be maintained. As this is i's turn to act, i *unilaterally* verifies whether the peering requirements are satisfied. Network i also verifies if it needs to peer with k due to necessity. If these peering criteria are not satisfied, then i *de-peers* k and exits the peering loop. If i and k are not peers, then i examines whether it is possible to establish a new peering link with k. This is a bilateral decision, and hence the peering criteria of both i and k must be satisfied in order for a peering link to be established. If the peering link is formed, then i again executes the procedure for removing providers that are in the customer tree of k (see step 2. If the peering link is formed, i exits the peering

[1] A corner case can occur when i needs providers to reach networks that are not in the customer tree of j, but all of i's providers are also in the customer tree of j. Rather than selecting arbitrarily which provider to keep, we impose the condition that i *keeps both* k *and* l.

loop. Note that in one move, i may add or remove only one peering link.

Note that all the actions performed by a network in each move are *completely deterministic*. This is in sharp contrast to previous evolutionary models of Internet topology (such as those based on preferential attachment [6]). Those models effectively generate a random graph that has certain structural properties such as a desired degree distribution. Our model is not intended to be a topology generator that matches certain structural properties of the Internet topology. Instead, our model attempts to *explain* the topology dynamics by modeling the optimizations made by networks. These optimizations are largely deterministic in nature, as each network attempts to unilaterally maximize its fitness.

J. Canonical model (CN)

So far, we have described the key components of our model in general terms, without mentioning any parameter values. In this section, we describe a model, referred to as the *canonical model* (CN), which we view as representative of the real Internet. This is essentially a parameterization of the model described previously, and it includes values for the fraction of P2P traffic, transit, peering and local cost parameters, multihoming degrees, geographical coverage and provider and peer selection strategies employed by different classes of networks.

Unless otherwise stated, in this paper, we work with a network consisting of 180 ECs, 20 providers (4 LTPs and 16 STPs) and 10 CPs. (the scale issue is further discussed in Section III-K). We assign 80% of ECs to be content sinks, and 20% to be content sources ($s = 0.8$). Each content sink i consumes an amount of traffic I_i, where I_i is drawn from a Pareto distribution with mean 2000Mbps and shape parameter 1.1. This produces the effect where certain sinks are "heavy-hitters" in terms of the amount of traffic consumed. In the canonical model, the traffic matrix is neither completely client-server, or completely P2P. Instead, the we set $c = 0.8$ to represent a "predominantly CS" traffic matrix, where 80% of the traffic consumed by a sink is CS traffic. The parameter α of the Zipf distributions used for creating the interdomain traffic matrix is chosen as 0.8. This produces a distribution in which the traffic from different sources to a particular destination is significantly skewed.

As far as possible, we parameterize the economic model using real-world data. Chang et al. [5] report that the exponent for the transit pricing functions of providers e_t is around 0.75, while the peering cost exponent e_r is around 0.25. The transit price multipliers $m_{t,i}$ of STPs are between 30 and 140, while those of LTPs are between 80 and 150. These values are based on data reported by Norton [4] in 2006. The peering cost multipliers $m_{r,i}$ are between 300 and 400. In the absence of data about local costs, the local cost exponent e_l is set to 0.5. The local cost multipliers are set differently for STPs and LTPs. These are between 100 and 200 for STPs and between 300 and 400 for LTPs. The traffic independent costs for LTPs are greater than those for STPs. This reflects the fact that LTPs have larger networks (due to a larger geographical scope), and hence need to spend more to maintain their network. The local cost parameters are assigned so that the traffic-dependent and traffic-independent local costs account for roughly equal fractions of the total local costs incurred by a provider. The transit, peering and local cost parameters are assigned in such a way that for the same traffic volume, peering costs are the lowest, followed by traffic-dependent local costs; transit costs are the highest.

We construct the initial topology in such a way as to match certain known properties of the Internet's interdomain topology. A recent study [2] measured the provider preference of different classes of networks in the Internet, and found that 60% of the providers of ECs are STPs and 40% are LTPs. On the other hand, approximately half of the providers of STPs and CPs are STPs. In our initialization, LTPs are assumed to be present in each geographical region, and are fully meshed by peering links. This is similar to the well-known clique of Tier-1 providers in the Internet. STPs are deterministically connected to other STPs and LTPs, in such a way that the number of links between STPs and LTPs is approximately equal to that between STPs and STPs. To connect ECs and CPs, we follow a procedure that simulates preferential attachment. We add ECs and CPs sequentially, choosing a provider (either STP or LTP according to the provider preference of ECs and CPs) with a probability that is proportional to the existing customer degree of that provider.

In the canonical model, ECs use PR provider selection and do not peer. CPs also use PR provider selection, but peer selectively based on CB analysis. For STPs and LTPs, we consider two possibilities of provider selection: PR and SEL. We have three possible peer selection methods for STPs and LTPs: NC, TR and CB. The combination of a provider and a peer selection method defines a *strategy* for STPs and LTps, which gives a total of 6 possible strategies for each transit provider. We define a *scenario* as a specification of the provider and peer selection strategies that STPs and LTPs follow. In a scenario, we assume that all providers belonging to the same class follow the same strategy. For example, the notation

$$\{CN, (SEL, TR), (SEL, NC)\}$$

represents the scenario with the canonical model (CN), where STPs use SEL provider selection and peer using TR, while LTPs select providers using the SEL rule and peer based on NC.

Tables I and II summarize the parameterization of our canonical model.

K. Computing the steady-state network

Our objective is to determine what happens to the network, in terms of topology, traffic flow, economics and performance, when STPs and LTPs use different strategies for provider and peer selection. For this purpose, it is necessary to "solve" the model, computing the steady state-network given the initialization and the strategy of each network. The steady-state is a situation where no network has the incentive to

metric	EC	STP	LTP	CP
Number of networks	180	16	4	10
Number of regions	1	2	5	1
Max. multihoming degree ($d_{p,i}$)	1	2	3	3
CS popularity ($p_{c,i}$)	[0,100]	-	-	[0,100]
P2P popularity ($p_{p,i}$)	[0,100]	-	-	[0,100]
Transit multiplier ($m_{t,i}$)	-	[40,130]	[80,150]	-
Peering multiplier ($m_{r,i}$)	-	[300,400]	[300,400]	-
Local cost multiplier ($m_{l,i}$)	-	[100,200]	[300,400]	-
Traffic-independent local cost (l_i)	-	5000	50000	-
Incoming traffic (I_i)	[0,2000]	-	-	-

TABLE II
CANONICAL MODEL PARAMETERIZATION: PER-NODE METRICS

global metric	notation	value
Fraction of client-server traffic	c	0.8
Fraction of content stubs	s	0.2
Source popularity parameter	α	0.8
Total number of regions		5
Transit price exponent	e_t	0.75
Peering cost exponent	e_r	0.25
Local cost exponent	e_l	0.5
Peering traffic ratio	r_t	2

TABLE I
CANONICAL MODEL PARAMETERIZATION: GLOBAL METRICS

unilaterally change its set of providers or peers. We solve this model computationally, as it is too complex to solve analytically. Solving the model involves iteratively allowing a network to play (according to its pre-defined strategy in each move), until we reach a stage where no network has the incentive to unilaterally make changes to its connectivity. This state, analogous to the concept of Nash Equilibrium in game theoretic models, represents the "steady-state" of the network. We assume that nodes play in a particular sequence, with a randomly chosen starting node.

1) Pick the next network i in the playing sequence.
2) Complete the move of network i, as described in section III-I
3) If the move of network i causes the topology to change, recompute the routing tables, traffic flow and fitness function of each network.
4) Check termination criteria. If each network has had a chance to play and has not made any change to its connectivity, then stop.

An important question is whether this model always reaches steady-state. Several aspects of our model are explicitly designed to avoid oscillations, e.g., the part of the model that focuses on removing unnecessary providers, and the fact that a provider adds or removes just a single peering link in each move. In practice, we find that in most cases, the network reaches a steady-state. In the cases where steady-state is not reached, the oscillations are due to a small number of networks, and even if the topology does not remain constant, the traffic flow and fitness distribution is practically constant. We plan to study the convergence properties of the model in future work.

Another important issue is the uniqueness of the steady state. We find that for a given initial topology and set of strategies for networks, *the order in which networks play can affect the steady-state network*. In some cases networks happen to make the "right move at the right time", such as forming a particular peering link or choosing a certain provider, causing different steady-state networks. The presence of multiple steady-states is analogous to the concept of games where the Nash equilibrium is not unique. To account for this uncertainty, we run multiple simulations for a particular initial topology and set of strategies by changing the order in which networks play. We then study the expected value of the steady-state properties. For example, the expected fitness for network i is the fitness of network i at steady-state, averaged over a number of permutations with different orders of play.

An important issue is the time complexity involved in determining the steady-state using agent-based simulations. Figure 3 shows the simulation time [2] for the scenario {CN,(SEL,TR),(SEL,NC)} as the number of networks grows, keeping the relative proportions of different network types fixed. Clearly, the running time of the model scales super-linearly with the number of networks. The main reasons for this are the complexity of computing the interdomain traffic flow, and the number of iterations to reach steady-state. Due to the fact that the running time of the model scales super-linearly with the number of networks, it is computationally infeasible to run the model at a scale larger than a few hundred networks, particularly as we need to run multiple simulations to investigate a wide parameter space and different variations of the canonical model.

L. How do we use this model?

We emphasize that the proposed model is not meant to be an Internet topology generator. We do not aim to produce a topology that matches structural properties of the actual Internet. Instead, the applications of the model lie in the ability to answer "what-if" questions. Here, we present specific questions that we plan to investigate in future work. We believe

[2]These simulations were run on a machine with with a 3GHz Intel Xeon processor and 2GB of memory.

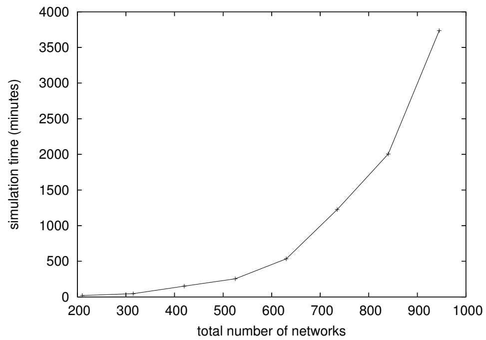

Fig. 3. Simulation time as a function of the number of networks.

that answers to these questions can give important insights into the possible evolution paths of the Internet ecosystem.

1) First, we plan to examine the steady-state networks that result from the canonical model. As described in Section III-J, we focus on STPs and LTPs. For each of these classes, we have 6 strategies of provider and peer selection, giving us a total of 36 combinations for these two classes. For each of these, we can examine the output metrics described in Section III-N. It is then possible to determine which combination of strategies leads to "favorable" steady-states, in terms of economics or performance.
2) We can determine the *optimal strategy* that each AS should use in order to obtain the maximum expected payoff at the steady-state. We can then investigate what happens when each AS plays its optimal strategy.
3) Our model accounts for various external factors such as the inter-AS traffic matrix, transit, peering and local cost structures, and the provider and peer selection strategies of ECs and CPs. We plan to examine the properties of the steady-state networks for different values of these factors. For example, what happens to the fitness of LTPs when the amount of p2p traffic in the Internet increases? Do LTPs make more or less money than if the traffic matrix was mostly client-server in nature?

M. Relation to game theory

Game theoretic models have been successfully applied to study other systems consisting of selfish, interacting agents. Here, we discuss the similarities and differences of our model with some popular game theoretic models. A large class of game theoretic models called static games are relevant for systems where agents make their moves simultaneously, and the potential payoffs are known in advance. In reality, networks do not make their decisions simultaneously, and may not always know the future payoffs. *Iterated games* consist of a number of repetitions of a base static game, and are also not applicable here for the same reason. The closest game theoretic model is a sequential game with incomplete and imperfect information. Networks are selfish and act in such a way as to optimize their own fitness. Networks play sequentially, and can potentially see the effects of previous moves made by other networks. Further, a network may not know about or anticipate the effects of possible moves that other networks make. A sequential game can be represented as an extensive-form game tree. Nodes in the tree represent every possible state of play for the game. The game begins at a unique initial node, and progresses through the tree along a path determined by the players actions until a terminal node is reached, where the game ends and payoffs are assigned to all players. At each non-terminal node, a player chooses among the possible moves at that node, which is an edge leading from that node to another node. Variations of the basic sequential game involve incomplete information (a player does not know the type of other players) or imperfect information (a player does not know the state of the game when she plays) have been studied in the literature. A key feature of each of these models, however, is that *payoffs obtained at each terminal node in the tree are known*. In this case it is possible to use methods such as *backwards induction* to determine the optimal move by a player at each stage of the game. In the case of the Internet ecosystem, however, it is hard to determine the payoffs that would be obtained from each combination of moves by networks. *These payoffs have to be determined numerically, by simulating, for each network, every possible move at each stage of the game.* Given the number of networks and the possible provider and peer selection options for each network, it is infeasible to determine the payoffs in this manner. Instead, we assign a particular provider and peer selection strategy to each *class of networks*, and an network plays this strategy at each stage of the game. It is then possible to determine the payoffs obtained from each of these strategies by running the model to completion.

N. Steady-state properties

Once we run the model as described previously, we characterize the resulting steady-state network. Here, we are interested not only in topological properties, such as degree distributions, but also metrics that pertain to economics and performance. We measure a set of output metrics that measure the properties of the steady-state network.

Topological properties: Examples of such properties are the degree distribution, the number of peering links, and customer-provider links between different types of networks. We measure the average path length between any two networks, which is an indicator of the global performance level. This metric can also be weighted by the traffic exchanged by those networks to get a *weighted path length*. We also measure properties related to traffic flow, such as the fraction of the total end-to-end traffic that flows over peering links.

Network-specific properties: Examples of such properties are the fitness of individual networks, the number of providers that have positive fitness at the steady-state (meaning that these providers are profitable), the number of customers for each provider, and the number of peering links established between different classes of providers.

Economic properties: A metric such as the total fitness of different classes of networks can be used to measure the economic performance of different classes of networks. (*e.g.,* are LTPs as a class profitable or not?). We also measure metrics that relate to the level of competition in the Internet Ecosystem, such as the number of profitable providers.

IV. Related Work

A major research effort aimed to characterize the AS-level topology during the last decade. One of the most well cited papers, by Faloutsos *et al.* [7], argued that the Internet AS-level topology is "scale-free". The observation that the degree distribution follows a power-law led to several topology generation models that could produce such distributions. These models focused on "growing" a topology that could match the Internet topology with respect to certain measurable graph metrics. The most well known work in this area is the preferential attachment model of Barabasi *et al.* [6]. Several variants and comparisons of preferential attachment models were later proposed [8], [9], [10], [11], [12], [13], [14] The models in this research thread have been mostly descriptive, meaning that they attempt to reproduce certain known structural characteristics of the Internet. They do not, however, attempt to explain how these properties emerge, or the domain specific interpretation of the various tunable parameters in those models.

The previous descriptive models received considerable criticism (for instance, see [15], [16]) because they mostly focus on the degree distribution and clustering, ignoring important characteristics of the Internet topology such as hierarchy or the presence of links of different types (transit versus peering). Further, the previous models do not explain how the Internet topology is evolving. This led to new models that view the Internet topology as a side effect of optimization-driven activity by individual ASes. These concepts were first introduced by Carlson and Doyle in [17], and later applied in the context of the Internet in [18] and elsewhere. Chang *et al.* [19] aimed to model AS interconnection practices, considering the effects of AS geography, AS business models and AS evolution. Norton [1] discusses, mainly using anecdotal evidence, how economic and competitive interests influence peering and transit connectivity in the Internet. Economides [20] discusses the economics of the Internet backbone (without looking at topology dynamics).

The body of work closest in spirit to ours is that of Chang et al. [5]. In that work, the authors focused on developing a model for the provider and peer selection behavior of ASes, taking into account the economics of transit and peering relationships and practical constraints such as geography. The focus was on "growing" a network using the local interactions between ASes. In this work, we focus mainly on the *rewiring* of links between existing ASes, and study the properties of the steady-state that results from the local optimizations of each AS. Also related is the work of Holme *et al.* [21], which developed an agent-based simulation model where the agents are individual ASes with economic incentives. In their model, each AS attempts to connect in such a way as to maximize its utility under a set of constraints. Their model captures the effects of economics, geography, user population and traffic flow in AS interconnection. They do not, however, model the presence of different classes of ASes with different incentives and business functions. Our model for network interactions accounts for different AS classes, policies in inter-AS relationships, realistic traffic flows and a detailed model for economic fitness of ASes. Corbo et al. [22] propose an economically-principled model that is able to create the observed structure of the AS graph. Their model considers the economic utility of an AS, and focuses on growing a network where each new AS tries to maximize its utility from connecting to the Internet. The goal of their work is mainly to derive, from first principles, a model that reproduces certain characteristics of the AS graph.

A series of papers [23], [24], [25] advocate the use of the Shapley value for revenue distribution between ISPs. They show that if profits were shared according to the Shapley value, the set of desirable "fair" properties inherent to the Shapley solution exist, and the selfish behavior of ISPs leads to globally optimal routing and interconnecting decisions.

A body of work known as "network formation games" [26], [27], [28] takes a *game theoretic* approach towards understanding the creation of bilateral contracts (interdomain links) between autonomous networks. These papers formulate a game where the nodes of the network (which can be Autonomous Systems) form a graph to route traffic between themselves. Variants of these models assign costs for routing traffic, as well as for a lack of end-to-end connectivity. The goal is for each node to create the set of links that maximizes its utility. A key difference of these models with ours is that they are *static* in nature; they model one-shot games where a node is able to predict the effects of creating a particular link. In contrast, we model the *dynamics* of network formation. Further, we consider the more realistic case where ASes do not play simultaneously, are able to observe the moves made by other players, and also the effects of those moves. Also, we assume that a network does not have the ability to predict the long-term effects of its actions.

References

[1] W. B. Norton, "The Evolution of the U.S. Internet Peering Ecosystem," *Equinix white papers*, 2004.

[2] A. Dhamdhere and C. Dovrolis, "Ten Years in the Evolution of the Internet Ecosystem," in *In the Proceedings of ACM SIGCOMM Internet Measurement Conference (IMC)*, 2008.

[3] L. Gao and F. Wang, "The Extent of AS Path Inflation by Routing Policies," in *Proceedings of IEEE Global Telecommunications Conference, GLOBECOM*, 2002.

[4] W. B. Norton, "Transit Cost Survey," *www.nanog.org/mtg-0606/pdf/bill.norton.2.pdf*, 2006.

[5] H. Chang, S. Jamin, and W. Willinger, "To Peer or Not to Peer: Modeling the Evolution of the Internet's AS-Level Topology," in *Proceedings of IEEE Infocom*, 2006.

[6] A. L. Barabasi and R. Albert, "Emergence of Scaling in Random Networks," *Science 286 509512*, 1999.

[7] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On Power-law Relationships of the Internet Topology," in *Proceedings of ACM SIGCOMM '99*, 1999.

[8] R. Albert and A. L. Barabasi, "Topology of Evolving Networks: Local Events and Universality," *Physical Review Letters 85, 5234*, 2000.

[9] T. Bu and D. Towsley, "On Distinguishing Between Internet Power Law Topology Generators," in *Proceedings of IEEE Infocom*, 2002.

[10] S. Park, D. M. Pennock, and C. L. Giles, "Comparing Static and Dynamic Measurements and Models of the Internet's AS Topology," in *Proceedings of IEEE Infocom*, 2004.

[11] M. A. Serrano, M. Boguna, and A. D. Guilera, "Modeling the Internet," *The European Physics Journal B*, 2006.

[12] X. Wang and D. Loguinov, "Wealth-Based Evolution Model for the Internet AS-Level Topology," in *Proceedings of IEEE Infocom*, 2006.

[13] S. H. Yook, H. Jeong, and A. L. Barabasi, "Modeling the Internet's Large-scale Topology," *Proceedings of the National Academy of Sciences*, 2002.

[14] S. Zhou, "Understanding the Evolution Dynamics of Internet Topology," *Physical Review E, vol. 74*, 2006.

[15] E. F. Keller, "Revisiting "Scale-free" Networks," *BioEssays 27, Wiley Periodicals Inc.*, 2005.

[16] D. Krioukov, kc claffy, M. Fomenkov, F. Chung, A. Vespignani, and W. Willinger, "The Workshop on Internet Topology (wit) Report," *ACM SIGCOMM CCR*, vol. 37, no. 1, 2007.

[17] J. M. Carlson and J. Doyle, "Highly Optimized Tolerance: A Mechanism for Power Laws in Designed Systems," *Physical Review E 60*, 1999.

[18] A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou, "Heuristically Optimized Trade-Offs: A New Paradigm for Power Laws in the Internet," in *Proceedings of ICALP*, 2002.

[19] H. Chang, S. Jamin, and W. Willinger, "Internet Connectivity at the AS-level: An Optimization-Driven Modeling Approach," in *Proceedings of ACM SIGCOMM Workshop on MoMeTools*, 2003.

[20] N. Economides, "The Economics of the Internet Backbone," *Handbook of Telecommunications Economics Ed. S. Majumdar, I. Vogelsang, M. Cave. Amsterdam: Elsevier Publishers*, 2006.

[21] P. Holme, J. Karlin, and S. Forrest, "An Integrated Model of Traffic, Geography and Economy in the Internet," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 3, pp. 5–16, 2008.

[22] J. Corbo, S. Jain, M. Mitzenmacher, and D. Parkes, "An Economically Principled Generative Model of AS Graph Connectivity," in *Proceedings of the International Joint Workshop on The Economics of Networke Systems and Incentive-Based Computing*, 2007.

[23] R. T. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein, "Interconnecting Eyeballs to Content: A Shapley Value Perspective on ISP Peering and Settlement," in *Proceedings of the ACM SIGCOMM 2008 Workshop on Economics of Networked Systems (NetEcon)*, 2008.

[24] ——, "On Cooperative Settlement Between Content, Transit and Eyeball Internet Service Providers," in *Proceedings of the ACM Conference on Emerging network experiment and technology (CoNEXT)*, 2008.

[25] R. T. B. Ma, D. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein, "Internet Economics: the use of Shapley Value for ISP Settlement," in *Proceedings of the ACM Conference on Emerging network experiment and technology (CoNEXT)*, 2007.

[26] E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler, "Near-optimal network design with selfish agents," in *STOC '03: Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*. New York, NY, USA: ACM, 2003, pp. 511–520.

[27] E. Anshelevich, B. Shepherd, and G. Wilfong, "Strategic network formation through peering and service agreements," in *FOCS '06: Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*. Washington, DC, USA: IEEE Computer Society, 2006, pp. 77–86.

[28] Ramesh Johari and Shie Mannor and John Tsitsiklis, "A contract-based model for directed network formation," *Games and Economic Behavior*, vol. 56, no. 2, pp. 201–224, August 2006. [Online]. Available: http://ideas.repec.org/a/eee/gamebe/v56y2006i2p201-224.html