A Model of Interdomain Network Formation, Economics and Routing

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

ABSTRACT

The Internet at the interdomain level is highly dynamic, as autonomous networks change their connectivity to optimize either monetary cost, profit and/or performance. Internet Service Providers (ISPs), for example, are mainly concerned with maximizing their profits, and they attempt to do so by changing their set of providers or peers. It is not well understood, however, what the properties of the resulting internetwork are, in terms of topology, economics and performance. In this paper, we propose ITER, a first-principles model of interdomain network formation that incorporates the effects of economics, interdomain traffic flow, geography, pricing/cost structures and interdomain routing policies. We use an agent-based compuational method (treating networks as selfish agents) to find the equilibrium that results as each network uses a certain provider and peer selection strategy (such as "peer by traffic ratios" or "peer by necessity"). We study the properties of this equilibrium in terms of topology, traffic flow and economics. We also investigate the effect of factors such as the interdomain traffic matrix, geography, and customer preferences on the properties of the equilibrium network.

1. INTRODUCTION

The Internet at the interdomain level is a system of interacting autonomous networks (ANs)[1] that connect to each other to provide end-to-end connectivity and access to various forms of content. The "Internet ecosystem" is dynamic, as ANs attempt to optimize, in a distributed manner, utility functions such as monetary profit, cost or performance. The utility that ANs are able to achieve depends both on "environmental" factors (transit prices, peering costs or the popularity of new Internet applications) and on their choice of providers and peers[2]. In practice, however, the process of provider and peer selection is often treated as "black art", even by network operators of large ISPs. These ISPs select their providers and peers using rules of thumb such as "peer by traffic ratios" or "peer restrictively". *A systematic evaluation of various provider/peer selection strategies is missing from the literature.*

The motivation behind this work was to create a framework that can be used to study the effects of provider/peer selection strategies used by different types of ANs. Note that the goal is not to recommend which exact networks an AN, i should choose as its providers or peers. That would require a precise knowledge of the strategy of every other AN, the interdomain traffic matrix, and pricing/cost parameters. Instead, the goal is to evaluate the effect of strategies such as "AN i peers with any network with which it exchanges roughly equal traffic". In this case, the lack of precise information does not prevent us from gaining insights into the effects of these strategies. We are interested in both *local effects* (how these strategies affect the ANs that use them), and *global effects* (how they affect the overall Internet). Studying the effects of provider and peer selection by ISPs is interesting for several reasons. First, individual networks would like to know which strategy would maximize their utility (either monetary profits or performance). Second, we would like to know the effects of these strategies on the global Internet, in terms of topological structure, profitability of various network types, and the risk of emerging monopolies or oligopolies. Third, it is important for ANs to understand how their provider/peer selection strategies perform under different conditions, such as diverse traffic characteristics and application popularity, different pricing structures, or new technology (*e.g.*, inexpensive transmission capacity).

The main contribution of this paper is a model, ITER, that provides the framework for answering questions of the aforementioned type. ITER is based on first-principles, and models the provider and peer selection process for different classes of ANs – Enterprise Customers (EC), Small and Large Transit Providers (STP and LTP), and Content Providers (CP). ITER takes as input the interdomain traffic matrix, routing policies, geographical constraints, and the economics of transit, peering and local costs. ITER models the interdependence between traffic flow, topology and the provider/peer selection strategies of ANs, as shown in Figure 1. The inter-

---

[1] ANs are similar to Autonomous Systems in BGP in the sense that they are independently operated, except that they also include networks that do not have AS numbers.

[2] A "provider" is a network that provides transit, or access to the rest of the Internet, to its customers. Two networks are "peers" if they engage in settlement-free interconnection, whereby they provide access to each other's customers for free.

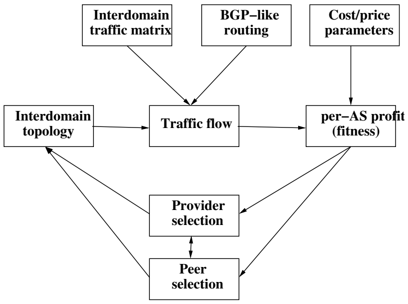

Figure 1: The interdependence between topology, traffic flow and per-AN fitness in the Internet ecosystem

domain traffic matrix, topology and routing policies determine the flow of traffic in the Internet. The traffic flow and economic factors together determine the utility of each AN (profit for transit providers and monetary cost/performance for ECs). ANs optimize their utility by changing their set of providers and peers, effectively changing the topology, which in turn can affect the utility of other ANs. The question we try to answer is, "Does this process to converge, and if so, where?", *i.e.*, we are interested in "solving" the model to find a state where no AN has the incentive to make further changes to its connectivity (if such a state exists). As ITER is intractable to solve analytically, we devise a method to solve it computationally, using agent-based simulations. We also study the existence and uniqueness of the resulting equilibrium. We emphasize that *ITER is not an evolutionary model*; it does not model the long-term evolution of the Internet ecosystem as ANs are born and die, application popularity changes and economic conditions fluctuate. Though it is important to study the evolution of the Internet, we argue that for the purpose of evaluating provider/peer selection strategies of ANs, the equilibrium of the static ITER model can give valuable insights.

In this paper, we focus on a first practical application of the ITER model, that of studying the properties of the equilibrium internetwork, given different provider/peer selection strategies used by ANs. In particular, we focus on two provider selection strategies (choose cheapest providers, or choose providers that are not competitors), and three peer selection strategies (peer only when necessary to maintain reachability, peer by traffic ratios, and peer when the potential benefit of peering is larger than the estimated cost) for small and large transit providers. We measure properties of the resulting network in terms of topology (*e.g.*, path lengths and diameter), traffic flow, profitability of different types of providers, and the number of providers that are profitable. We also analyze the effect of factors such as the interdomain traffic matrix, geography, and customer preference on the resulting internetwork. Specifically, we investigate what happens when the interdomain traffic matrix consists of mostly peer-to-peer (P2P) traffic, or if ANs at the edge of the Internet choose their providers based on path lengths, or if content providers replicate their content in all geographical regions. We envision several other applications of ITER, discussed in Section 10, which we plan to pursue in future work. We summarize the main findings in this paper:

- We find that if networks at the edge are price-conscious, then LTPs can benefit by peering with CPs, and can significantly harm the profitability of STPs; this comes at the cost of longer end-to-end paths (Section 5).
- We find that the STP strategy of peering using "balanced traffic ratios" is profitable only if they also use price-based provider selection. In this case, STPs should peer avoid peering with content providers. The choice of the best peering strategy for STPs is heavily influenced by their provider selection strategy (Section 5).
- We find that two conditions that are quite plausible in the future Internet – an interdomain traffic matrix with mostly P2P traffic , and content providers that replicate their content in all regions – result in increased profitability for STPs (Sections 6 and 8).
- We find that performance-aware provider selection by edge networks results in a situation where end-to-end paths are short and LTPs are profitable (Section 7).

The rest of this paper is organized as follows. Section 2 presents the details of the ITER model. Section 3 describes our approach for solving ITER using agent-based simulations. In Section 4 we validate the model against some well-known static and dynamic properties of the Internet. In Section 5, we present results for the default model, which we view as the current state of the Internet. In Section 6, we evaluate a deviation of the default model with a predominantly P2P traffic matrix. We evaluate a deviation where edge networks choose their providers based on performance in Section 7, and a deviation where Content Providers replicate their content in all geographical regions in Section 8. We survey the related work in Section 9, and conclude in Section 10 with a discussion of future applications of ITER.

2. MODEL DESCRIPTION

In this section, we summarize the key features of ITER[3]. Table 1 lists all the acronyms and that will be used here and in subsequent sections.

2.1 Network types

Enterprise Customers (EC): ECs are stub networks that normally act as either mostly sources of traffic (*e.g.*, web hosting companies), or mostly sinks of traffic (*e.g.*, campus, corporate or residential access networks). In ITER, ECs

---

[3] An earlier version of the ITER model was described in a 6-page invited (not peer-reviewed) conference paper that we do not cite to preserve anonymity. That paper does not include any quantitative results.

acronym	definition
AN	Autonomous Network
EC	Enterprise Customer
STP	Small Transit Provider
LTP	Large Transit Provider
CP	Content Provider
CS	Client-Server
P2P	Peer-to-Peer
PR	Price-based provider selection
PF	Performance-based provider selection
SEL	Price-based Selective provider selection
NC	Peering by necessity
TR	Peering by traffic ratios
CB	Peering by cost-benefit analysis
DF	Default Model
P2P	Deviation: P2P traffic matrix
EP	Deviation: edge networks use performance based provider selection
GEO	Deviation: content providers present in each geographical region

Table 1: Definitions of acronyms used

do peer and they do not have customers; their only action is provider selection. We model a fraction of ECs as traffic sinks (*sink-ECs*), while the remaining as traffic sources (*source-ECs*).

Content providers (CP): CPs are also stub networks that differ from ECs in two ways. First, they are sources of traffic (*e.g.*, Yahoo!, Google). Second, they can engage in peering relations, following an "open peering" policy (peer with any network that agrees to peer with them).

Small Transit Providers (STP) and Large Transit Providers (LTP): Transit providers are networks whose main business function is to provide Internet connectivity to their customers. In ITER, transit providers do not act as sources or sinks of traffic; they only carry traffic on behalf of other networks. Transit providers aim to maximize their profit and so they select their providers and peers with this economic objective. Their peering policies are often described as "restrictive" or "selective", in practice. STPs are transit providers with limited geographical presence (*e.g.*, Rogers Telecom or China Telecom), while LTPs are transit providers with practically global presence (*e.g.*, AT&T or Level3).

In the default ITER model, we simulate 180 ECs, 10 CPs, 16 STPs, and 4 LTPs. 20% of the ECs act as source-ECs, while the rest are sink-ECs. This 210-node internetwork is of course small compared to the real Internet (the number of Autonomous Systems is about 30,000 today) to keep the simulation time tractable; we will return to this scalability issue in Section 3.

2.2 Traffic model

The traffic model concerns the generation of an inter-AN traffic matrix. This matrix determines the amount of traffic sent from each AN to every other AN. In ITER, we consider two types of traffic: Client-Server (CS) traffic flows from traffic sources, which are either CPs or source-ECs, to sink-ECs (*e.g.*, YouTube or RapidShare). Peer-to-Peer (P2P) traffic flows between sink-ECs (*e.g.*, BitTorrent). Without showing the actual mathematical expressions, the key points of the traffic model are the following. The total traffic volume (both CS and P2P) destined to each traffic sink is heavy-tailed (Pareto distributed with shape parameter=1.1), *i.e.*, few sink-ECs are much larger traffic consumers than most other sink-ECs. Traffic sources are ranked based on a *popularity index*. CPs have higher popularity index than source-ECs. The fraction of traffic from a given source to any sink-EC follows a Zipf distribution (with exponent 0.8), determined by the previous popularity ranking. The Zipf distribution implies that few traffic sources, mostly CPs, are much heavier traffic producers than most other sources. For simplicity, we assume that the popularity of a source is the same for all sink-ECs, ignoring any regional content preferences. A similar popularity index for each sink-EC determines the distribution of P2P traffic between sink-ECs. In the default ITER model, 80% of the overall traffic is CS while the rest is P2P.

2.3 Geographical constraints

In ITER, each AN is geographically present in a certain set of locations (*e.g.*, exchange points or "GigaPoPs"). Two ANs cannot establish a customer-provider or peering relation unless they are present in a common location. In the default ITER model, 210 ANs are distributed in 5 locations. ECs and CPs are present in one location, STPs in 2, and LTPs in all 5 locations.

2.4 Routing and traffic flow

ITER captures the salient features of interdomain routing. Specifically, traffic follows the "no-valley" policy, (traffic from a provider cannot be sent to another provider, and traffic from a peer cannot be sent to another peer), as well as the "prefer-customer" policy (prefer a route that goes through a customer; if not available, prefer a route that goes through a peer; otherwise route through a provider). Whenever multiple preferred neighbors offer a route, choose the shortest path; break ties deterministically based on the neighbor's AN number. It should be noted that calculating policy-compliant shortest paths between all pairs of nodes is computationally expensive (O(N^3), where N is the number of ANs in the internetwork). We have simplified the routing computation, without violating the previous policies, with an algorithm inspired by the method proposed by Gao and Wang [16]. 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 time $O(N_pE_p)$, 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 time $O(N_pN_sd_s)$, where d_{p_s} is the multihoming degree of stubs. Finally, each stub s determines the best path towards stub s'.

To do this, s chooses the provider p from among its set of providers that gives the shortest path towards s'. The final step can be done in time $O(N_s^2 d_s)$.

Given the inter-AN traffic matrix, the interdomain topology and the routing model, we can then calculate the traffic flow in the internetwork. The traffic flow determines the aggregate amount of traffic that flows over each link and AN. These per-link traffic loads are then used by the economic model, described next.

2.5 Economic model

The economic component of ITER focuses on the profit of transit providers. STPs and LTPs adjust their provider and peering selections so that they maximize their profit. The profit of a transit provider is calculated as the total revenue from its customers, minus the transit fees to its providers (if any), the peering costs (if any), and the local costs to maintain and operate its network. Let $\mathcal{C}_i$ be the set of customers, $\mathcal{P}_i$ the set of providers and $\mathcal{R}_i$ the set of peers of a transit provider i. Its profit f_i is:

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

$T_i(v)$ represents the pricing function used by provider i for a transit volume v (*i.e.*, volume-based pricing). $T_i(v_{ic})$ gives the transit payment made by customer c to provider i when the aggregate traffic exchanged by the two networks is v_{ic}. $T_p(v_{pi})$ is the transit payment made by i to its provider p for the traffic volume v_{pi}. $R_i(v_{ir})$ is the cost of maintaining a peering link between i and its peer r when the corresponding traffic volume is v_{ir}. This fee is not paid by one peer to the other; rather, it represents costs to setup (amortized over time) or maintain that peering link. $L_i(v_i)$ determines local costs incurred by AN i (such as operations, staff, equipment) when the aggregate traffic handled by i is v_i.

In practice, transit prices show *economies of scale* meaning that the per-bit cost of Internet transit decreases as the volume of traffic increases. In ITER, we use concave increasing functions for transit, peering and local cost functions. Specifically, the pricing function of a transit provider p for traffic volume v is given by

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

The exponent e_t controls the extent of the economies of scale associated with the various costs; a lower value of the exponent results in larger economies of scale. All transit providers have the same exponent e_t but they differ in the multipliers $m_{t,p}$. This is consistent with pricing data we collected from Norton [25] and Chang [10]. Similarly, peering costs are calculated as:

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

while local costs also include a traffic-independent term l_i:

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

All transit providers are assigned the same exponents for their peering and local cost functions, but they differ in the multipliers $m_{r,i}$, $m_{l,i}$, and in the l_i term.

To the extent possible, we parameterized the economic model using real-world data. Chang et al. [10] report that the exponent for the transit pricing functions e_t is around 0.75, while the peering exponent e_r is around 0.25. The transit price multipliers $m_{t,i}$ of STPs vary between [30,140], while those of LTPs vary between [80,150], *i.e.*, LTPs tend to be more expensive than STPs, but not always. These values are based on data reported by Norton [25] in 2006. The peering cost multipliers $m_{r,i}$ vary in [300,400]. The local cost exponent e_l is set to 0.5, while the local cost multipliers are set differently for STPs and LTPs: [100,200] for STPs and [300,400] for LTPs. The traffic-independent costs for LTPs are greater than those for STPs; this reflects that LTPs have larger networks, and hence larger operational costs. The local cost parameters are assigned so that the traffic-dependent and traffic-independent costs account for roughly equal fractions. The transit, peering and local cost parameters are assigned so that, for the same traffic volume, peering costs are the lowest, followed by traffic-dependent local costs, while transit costs are the highest.

2.6 Provider selection methods

The interdomain topology is formed when each AN selects its provider(s), and potentially its peers. In ITER, we consider three provider selection methods and three peer selection methods. Even though these methods are, to some degree, abstractions of a wide diversity of service agreements in the Internet, we believe that they capture the most common practices.

Regarding provider selection, an AN i first determines the set of candidate providers. These are transit providers (STPs or LTPs) that have at least one region in common with i and that are *not* in the customer tree of i. Then, i uses one of the following three methods to select the final provider (or set of providers, in case of multihoming):

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

Selective price-based (SEL): A transit provider i would not want to select as provider a network that may become its peer or customer in the future. In ITER, this implies that an STP would not want to select another STP as provider, and so it would choose only among LTPs. Similarly, an LTP would not select an STP as provider, even if it is cheaper than LTP candidate providers. Among the remaining candidates, i would again select provider(s) based on price. SEL is applicable only to STPs and LTPs.

Performance-based (PF): A network may select providers based on the performance they offer. In ITER, we consider a performance metric that is related to the weighted path length from i to all sources and destinations of its traffic. This method is applicable only to ECs and CPs, not to transit providers (the latter would certainly not ignore pricing). 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. The performance metric associated with provider k is given by $L_i(k) = \sum_j A_{ij} l_{kj} / \sum_j A_{ij}$.

2.7 Multihoming

Multihoming, which refers to the practice of choosing multiple transit providers, is increasingly used, particularly by transit providers [13]. In ITER, AN i is assigned a *Maximum Multihoming Degree* (MMD), *i.e.*, a maximum number of providers, depending on its type. This upper bound is typically determined by the desired redundancy level. In practice, it may not be possible to always find MMD candidate providers. AN i ranks its set of candidate providers, based on one of the previous three selection methods, and selects up to MMD providers. In the default ITER model, we set the MMD to 1 for ECs, 3 for CPs, 2 for STPs and 3 for LTPs.

2.8 Peer selection methods

For any AN, the objective for peering is to save transit costs, by reducing the traffic volume that needs to be routed through providers. Further, peering is required in some cases to maintain reachability with the rest of the Internet. We consider three peer selection methods, modeling the most common approaches found in practice.

Peering by necessity (NC): With NC, networks i and j peer only if that is necessary to maintain global reachability; otherwise i will not be able to reach some of j's customers and vise versa. Neither AN can "force" the other to become its customer. Also, in some cases i and j would choose each other as provider based on their provider selection method. When that is the case, they decide to peer instead.

Peering by traffic ratios (TR): A common approach for peering is to rely on "traffic ratios". Here, two ANs i and j agree to 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 close to one (within a factor of 2 in default ITER), the two ANs agree to peer.

Peering by cost-benefit analysis (CB): Here, AN i assesses both the costs associated with a given peering link and the potential benefits that can be achieved by that link. The costs associated with peering are due to the fixed and traffic-dependent costs of establishing a peering link. The benefits are due to reduced transit fees. AN i chooses to peer with j if the estimated benefits are greater than the estimated costs. In practice, i would need to estimate the "peerable traffic volume" with network j to use CB.

2.9 Initialization

We construct the initial internetwork so that it matches certain known properties of the Internet's interdomain topology. First, LTPs are assumed to be present in each geographical region and are fully-meshed with peering links. This is similar to the well-known clique of Tier-1 Internet providers. These are the only peering links in the initial topology. Regarding the initial customer-provider links, a recent study [13] measured the provider preference of different network types in the Internet and found that 60% of the providers of ECs are STPs while 40% are LTPs. On the other hand, approximately half of the providers of STPs and CPs are STPs. So, we connect STPs to other STPs and LTPs so that the number of links between STPs and LTPs is the same with that between STPs and STPs. To connect ECs and CPs, we follow a procedure that is similar to preferential attachment. We add ECs and CPs sequentially, choosing a provider (STP or LTP) with a probability that is proportional to the existing customer degree of that provider.

We define a *scenario* as a specification of the provider and peer selection strategies used by STPs and LTPs. In a scenario, we assume that *all providers belonging to the same class follow the same strategy*. For example, the notation

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

represents a scenario with the default ITER model (DF), STPs use SEL provider selection and TR-peering, and LTPs use SEL-provider selection and NC-peering.

3. SOLVING THE MODEL

Our goal is to to "solve" the model, determining the internetwork that results as each AN changes its set of providers and peers to optimize a certain utility function. ANs play *sequentially*, and each AN i can observe how the actions of previous ANs affect i's traffic flow and economics.

3.1 AN actions

We present the steps used by an AN in each move.

1. **Provider selection:** First, an AN 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 AN i does not engage in peering, 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. AN 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. [4]

---

[4] A corner case can occur when i needs providers to reach ANs that

3. **Check for potential peering candidates:** AN 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 *unilaterally* verifies whether the peering requirements with i are satisfied. AN 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 for a peering link to be created. 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 loop. Note that in one move, i may add or remove only one peering link.

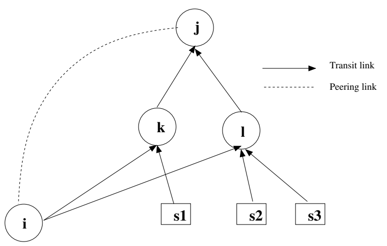

Figure 2: AN i can remove providers k and l after forming a peering link with provider j.

Note that all the actions performed by an AN in each move are *completely deterministic*. This is in contrast to previous evolutionary models of Internet topology (such as those based on preferential attachment [6]). Those models generate a random graph that has certain structural properties such as a desired degree distribution. The ITER model is not intended to be a topology generator. Instead, ITER models the optimizations performed by ANs, in terms of selecting providers and peers. These optimizations are essentially deterministic in nature, as each AN attempts to unilaterally maximize its utility function.

3.2 Computing equilibrium

Our goal is to "solve" the ITER model, computing an equilibrium, given the initialization and the strategy of each AN. An equilibrium, if it exists, is a situation where no AN has the incentive to unilaterally change its set of providers or peers. We solve ITER computationally, as it is too complex to solve analytically. Solving ITER involves iteratively allowing an AN to play (according to its pre-defined strategy in each move), until we reach a stage where no AN has the incentive to change its connectivity. This state is analogous to the concept of Nash Equilibria (or pairwise stable equilibria when bilateral peering contracts are involved) in game theoretic models. We assume that nodes play in a particular sequence, with a randomly chosen starting node. We use the following procedure to compute the equilibrium for ITER.

1. Pick the next AN i in the playing sequence.
2. Complete the move of AN i, as described in section 3.1.
3. If the move of AN i causes the topology to change, recompute the routing tables, traffic flow and fitness function of each AN.
4. Check termination criteria. If each AN has had a chance to play and has not changed its connectivity, then stop.

An important issue is the time complexity involved in finding an equilibrium using agent-based approach described above. Figure 3 shows the simulation time [5] for the scenario $\{$DF, (SEL,TR), (SEL,NC)$\}$ as we increase the number of ANs, keeping the relative proportions of different AN types fixed. We find that the running time of the model scales super-linearly with the number of ANs. The main reasons for this are the complexity of computing the interdomain traffic flow, and the number of iterations to reach equilibrium. As a result, it is computationally infeasible to run the model at a scale larger than a few hundred ANs, particularly as we need to run multiple simulations to investigate a wide parameter space and different variations of the default model.

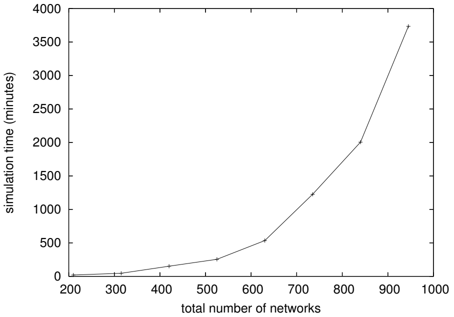

Figure 3: Simulation time to find an equilibrium vs. the number of ANs.

3.3 Existence of equilibrium

An important question is whether the agent-based simulation described in Section 3.2 is always able to find an equilibrium for ITER. We find empirically that in more than 95% of

---

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.

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

the simulation instances, we are able to solve ITER to find an an equilibrium. We find that 80% of the cases where we cannot find an equilibrium occur when STPs use CB-peering. In cases where we cannot find an equilibrium, the oscillation is caused by a small number of ANs, and *this oscillation is an expected outcome of the interaction between provider and peer selection, and traffic flow, and performance in the internetwork*. Next, we present some cases where ITER does not have an equilibrium, focusing on the fundamental reasons behind the oscillations.

In figure 4(a), AN 25 (a content provider) is connected to its preferred providers 1, 5 and 10, and the peering link with 12 does not exist. AN 25 uses its provider link to 10 to reach ANs in the customer tree of 10. When 12 uses CB-peering, it finds that peering with 25 leads to a higher fitness. This is because 25 now uses the (free) peering link with 12 to reach ANs in the customer tree of 10, due to which 12 earns revenues from 10. After the peering link between 25 and 12 is formed, 25 no longer needs 10 as a provider, and removes the link to provider 10. When 25 removes the provider link to 10, 12 no longer sees a benefit in peering with 25, and de-peers 25. As the peering link betwen 12 and 25 is removed, 25 is again able to choose its preferred providers, which includes AN 10. The above sequence then repeats. The fundamental factor that causes this oscillation is the interaction between provider and peer selection. An AN that creates a peering link with a provider does not need to retain providers that are in the customer tree of peers.

In figure 4(b), AN 8 and 10 both use TR-peering. Content stubs 110, 169 and 177 use PF provider selection, and are initially not connected to 8. In this situation, the traffic ratio between 8 and 10 is balanced, and 8 is able to peer with 10. Due to this peering link, 8 obtains shortcut paths to nodes in the customer tree of 10, and becomes more attractive for content stubs 110,169 and 177 due to shorter weighted path lengths. These content stubs connect to 8 as customers. This affects the traffic flow between 8 and 10, whereby 8 sends more traffic 10 on the peering link. When 10 evaluates the peering link, it finds that the traffic ratios are no longer balanced. This causes 10 to de-peer 8. Consequently, 8 loses the advantage (attractiveness for performance-oriented customers) from peering with 10, and the content stubs 110, 169 and 177 no longer prefer to connect to 8 as provider. After the content stubs depart, the traffic ratio between 8 and 10 is again balanced, and 8 can peer with 10. The above sequence then repeats. The fundamental reason for this oscillation is that the creation of a peering link between two providers can improve (or harm) the weighted path lengths that either provider can offer to customers. The peering criterion (either traffic ratio or cost-benefit analysis) between the two peers could now fail as customers are attracted (or repelled) from this provider.

In a third example, the topology is as shown in figure 4(c). STPs 9 and 10 both use CB-peering, and initially, the peering links 9-10 and 10-19 are not present. Traffic from customers of 10 to the common customers of 9 and 19 (such as C) initially follows the path 10-0-9-C. Using CB-peering, STP 10 adds 19 as a peer, as both see a benefit. Now traffic from 10 to C flows over the peering link between 10 and 19 (path 10-19-C). This causes traffic to shift away from 9, leading to a loss of revenue. Using CB-analysis, STP 9 finds that creating a peering link with 10 would serve to bring traffic back to 9, leading to better fitness. Consequently, 9 and 10 form a peering link using CB-peering. Once the peering link between 9 and 10 is formed, 10 does not see a benefit in keeping the peering link to 19. After the link 10-19 is removed, 9 finds that it would achieve better fitness without the peering link with 10. Hence, 9 de-peers 10. The above sequence then repeats. The underlying reason for this oscillation is that the creation of a peering link alters the traffic flow, affecting the profitability of other networks and leading to the creation/removal of other peering links.

3.4 Uniqueness of equilibrium

An important issue is the uniqueness of the equilibrium that results from solving ITER using the method described in Section 3.2. We find that for a given initial topology and set of AN strategies, *the equilibria can depend on the order in which ANs make their moves*. In some cases ANs make the "right move at the right time", such as forming a particular peering link or choosing a certain provider, causing different equilibria. The presence of multiple equilibria is analogous to game theoretic models where the Nash equilibrium is not unique. To account for this uncertainty, we run multiple simulations for a given initial topology and set of strategies by changing the order of play for ANs. We then study the expected value of the properties of the resulting equilibrium network. For example, the expected fitness for AN i is the fitness of AN i at equilibrium, averaged over a number of permutations with different orders of play.

4. MODEL VALIDATION

A major problem with any model that aims to capture, not only the interdomain topology, but also the economics and the traffic flow in the Internet, is how to validate it. ISPs are secretive about their economic and traffic data, while the ground truth for the Internet topology remains elusive (especially for peering links) [11]. In this section, we present a "best-effort" approach to validate ITER, comparing its predictions with known quantitative and qualitative characteristics of the Internet. These characteristics span both static and dynamic topological properties, as well as some basic facts about Internet economics and distribution of traffic load. Clearly, however, the following results cannot be viewed as a definitive validation, given that other models may also be able to reproduce the same properties.

The following results are based on the following scenarios, {DF, (SEL,CB), (SEL,NC)}, {DF, (PR,TR), (SEL,NC)} and {DF, (PR,CB), (SEL,NC)}, which we view as the most common provider/peer selection scenarios in practice. The

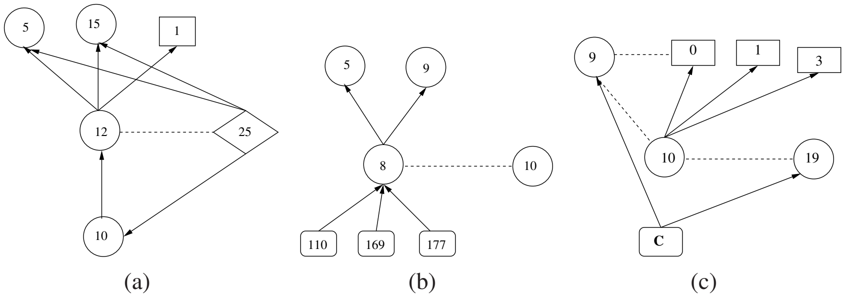

Figure 4: Examples of cases that lead to oscillations

differences between these three scenarios, in terms of the following observations, are minimal.

Degree distribution: Figure 5 shows the complementary CDF (C-CDF) of the degree distribution for the scenario $\{$DF, (SEL,CB), (SEL,NC)$\}$ with 945 networks. Even though it is not possible to be rigorous about the presence of a power-law in such a small scale, it is clear that the degree distribution is heavy-tailed. Of course this should not be surprising. In the default ITER, we set the multihoming degree of ECs and CPs to 1-3 providers, while STPs and LTPs can attract many customers at their regions, and so few of them will necessarily end up with large degrees. We also see the presence of networks with intermediate degrees, indicating that a single "attractor" network does not end up with all other networks as its customers or peers.

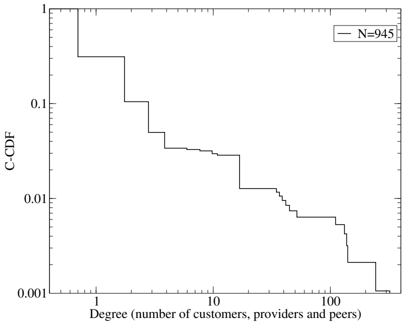

Figure 5: Degree distribution for an internetwork with 945 ANs $\{$DF, SEL,CB), (SEL,NC)$\}$.

Average path length: Another property of the Internet is that the average path length, in terms of AS links, has remained almost constant (at about 4 AS hops) during the last decade [21, 13]. We have reproduced the same behavior in ITER. Figure 6 shows the average path length in the network for the scenario $\{$DF, (SEL,CB), (SEL,NC)$\}$ as the number of ANs is increased from 210 to 945. We find that the average path length between any two ANs remains close to 4 hops (with a variation range between 3 to 5 hops, which also does not vary with the size of the internetwork).

Activity frequency: We next examine the dynamics as each AN executes its provider and peer selection methods. To do so, we analyze the transient phase of the simulation, before the internetwork converges. The *activity frequency* of an AN is defined as the ratio of the number of times that that AN changed its connectivity to the number of times that AN "played". A recent measurement study found that ECs are the least active in changing their connectivity, while providers (STPs, LTPs) and CPs are the most active [13]. The ITER simulations show the same qualitative trends: the activity frequency of ECs is less than that for STPs, LTPs and CPs, as shown in Figure 7. The higher activity frequency for providers results in part from the interaction between provider and peer selection in these network types. ECs do not peer, and they do not have a reason to change their upstream providers as often as other network types.

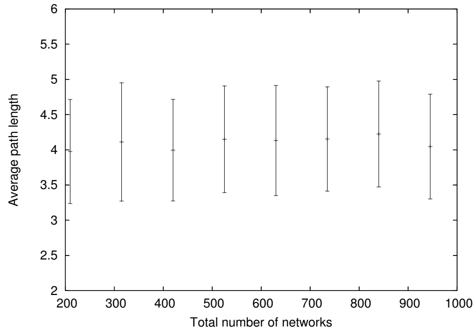

Figure 6: Average path length as the number of ANs is increased for scenario $\{$DF, (SEL,CB), (SEL,NC)$\}$.

Economic structure of transit market: We also examine the profitability of transit providers in the resulting ITER internetwork. We find that a significant fraction of STPs and LTPs fail to attract enough customers, and so they end up with negative "profits". In an evolutionary version of ITER, these ANs would be removed as bankrupt or "dead", similar to what often happens in the real Internet. On the other hand, there are several profitable STPs and LTPs, meaning that the ITER internetwork does not converge to a monopoly or oligopoly. This is in agreement with a recent measurement study [13] which showed that the number of transit providers that are active (meaning that they attract customers) is significant, indicating that the Internet transit market is not head-

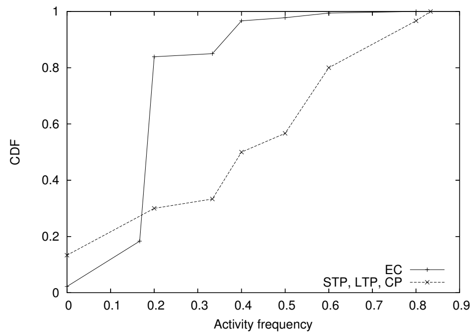

Figure 7: Activity frequency of each class of ANs for scenario $\{$DF, (SEL,CB), (SEL,NC)$\}$.

ing towards a monopoly or oligopoly.

Distribution of link load: We also measure the traffic volume carried by each link in the ITER internetwork. Figure 8 shows the C-CDF of the link loads on each interdomain link for the scenario $\{$DF, (PR,TR), (SEL,NC)$\}$. Most links carry small traffic loads; these are links mostly at ECs and CPs at the edge of the Internet. On the other hand, there are few links that carry very large traffic volumes; these are customer-provider and peering links between transit providers. Akella et al [2] observed a qualitatively similar phenomenon in the Internet. They reported that links between transit providers high in the hierarchy are typically of higher capacity than those between providers lower in the hierarchy.

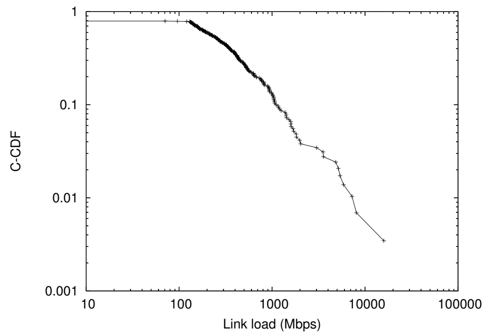

Figure 8: C-CDF of traffic volume on each link for scenario $\{$DF, (PR,TR), (SEL,NC)$\}$.

5. THE DEFAULT MODEL

In the rest of this paper, our goal is to understand the impact of different provider/peer selection methods on the topology, traffic flow, economics and performance of the resulting internetwork. In this section, we focus on the Default ITER model. In the following three sections, we consider a number of deviations from the Default model, in terms of the traffic matrix, the edge network provider preferences and the geographical presence of CPs.

In the default ITER model, ECs use PR provider selection and they do not peer with other ANs. CPs also use PR provider selection, but they peer using the CB method. For STPs, as well as for LTPs, we consider two provider selection methods, PR and SEL, and three peer selection methods: NC, TR and CB. All ANs of the same type choose the same provider and peer selection method. This agrees with what we see in the Peering Database [1], for instance, where networks of the same business function and size tend to use the same type of peering policy.

An *ITER scenario* refers to the selection of a specific pair of provider and peer selection methods for STPs and of another such pair for LTPs. Since we have 6 provider/peer combinations for STPs and 6 identical combinations for LTPs, the total number of scenarios we need to consider is 36. Table 2 shows the output metrics for each of these 36 scenarios in the default ITER model. For each scenario, we run 20 ITER simulations. In each simulation, we use a different random permutation of the sequence in which ANs move during the ITER transient phase.

We measure several metrics that characterize the equilibrium network: The average path length between each pair of ANs (unweighted as well as weighted by the traffic that flows between those ANs), the aggregate fitness of STPs and LTPs, the number of fit STPs and LTPs, the fraction of peering links and the fraction of total traffic that flows over peering links. The results in Table 2 are averaged over that subset of the 20 runs in which ITER converged to a stable internetwork. The standard error for each metric is also shown. We compare various scenarios only when the corresponding confidence intervals are non-overlapping.

5.1 Path Lengths

We report the weighted path length (column "wPL" in Table 2) and the unweighted path length (column "uPL" in Table 2) for each scenario of the Default model. Note that the average path length in the resulting internetwork is close to 4 hops for all scenarios except when LTPs use CB; paths tend to be longer when LTPs use CB-peering. In particular, the scenario $\{\{$DF, (PR,NC), (SEL,CB)$\}$ results in average path length of 4.2, compared to 3.9 in other scenarios. When LTPs peer with CPs the traffic from CPs goes through peering links to LTPs, and from there to ECs potentially through one or more STPs. Figure 9 illustrates this case for scenario $\{$DF, (PR,NC), (SEL,CB)$\}$. We see paths of the following nature: LTPs peer with several CPs using the CB method. The path from these CPs to destination ECs (which are customers of say STP A) is of the form CP-LTP-STP_B-STP_A-EC. If LTPs do not use CB, they will not form peering links with CPs (TR peering would not work because CPs always generate much more traffic than they consume). The CPs would then probably choose STPs as providers, as they tend to be less expensive than LTPs. This leads to paths of the form CP-STP_A-EC or CP-STP_A-STP_B-EC that are shorter than the path observed when the

sc	STP str	LTP str	wPL	uPL	diameter	profit all STP ($k)	profit all LTP ($k)	profit fit STP ($k)	profit fit LTP ($k)	num fit STP	num fit LTP	num UA	Traf UA	%PP	Traf PP
s1	PR,NC	PR,NC	3.9	3.9	6.1	331	409	446	527	4.4	1.6	1.6	0.1	2.3	0.1
s2	PR,NC	PR,TR	3.9	3.9	6.1	331	409	446	527	4.4	1.6	1.6	0.1	2.3	0.1
s3	PR,NC	PR,CB	4.2	4.0	6.6	40	439	180	521	4.0	2.0	2.1	0.2	6.3	0.4
s4	PR,NC	SEL,NC	3.9	3.9	6.1	355	368	465	495	4.8	1.5	1.4	0.0	2.9	0.1
s5	PR,NC	SEL,TR	3.9	3.9	6.1	355	368	465	495	4.8	1.5	1.4	0.0	2.9	0.1
s6	PR,NC	SEL,CB	4.2	4.0	6.6	39	441	179	523	3.9	2.0	2.2	0.2	6.3	0.4
s7	PR,TR	PR,NC	3.9	3.9	6.2	335	393	451	504	4.5	1.7	1.6	0.0	2.5	0.1
s8	PR,TR	PR,TR	3.9	3.9	5.9	317	426	433	544	4.2	1.6	1.6	0.0	2.2	0.2
s9	PR,TR	PR,CB	4.1	3.9	6.3	55	458	200	545	3.0	2.1	1.9	0.2	6.0	0.5
s10	PR,TR	SEL,NC	3.9	3.9	6.2	347	369	459	491	4.9	1.5	1.4	0.0	3.0	0.1
s11	PR,TR	SEL,TR	3.9	3.9	6.0	301	431	416	546	4.3	1.7	1.4	0.0	3.1	0.2
s12	PR,TR	SEL,CB	4.1	3.9	6.3	24	480	173	554	3.1	2.2	2.0	0.2	6.7	0.4
s13	PR,CB	PR,NC	3.9	3.9	6.0	333	392	445	502	4.5	1.7	1.5	0.0	3.3	0.2
s14	PR,CB	PR,TR	3.9	3.9	5.8	229	498	344	602	4.1	1.9	1.0	0.0	3.4	0.2
s15	PR,CB	PR,CB	3.9	3.9	6.0	63	472	209	538	2.5	2.5	1.4	0.2	7.4	0.5
s16	PR,CB	SEL,NC	3.9	3.9	6.0	243	471	352	576	4.6	1.9	1.3	0.0	4.0	0.2
s17	PR,CB	SEL,TR	3.9	3.9	5.9	226	501	340	605	4.2	1.9	1.0	0.0	3.9	0.2
s18	PR,CB	SEL,CB	3.9	3.9	5.9	33	492	183	551	2.4	2.5	1.7	0.1	8.4	0.5
s19	SEL,NC	PR,NC	3.9	3.9	5.0	-48	851	92	951	3.0	2.0	1.0	0.0	2.1	0.0
s20	SEL,NC	PR,TR	3.9	3.9	5.0	-48	851	92	951	3.0	2.0	1.0	0.0	2.1	0.0
s21	SEL,NC	PR,CB	4.0	3.9	5.0	-185	787	2	873	1.0	2.0	3.0	0.5	5.9	0.3
s22	SEL,NC	SEL,NC	3.9	3.9	5.0	-48	851	92	951	3.0	2.0	1.0	0.0	2.1	0.0
s23	SEL,NC	SEL,TR	3.9	3.9	5.0	-48	851	92	951	3.0	2.0	1.0	0.0	2.1	0.0
s24	SEL,NC	SEL,CB	4.0	3.9	5.0	-185	787	2	873	1.0	2.0	3.0	0.5	5.9	0.3
s25	SEL,TR	PR,NC	3.9	3.9	5.0	-48	851	92	951	3.0	2.0	1.0	0.0	2.1	0.0
s26	SEL,TR	PR,TR	3.9	3.9	5.0	-6	799	134	899	2.9	2.0	1.1	0.0	2.0	0.1
s27	SEL,TR	PR,CB	3.9	3.9	5.0	-104	701	83	772	1.2	2.3	2.5	0.4	5.7	0.4
s28	SEL,TR	SEL,NC	3.9	3.9	5.0	-48	851	92	951	3.0	2.0	1.0	0.0	2.1	0.0
s29	SEL,TR	SEL,TR	3.9	3.9	5.0	-10	806	129	906	3.0	2.0	0.9	0.0	2.3	0.1
s30	SEL,TR	SEL,CB	3.9	3.9	5.0	-110	709	76	782	1.0	2.3	2.7	0.5	6.1	0.4
s31	SEL,CB	PR,NC	3.8	3.9	5.0	65	734	182	834	4.0	2.0	1.0	0.0	3.6	0.1
s32	SEL,CB	PR,TR	3.8	3.9	5.0	113	680	233	776	3.7	2.0	1.3	0.0	3.1	0.2
s33	SEL,CB	PR,CB	3.9	3.9	5.0	-40	597	122	653	2.0	2.7	1.4	0.2	6.2	0.5
s34	SEL,CB	SEL,NC	3.8	3.9	5.0	65	734	182	834	4.0	2.0	1.0	0.0	3.6	0.1
s35	SEL,CB	SEL,TR	3.8	3.9	5.0	115	680	231	780	3.8	2.0	1.3	0.0	3.7	0.2
s36	SEL,CB	SEL,CB	3.9	3.9	5.0	-46	622	115	687	2.0	2.4	1.6	0.2	6.7	0.5
standard error			0.02	0.01	0.07	24	29	23	28	0.13	0.07	0.14	0.01	0.13	0.02

Table 2: Default model, DF

LTP peers with CPs. Also, these longer paths are from the major sources of traffic (CPs) to their destinations (ECs). So, the weighted path length (4.2) is longer than the unweighted path length (4.0).

5.2 Peering links

As expected, we see that there is a positive correlation between the percentage of peering links ("%PP") and the fraction of the total traffic flow that traverses at least one peering link ("Traf-PP"). Both of these metrics are maximized when STPs and LTPs both use CB-peering. In those scenarios, 6-8% of all links are PP links, and 50% of the total end-to-end traffic flows over those links. In those scenarios, both STPs and LTPs are able to peer with CPs. The large traffic volume from CPs to ECs now flows through those peering links.

5.3 "Unprofitable-but-Active" (UA) providers

We evaluate a metric that measures the long-term economic stability of the resulting internetwork. Some transit providers attract customers due to either lower prices or better performance, but are not profitable because their local and transit costs are higher than their revenues. Such an economic situation would not be sustainable in the long-term, as these providers would either go bankrupt or they would have to increase their prices. We measure the number of providers that are Unprofitable-but-Active ("num UA" in Table 2). We also measure the maximum traffic volume (as a fraction of the total traffic flow) carried by UA transit providers ("Traf UA"). First, note that the two metrics are positively correlated: a larger number of UA providers results in a larger traffic volume handled by UA providers. Second, we have more UA providers when LTPs peer with CPs (see, for example, scenario $\{$DF, (PR,NC), (SEL,CB)$\}$). In that scenario, traffic from CPs to ECs flows through a hierarchy of STPs. STPs at the top of the hierarchy can become UA providers, as they pay large transit fees to the generally more expensive LTPs. The largest number of UA providers results when STPs use SEL provider selection and peer using either NC or TR (up to 3 UA providers, and 50% of the total traffic carried by those providers). Then, STPs cannot peer with CPs and they also choose only LTPs as providers. All traffic from CPs to ECs flows through customer-provider links between STPs and LTPs, and this creates even more UA STPs.

5.4 Provider profitability when STPs use PR:

LTPs can harm STP profitability by peering with CPs:
When most edge networks choose providers based on price, cheaper STPs are able to attract a large fraction of the edge networks. Due to the overlapping prices of STPs and LTPs, however, LTPs can also attract some edge networks as customers. When STPs use PR provider selection, a hierarchy of STPs is formed. When LTPs use PR as well, they may be forced to connect to STP providers, and the peering clique of LTPs may no longer be sustainable. In such situations, both STPs and LTPs see approximately equal aggregate fitness. LTPs can, however, significantly harm the aggregate profits of STPs by using CB-peering, which allows them to peer with CPs. In this case some LTPs are able to form a large number of peering links with CPs. Consequently, these CPs reach most of their destinations through peering links with LTPs, followed by a hierarchy of STPs. In the default model, CPs source a large fraction of the traffic that goes to ECs. Consequently, the LTPs can significantly reduce the fitness of STPs by engaging in CB-peering. The conventional wisdom for LTPs is to only peer with other LTPs. This result shows that *CB-peering by LTPs can lead to a situation where LTPs significantly reduce STP profits, and are able to increase aggregate LTP fitness.*

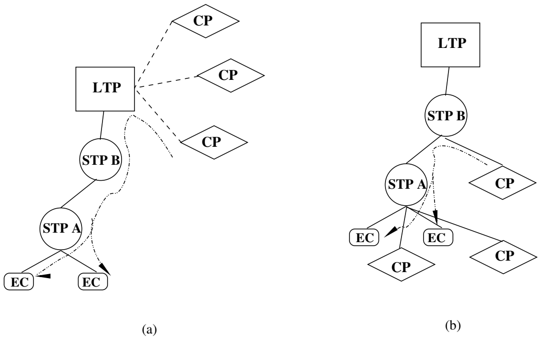

Figure 9: Peering between LTPs and CPs increases LTP profitability, but also increases weighted path lengths. The arrows indicate the paths followed by large traffic flows.

STPs should use TR-peering:
We find that the best peering strategy for STPs depends on the peering method used by LTPs. First consider the case where LTPs do not peer with CPs. In this scenario, we find that STPs achieve higher aggregate fitness by using TR-peering (though the total number of fit STPs is smaller). For example, the aggregate STP fitness in scenario $\{$DF, (PR,TR), (PR,TR)$\}$ is \$317k, while it is \$229k in scenario $\{$DF, (PR,CB), (PR,TR)$\}$. This indicates that the conventional wisdom of TR-peering results in higher STP fitness, *when STPs use price-based provider selection*. The reason for this is as follows. If STPs use CB-peering, then some CPs become their peers. On the other hand, if STPs use TR-peering, they cannot peer with CPs, as CPs always generate more traffic than they consume. These CPs would eventually become customers of STPs, as most edge networks choose providers based on price. This increases the fitness of STPs. On the other hand, when STPs use CB-peering, they can peer with CPs. In this case the traffic flow is of the form CP-STP-EC; less traffic flows on the customer-provider links in the hierarchy of STPs, leading to lower aggregate fitness for STPs.

Next, consider the case where LTPs use CB-peering. In this case, STPs are more profitable by using CB-peering than with TR-peering. If CPs peer with LTPs, then they do not need to choose STPs as providers. Given that these CPs will not become their customers, STPs can improve their profitability by peering with them. This can happen only if STPs use CB-peering.

5.5 Provider profitability when STPs use SEL

STP fitness is determined by LTP prices:
In this scenario, STPs do not choose other STPs as providers, because they consider them as potential peers or competitors. STPs still attract the price-conscious ECs and CPs. All STPs connect directly to LTPs due to SEL provider selection. This results in higher fitness for LTPs than the scenarios where STPs use PR provider selection. In case STPs peer only by necessity, the aggregate STP fitness can be negative, and there are no fit STPs. As these STPs carry traffic to/from their customers, we see a larger number of UA providers, and a larger fraction of traffic flowing through such providers. The aggregate fitness of STPs depends on the relative prices of STPs and LTPs. In our simulation setting, LTP prices are slightly higher than those of STPs, leading to a situation where STPs pay more in transit prices than they can recover from their customers. If LTP and STP prices are comparable, the aggregate STP profit can still be positive. The key point is that *if STPs use SEL provider selection, LTPs are in a position to use their market power to charge higher prices, and potentially make STPs unprofitable.*

STPs should use CB-peering:
When STPs use SEL provider selection, they achieve higher profits using CB-peering than TR-peering (*e.g.*, the aggregate STP profit is \$65k in scenario $\{$DF, (SEL,CB), (PR,NC)$\}$, while it is \$-48k in $\{$DF, (SEL,TR), (PR,NC)$\}$). This is in contrast to the case where STPs use PR provider selection, where they are better off using TR-peering. The reason for this is as follows. When STPs use SEL provider selection, they only connect to LTPs. Due to the higher prices of LTPs, it is beneficial for STPs to send as little traffic as possible to their upstream providers. If an STP S peers with a CP C, S only carries traffic destined from C to ECs in the customer tree of S. S does not send any of this traffic to its providers, and this traffic is profit-generating. This allows the STP to remain profitable even if LTPs charge high transit prices. A further benefit of CB-peering is that it allows "content-heavy" and "access-heavy" STPs to peer. Content-heavy STPs have many CPs as customers, while access-heavy STPs have many ECs as customers. These two types of STPs can peer only with CB-peering (traffic ratios will always be unbalanced), and results in increased fitness for STPs. This makes the case that *content and access heavy STPs should peer with each other to be profitable.*

6. DEVIATION-1: P2P TRAFFIC MATRIX

In the default model, the interdomain traffic matrix consists mostly of CS traffic (80%). In this section, we consider a deviation where the traffic matrix consists mostly (80%) of P2P traffic. Edge networks still choose their providers using PR, as in the default model. We call this deviation "P2P". The tables with the detailed results for P2P and subsequent deviations are in the appendix.

Peer-to-peer traffic helps STPs:
In the default model, most edge networks choose providers based on price. In Section 5, we observed that LTPs can significantly diminish the aggregate profit of STPs by using CB-peering. When the traffic matrix consists mostly of P2P traffic, the traffic volume from CPs to ECs is relatively smaller. As a result, the benefit for LTPs from peering with CPs is lower. The aggregate fitness of STPs is \$187k with scenario $\{$P2P, (PR,NC), (PR,CB)$\}$, while it is \$40k for the scenario $\{$DF, (PR,NC), (PR,CB)$\}$. *A traffic matrix that consists of mostly P2P traffic thus benefits STPs.*

Smaller increase in weighted path lengths when LTPs peer with CPs:
In the default model, we observed that when LTPs peer with CPs, weighted path lengths are longer than unweighted path lengths. For example, the weighted path length is 4.2 in scenario $\{$DF, (PR,NC), (PR,CB)$\}$, while the unweighted path length is 4.0. For scenario $\{$P2P, (PR,NC), (PR,CB)$\}$ the weighted path length is 4.1, while the unweighted path length is 4.0, *i.e.*, we observe a similar phenomenon with the P2P traffic matrix, though the difference between the weighted and unweighted path lengths is smaller. When the traffic matrix is predominantly P2P, the volume of traffic flowing from CPs to ECs (over the long paths caused when LTPs peer with CPs) is smaller than in the default model.

TR-peering is more profitable for STPs:
In the default model, when STPs use SEL provider selection, we found that either NC or TR-peering led to negative aggregate STP fitness (*e.g.*, aggregate STP profit is \$-104k in $\{$DF, (SEL,TR), (PR,CB)$\}$). This is because the traffic matrix has mostly CS traffic, and only a small fraction of the traffic flows between ECs (which are customers of STPs). Consequently, peering by STPs does not give significant benefit. With the P2P traffic matrix, however, a larger fraction of the end-to-end traffic flows between ECs. STPs can save significant transit fees if they use TR-peering (aggregate STP profit is \$58k in $\{$P2P, (SEL,TR), (PR,CB)$\}$). The likelihood of STPs being able to peer using TR-peering depends also on the peering strategy of LTPs; in particular, whether LTPs peer with CPs. We illustrate this with a specific example in Figure 10. In subfigure (a), the LTP peers with CPs. The traffic between STP A and STP B is now balanced, allowing them to peer using TR-peering. Subfigure (b) shows the case where LTPs do not peer with CPs. These CPs become customers of STPs, which are cheaper than LTPs. This can lead to the emergence of "content-heavy" STPs (STPs with content customers) and "access-heavy" STPs (STPs with access customers). Content and access heavy STPs cannot peer with each other using TR-peering, as the traffic is always imbalanced (more traffic from CPs to ECs). Thus, *if the traffic matrix consists of mostly P2P traffic, then STPs can save*

significant transit costs with TR-peering. Further, *peering between LTPs and CPs favors STPs, as it results in more balanced traffic between STPs, giving them more opportunities to peer.*

Traffic flow over UA providers:
In the P2P model, the traffic flow through UA providers is reduced, particularly when STPs use PR provider selection. In the default model, if LTPs peer with CPs, a number of STPs become unprofitable. As stated earlier, P2P traffic helps STPs, and the ability of LTPs to decrease aggregate STP fitness is reduced. This leads to a smaller number of STPs that are "unprofitable but active". In particular, for scenario {DF, (PR,NC), (PR,CB)}, 20% of the end-to-end traffic flows over UA providers, while for scenario {P2P, (PR,NC), (PR,CB)}, this value is around 4%.

An exception to the above result is when STPs use (SEL,NC). In this case, STPs connect directly to LTPs, and do not peer with other STPs. With P2P traffic, a large amount of traffic flows from ECs to other ECs. When STPs use SEL provider selection, this traffic traverses the customer-provider links from STPs to LTPs. This results in a larger number of UA providers and a larger fraction of traffic handled by those UA providers; 30% of the total traffic flows over UA providers in {P2P, (SEL,NC), (PR,NC)}, while that number is close to 0 for {DF, (SEL,NC), (PR,NC)}.

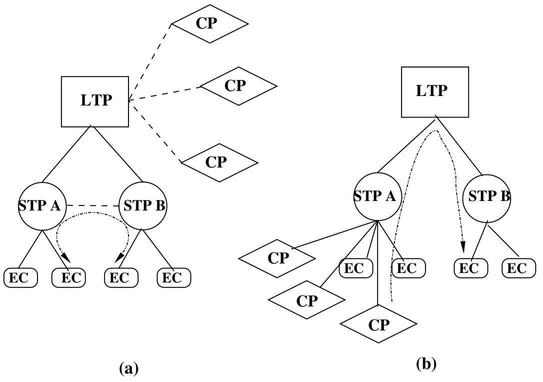

Figure 10: Peering between STPs more likely with P2P traffic and especially when LTPs peer with CPs. The arrows indicate the paths followed by large traffic flows.

7. DEVIATION-2: PF PROVIDER SELECTION BY EDGE NETWORKS

In the default model, 80% of ECs and CPs use PR provider selection. In this section, we consider a deviation where 80% of edge networks choose their providers using the PF method described in Section 2.6. We call this deviation "EP". The interdomain traffic matrix still consists of mostly (80%) CS traffic, as in the default model.

PF provider selection favors LTPs:
ECs and CPs that use PF provider selection are attracted to LTPs, and eventually, most ECs and CPs connect directly to LTPs. This is because LTPs can reach all destinations using links to their customers or peers, and so they provide the shortest paths. STPs can only attract the few ECs and CPs that use PR-provider selection. When STPs use SEL provider selection, there are no fit STPs (aggregate STP fitness is negative). Further, when STPs use SEL, no form of peering leads to positive aggregate fitness. In the default model, when STPs use SEL, they can be profitable by CB-peering. In the default model, STPs have a large customer base of ECs and CPs, and peering with CPs (or "content-heavy" STPs) can save significant transit expenses. In the EP model, however, STPs have a significantly smaller customer base. Consequently, peering does not increase the aggregate fitness of STPs.

Shorter paths:
When ECs and CPs use PF provider selection, the average path lengths in the network decrease. This is because networks at the edge are attracted to LTPs. End-to-end paths are of the form EC-LTP-EC or EC-LTP-LTP-EC, with no intermediate STPs. This results in shorter end-to-end paths. The unweighted path length is 3.3 for scenario {EP, (SEL,NC), (PR,NC)}, as opposed to 3.9 for scenario {DF, (SEL,NC), (PR,NC)}. *This implies that the EP model leads to a situation that is beneficial for the performance seen by edge networks, at the expense of STP profitability.*

Weighted paths shorter than unweighted paths:
In scenarios where STPs use (PR,NC), the weighted path lengths are *smaller* than the corresponding unweighted path lengths. For example, the weighted path length for scenario {EP, (PR,NC), (PR,NC)} is 3.3, while the unweighted path length is 3.5. This can be explained as follows: Networks at the edge (ECs and CPs) choose providers based on performance, and are thus attracted to LTPs. STPs, on the other hand, use PR provider selection and connect to other STPs. This creates a hierarchy of STPs, but with *a very small customer base connected to STPs*. Consequently, we see paths that traverse the hierarchy of STPs and LTPs, but only a small fraction of the total traffic flows on those paths. Most of the traffic flows on the short paths of the form EC-LTP-EC. As a result, *the weighted path length is smaller than the corresponding unweighted path length.*

More traffic over UA providers:
In the default model, significant traffic is carried by UA providers when STPs use PR provider selection and LTPs peer with CPs, *e.g.*, 20% of the total end-to-end traffic is carried by UA providers in scenario {DF, (PR,TR), (PR,CB)}. In other scenarios of the default model, negligible traffic is carried by UA providers. In the EP model, however, *10-20% of the total traffic is carried by UA providers in each scenario*. The reason for this is as follows. When edge networks use PF-provider selection, only a few ECs and CPs connect to STPs. Most of the traffic sourced/consumed by these customers of STPs comes from CPs and ECs that are connected to LTPs. Consequently, STPs send/receive large traffic volumes to their LTP providers, which leads to a larger number

of STPs that are unprofitable.

8. DEVIATION-3: CPS REPLICATE THEIR CONTENT IN EVERY REGION

In the default model, each CP is present in a single geographical region. A recent trend in the Internet is that CPs increasingly expand their geographical presence [17], either through the use of content distribution networks (CDNs), or by replicating their content at multiple locations. We present a deviation of the default model where CPs are present *in every geographical region*. Geographical expansion by CPs allows them to peer with networks in a larger number of regions, and also them to select providers from a larger number of regions. We call this deviation "GEO".

Larger STP profits:
In GEO, STPs obtain larger as compared to the default model. In GEO, most CPs use PR provider selection, but they are not restricted to choosing the cheapest provider from a single region. Instead, they can choose the cheapest (which are typically STPs) across all regions. This results in larger aggregate profits for STPs. For example, the aggregate STP fitness is \$331k with scenario $\{$DF, (PR,NC), (PR,NC)$\}$, while it is \$450k in scenario $\{$GEO, (PR,NC), (PR,NC$\}$.

STPs can be profitable even with SEL provider selection:
In the default model, we observed that if STPs use SEL, then the aggregate STP profits depend on the relative prices of STPs and LTPs. In the default model, STPs are unprofitable with SEL provider selection, unless they use CB-peering. In the GEO model, however, we find that STPs can be profitable even if they use SEL provider selection and NC or TR-peering. This is because most CPs use PR provider selection, and can select the cheapest STPs from all regions. This increases the aggregate profits of STPs.

Larger aggregate profit for STPs by TR-peering:
In the default model, when STPs use PR provider selection, their aggregate profit is larger with TR-peering than with CB-peering. This is because by using CB-peering, STPs peer with CPs, which would otherwise become their customers. We find that this effect is more pronounced when CPs are present in every region. The difference between scenarios $\{$DF, (PR,TR), (PR,TR)$\}$ and $\{$DF, (PR,CB), (PR,TR)$\}$ is \$88k (38%), while the difference between $\{$GEO, (PR,TR), (PR,TR)$\}$ and $\{$GEO, (PR,CB), (PR,TR)$\}$ is \$151k (45%).

Shorter paths:
As in the EP model, we find that several scenarios in GEO lead to weighted paths that are shorter than unweighted paths. For example, the weighted path length for scenario $\{$GEO, (PR,NC), (PR,NC)$\}$ is 3.7, while the unweighted path length is 3.9. In GEO, the traffic matrix has mostly CS traffic, with large traffic volumes from from CPs to ECs. Further, CPs can connect to STPs in multiple regions, and the number of CP-STP links is larger than in the default model. Consequently, we see a large number of "short" paths of the form CP-STP-EC, which bypass LTPs and also carry significant traffic. This leads to weighted paths that are shorter than unweighted paths.

9. 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.* [15], 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, starting with the preferential attachment model of Barabasi *et al.* [6]. Several variants and comparisons of preferential attachment models were later proposed [3, 7, 26, 28, 29, 30, 31] The models in this research thread have been mostly descriptive, meaning that they attempt to reproduce certain known structural characteristics of the Internet.

The previous descriptive models received considerable criticism (for instance, see [19, 20]) 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, those models do not explain how the Internet topology is evolving. This led to models that view the Internet topology as the effect of optimization-driven activity by individual ASes. These concepts were first introduced by Carlson and Doyle in [8], and later applied in the context of the Internet in [14]. Chang *et al.* [9] model AS interconnection practices, considering the effects of AS geography, AS business models and AS evolution.

The body of work closest in spirit to ours is that of Chang et al. [10]. That work 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. In this work, we focus mainly on studying the properties of the equilibrium that results as each AS uses certain provider and peer selection strategies. Also related is the work of Holme *et al.* [18], which developed an agent-based simulation model where the agents are individual ASes with economic incentives. 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, and their model is rather simplistic, ignoring some important domain-specific details about the Internet at the interdomain level. Corbo et al. [12] propose an economically-principled model that is able to create the observed structure of the AS graph. 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 [22, 23, 24] advocate the use of the Shapley value for revenue distribution between ISPs. They show that if profits are shared according to the Shapley value, the set of "fair" properties inherent to the Shapley solution exist, and the selfish behavior of ISPs leads to globally optimal routing and interconnection decisions. A body of work known as "network formation games" [4, 5, 27] takes a *game theoretic* approach to the creation of interdomain links between autonomous networks. These papers formulate a game where 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 of each AS is 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 an AS knows the payoff obtained from creating a particular link. We consider the realistic case where ASes do not play simultaneously, and are able to observe the moves made by previous players. Also, we assume that an AS cannot predict the payoffs it would obtain by choosing certain providers or peers.

10. CONCLUSIONS

In this paper, we proposed ITER, a detailed model of interdomain network formation that captures the interdependence between interdomain topology, traffic flow and provider and peer selection strategies of ANs. We present an approach to solve this model using agent-based simulations. As a first practical application of this model, we evaluate the effect of various strategies for provider selection ("choose cheapest providers", "choose higher-tier providers") and peer selection ("peer by necessity", "peer by traffic ratios" and "peer by cost-benefit analysis") on the profitability of small and large transit providers. We examine the effects of these strategies on the economics, topology and performance of the internetwork at equilibrium. There are several other applications of ITER, which we plan to pursue in future work:

1. In future work, we will use ITER to find the *optimal strategy* for each class of ANs, *i.e.*, the strategy that maximizes their profitability given the expected strategy of other ANs.

2. In this paper, we focused on ITER scenarios where each class of ANs uses the same strategies for provider/peer selection. In future work, we plan to examine scenarios where each AN can use a different strategy.

3. An important aspect of ITER is that the selection of providers is partly based on their prices. To study a basic framework for pricing decisions, we propose to extend ITER to the following extended model. Before the network formation stage, each transit provider decides its price. In the second stage, we solve ITER for this setting of prices. Using the equilibria resulting from the second stage, we can determine the equilibrium pricing strategies (first stage) so that each provider maximizes its own profit. We propose to extend ITER to include strategic pricing, allowing us to observe the dynamics of the joint process between pricing and network formation.

4. Each AN in the Internet tries to optimize its own utility function in a distributed manner. An important open question is "What problem does the Internet, at the global level, try to solve through such distributed optimizations?" In future work, we will attempt to answer this question by comparing the equilibrium that results from ITER with the network that results from centralized optimization of a particular metric.

11. REFERENCES

[1] Peering Database. `http://www.peeringdb.com`.

[2] A. Akella, S. Seshan, and A. Shaikh. An Empirical Evaluation of Wide-Area Internet Bottlenecks. In *Proc. ACM SIGCOMM IMC*, 2003.

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

[4] E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler. Near-optimal network design with selfish agents. In *Proc. the annual ACM symposium on Theory of computing (STOC)*, 2003.

[5] E. Anshelevich, B. Shepherd, and G. Wilfong. Strategic network formation through peering and service agreements. In *Proc. 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2006.

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

[7] T. Bu and D. Towsley. On Distinguishing Between Internet Power Law Topology Generators. In *Proc. IEEE Infocom*, 2002.

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

[9] H. Chang, S. Jamin, and W. Willinger. Internet Connectivity at the AS-level: An Optimization-Driven Modeling Approach. In *Proc. ACM SIGCOMM Workshop on MoMeTools*, 2003.

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

[11] H. Chang and W. Willinger. Difficulties Measuring the Internet's AS-Level Ecosystem. In *Proc. the 40th Annual Conference on Information Sciences and Systems*, 2006.

[12] J. Corbo, S. Jain, M. Mitzenmacher, and D. Parkes. An Economically Principled Generative Model of AS Graph Connectivity. In *Proc. International Joint Workshop on The Economics of Network Systems and Incentive-Based Computing*, 2007.

[13] A. Dhamdhere and C. Dovrolis. Ten Years in the Evolution of the Internet Ecosystem. In *Proc. ACM SIGCOMM IMC*, 2008.

[14] A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou. Heuristically Optimized Trade-Offs: A New Paradigm for Power Laws in the Internet. In

Proc. ICALP, 2002.

[15] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On Power-law Relationships of the Internet Topology. In *Proc. ACM SIGCOMM*, 1999.

[16] L. Gao and F. Wang. The Extent of AS Path Inflation by Routing Policies. In *Proc. IEEE GLOBECOM*, 2002.

[17] P. Gill, M. Arlit, Z. Li, and A. Mahanti. The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? In *Proc. Passive and Active Measurement Conference (PAM)*, 2008.

[18] P. Holme, J. Karlin, and S. Forrest. An Integrated Model of Traffic, Geography and Economy in the Internet. *SIGCOMM CCR*, 38(3), 2008.

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

[20] D. Krioukov, kc claffy, M. Fomenkov, F. Chung, A. Vespignani, and W. Willinger. The Workshop on Internet Topology (wit) Report. *ACM SIGCOMM CCR*, 37(1), 2007.

[21] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph Evolution: Densification and Shrinking Diameters. *ACM Transactions on Knowledge Discovery from Data (ACM TKDD)*, 2007.

[22] 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 *Proc. ACM SIGCOMM 2008 NetEcon Workshop*, 2008.

[23] R. T. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein. On Cooperative Settlement Between Content, Transit and Eyeball Internet Service Providers. In *Proc. ACM CoNEXT*, 2008.

[24] 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 *Proc. ACM CoNEXT*, 2007.

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

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

[27] Ramesh Johari and Shie Mannor and John Tsitsiklis. A contract-based model for directed network formation. *Games and Economic Behavior*, 56(2), August 2006.

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

[29] X. Wang and D. Loguinov. Wealth-Based Evolution Model for the Internet AS-Level Topology. In *Proc. IEEE Infocom*, 2006.

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

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

APPENDIX

Detailed results for deviations P2P, EP and GEO

sc	STP str	LTP str	wPL	uPL	diam-eter	profit all STP ($k)	profit all LTP ($k)	profit fit STP ($k)	profit fit LTP ($k)	num fit STP	num fit LTP	num UA	Traf UA	%PP	Traf PP
s1	PR,NC	PR,NC	4.0	3.9	6.2	266	419	390	540	3.8	1.5	2.0	0.0	2.4	0.1
s2	PR,NC	PR,TR	4.0	3.9	6.2	266	419	390	540	3.8	1.5	2.0	0.0	2.4	0.1
s3	PR,NC	PR,CB	4.1	4.0	6.5	187	405	312	512	4.0	1.6	2.0	0.0	6.1	0.2
s4	PR,NC	SEL,NC	4.0	3.9	6.1	295	388	415	515	4.0	1.4	1.9	0.0	2.9	0.2
s5	PR,NC	SEL,TR	4.0	3.9	6.1	295	388	415	515	4.0	1.4	1.9	0.0	2.9	0.2
s6	PR,NC	SEL,CB	4.1	4.0	6.5	187	405	312	512	4.0	1.6	2.0	0.0	6.1	0.2
s7	PR,TR	PR,NC	4.0	3.9	6.0	295	381	416	508	4.1	1.4	1.9	0.0	2.7	0.2
s8	PR,TR	PR,TR	4.0	3.9	6.1	384	284	506	407	4.0	1.4	2.2	0.0	2.1	0.2
s9	PR,TR	PR,CB	4.1	4.0	6.3	206	404	335	515	3.6	1.5	1.7	0.1	5.2	0.3
s10	PR,TR	SEL,NC	4.0	3.9	5.9	303	373	421	502	4.1	1.4	1.7	0.0	3.2	0.2
s11	PR,TR	SEL,TR	4.0	3.9	6.2	387	273	506	403	4.2	1.2	2.0	0.0	3.2	0.3
s12	PR,TR	SEL,CB	4.1	3.9	6.1	219	395	347	505	3.6	1.5	2.2	0.0	6.1	0.3
s13	PR,CB	PR,NC	4.0	3.9	6.1	379	293	500	421	4.1	1.3	2.3	0.0	2.6	0.2
s14	PR,CB	PR,TR	4.0	3.9	5.9	331	365	455	482	3.8	1.4	2.3	0.0	2.3	0.2
s15	PR,CB	PR,CB	4.0	3.9	6.2	185	433	320	503	3.0	2.2	1.1	0.1	6.7	0.4
s16	PR,CB	SEL,NC	4.0	3.9	6.2	353	311	472	440	4.4	1.3	1.9	0.0	3.7	0.2
s17	PR,CB	SEL,TR	4.0	3.9	5.8	320	356	439	479	4.1	1.3	1.9	0.0	3.4	0.3
s18	PR,CB	SEL,CB	4.0	3.9	6.0	185	436	316	509	3.1	2.2	1.1	0.0	7.5	0.5
s19	SEL,NC	PR,NC	4.1	3.9	5.0	-85	863	59	963	2.0	2.0	2.0	0.3	2.1	0.0
s20	SEL,NC	PR,TR	4.1	3.9	5.0	-85	863	59	963	2.0	2.0	2.0	0.3	2.1	0.0
s21	SEL,NC	PR,CB	4.1	3.9	5.0	-127	835	31	945	2.0	1.6	3.0	0.3	5.0	0.1
s22	SEL,NC	SEL,NC	4.1	3.9	5.0	-85	863	59	963	2.0	2.0	2.0	0.3	2.1	0.0
s23	SEL,NC	SEL,TR	4.1	3.9	5.0	-85	863	59	963	2.0	2.0	2.0	0.3	2.1	0.0
s24	SEL,NC	SEL,CB	4.1	3.9	5.0	-127	835	31	945	2.0	1.6	3.0	0.3	5.0	0.1
s25	SEL,TR	PR,NC	4.0	3.9	5.0	-11	784	125	884	3.0	2.0	1.0	0.0	2.5	0.1
s26	SEL,TR	PR,TR	4.0	3.9	5.0	29	742	163	845	2.9	1.9	1.1	0.0	2.3	0.1
s27	SEL,TR	PR,CB	4.0	3.9	5.0	58	623	199	711	2.8	2.0	1.6	0.0	5.3	0.3
s28	SEL,TR	SEL,NC	4.0	3.9	5.0	-11	784	125	884	3.0	2.0	1.0	0.0	2.5	0.1
s29	SEL,TR	SEL,TR	4.0	3.9	5.0	20	751	155	853	3.0	1.9	1.0	0.0	2.6	0.1
s30	SEL,TR	SEL,CB	4.0	3.9	5.0	58	621	192	712	3.0	2.0	1.7	0.0	6.0	0.4
s31	SEL,CB	PR,NC	4.0	3.9	5.0	1	770	126	870	3.0	2.0	2.0	0.0	3.6	0.1
s32	SEL,CB	PR,TR	4.0	3.9	5.1	66	705	196	809	2.8	1.8	2.3	0.1	2.8	0.2
s33	SEL,CB	PR,CB	4.0	3.9	5.1	92	578	228	652	2.6	2.1	1.6	0.1	5.7	0.4
s34	SEL,CB	SEL,NC	4.0	3.9	5.0	1	770	126	870	3.0	2.0	2.0	0.0	3.6	0.1
s35	SEL,CB	SEL,TR	4.0	3.9	5.0	56	717	183	821	2.8	1.9	2.4	0.1	3.6	0.2
s36	SEL,CB	SEL,CB	4.0	3.9	5.0	91	588	227	666	2.6	1.8	2.0	0.1	6.5	0.4

Table 3: Deviation-1 (P2P)

sc	STP str	LTP str	wPL	uPL	diam- eter	profit all STP ($k)	profit all LTP ($k)	profit fit STP ($k)	profit fit LTP ($k)	num fit STP	num fit LTP	num UA	Traf UA	%PP	Traf PP
s1	PR,NC	PR,NC	3.3	3.5	6.0	-154	1537	13	1656	0.8	1.6	4.2	0.1	2.7	0.0
s2	PR,NC	PR,TR	3.3	3.5	6.0	-154	1537	13	1656	0.8	1.6	4.2	0.1	2.7	0.0
s3	PR,NC	PR,CB	3.4	3.5	6.5	-204	1512	4	1620	0.4	1.6	5.3	0.2	5.2	0.1
s4	PR,NC	SEL,NC	3.3	3.5	6.0	-154	1537	13	1656	0.8	1.6	4.2	0.1	2.7	0.0
s5	PR,NC	SEL,TR	3.3	3.5	6.0	-154	1537	13	1656	0.8	1.6	4.2	0.1	2.7	0.0
s6	PR,NC	SEL,CB	3.4	3.5	6.5	-204	1512	4	1620	0.4	1.6	5.3	0.2	5.2	0.1
s7	PR,TR	PR,NC	3.3	3.4	5.8	-152	1498	8	1618	0.8	1.5	4.1	0.1	2.8	0.0
s8	PR,TR	PR,TR	3.3	3.4	6.0	-132	1455	24	1583	0.9	1.4	4.0	0.1	2.4	0.0
s9	PR,TR	PR,CB	3.3	3.4	6.2	-154	1469	28	1589	0.8	1.4	4.2	0.1	3.9	0.1
s10	PR,TR	SEL,NC	3.3	3.4	5.8	-152	1498	8	1618	0.8	1.5	4.1	0.1	2.8	0.0
s11	PR,TR	SEL,TR	3.3	3.4	5.9	-132	1456	24	1584	0.9	1.4	3.9	0.1	2.9	0.0
s12	PR,TR	SEL,CB	3.3	3.4	6.3	-154	1468	28	1587	0.9	1.5	4.0	0.1	4.6	0.1
s13	PR,CB	PR,NC	3.3	3.4	5.9	-165	1512	2	1633	0.5	1.4	4.6	0.1	4.3	0.0
s14	PR,CB	PR,TR	3.3	3.4	6.1	-149	1535	19	1653	0.5	1.6	4.2	0.1	3.6	0.1
s15	PR,CB	PR,CB	3.3	3.4	6.0	-175	1475	18	1586	0.4	1.6	4.0	0.2	5.0	0.1
s16	PR,CB	SEL,NC	3.3	3.4	5.9	-165	1512	2	1633	0.5	1.4	4.6	0.1	4.3	0.0
s17	PR,CB	SEL,TR	3.3	3.4	6.1	-143	1526	23	1646	0.7	1.5	4.4	0.1	4.2	0.1
s18	PR,CB	SEL,CB	3.3	3.4	6.1	-174	1474	18	1585	0.5	1.6	4.1	0.2	5.6	0.1
s19	SEL,NC	PR,NC	3.2	3.3	5.0	-183	1519	0	1620	0.0	2.0	3.0	0.2	2.1	0.0
s20	SEL,NC	PR,TR	3.2	3.3	5.0	-183	1519	0	1620	0.0	2.0	3.0	0.2	2.1	0.0
s21	SEL,NC	PR,CB	3.2	3.3	5.0	-193	1517	0	1617	0.0	2.0	4.0	0.2	2.9	0.0
s22	SEL,NC	SEL,NC	3.2	3.3	5.0	-183	1519	0	1620	0.0	2.0	3.0	0.2	2.1	0.0
s23	SEL,NC	SEL,TR	3.2	3.3	5.0	-183	1519	0	1620	0.0	2.0	3.0	0.2	2.1	0.0
s24	SEL,NC	SEL,CB	3.2	3.3	5.0	-193	1517	0	1617	0.0	2.0	4.0	0.2	2.9	0.0
s25	SEL,TR	PR,NC	3.2	3.3	5.0	-180	1514	0	1615	0.0	2.0	3.0	0.2	2.5	0.0
s26	SEL,TR	PR,TR	3.2	3.3	5.0	-178	1511	1	1611	0.1	2.0	2.9	0.2	2.4	0.0
s27	SEL,TR	PR,CB	3.2	3.3	5.0	-192	1510	0	1610	0.0	2.0	4.0	0.2	3.4	0.0
s28	SEL,TR	SEL,NC	3.2	3.3	5.0	-180	1514	0	1615	0.0	2.0	3.0	0.2	2.5	0.0
s29	SEL,TR	SEL,TR	3.2	3.3	5.0	-178	1512	1	1612	0.1	2.0	2.9	0.2	2.6	0.0
s30	SEL,TR	SEL,CB	3.2	3.3	5.0	-192	1510	0	1610	0.0	2.0	4.0	0.2	3.4	0.0
s31	SEL,CB	PR,NC	3.2	3.3	5.0	-192	1514	0	1615	0.0	2.0	3.0	0.2	3.8	0.0
s32	SEL,CB	PR,TR	3.2	3.3	5.0	-186	1500	3	1600	0.2	2.0	2.8	0.2	3.7	0.0
s33	SEL,CB	PR,CB	3.2	3.3	5.0	-198	1496	1	1593	0.1	2.0	3.6	0.2	4.4	0.0
s34	SEL,CB	SEL,NC	3.2	3.3	5.0	-192	1514	0	1615	0.0	2.0	3.0	0.2	3.8	0.0
s35	SEL,CB	SEL,TR	3.2	3.3	5.0	-180	1491	6	1599	0.3	1.9	2.7	0.2	3.9	0.0
s36	SEL,CB	SEL,CB	3.2	3.3	5.0	-197	1495	1	1593	0.1	2.0	3.7	0.2	4.5	0.0

Table 4: Deviation-2 (EP)

sc	STP str	LTP str	wPL	uPL	diameter	profit all STP ($k)	profit all LTP ($k)	profit fit STP ($k)	profit fit LTP ($k)	num fit STP	num fit LTP	num UA	Traf UA	%PP	Traf PP
s1	PR,NC	PR,NC	3.7	3.9	6.2	450	207	570	334	4.0	1.4	0.7	0.0	2.4	0.1
s2	PR,NC	PR,TR	3.7	3.9	6.2	450	207	570	334	4.0	1.4	0.7	0.0	2.4	0.1
s3	PR,NC	PR,CB	4.2	4.0	6.3	17	393	149	493	3.6	2.0	1.4	0.2	6.6	0.5
s4	PR,NC	SEL,NC	3.7	3.9	6.2	464	190	583	324	4.2	1.2	0.8	0.0	2.7	0.1
s5	PR,NC	SEL,TR	3.7	3.9	6.2	464	190	583	324	4.2	1.2	0.8	0.0	2.7	0.1
s6	PR,NC	SEL,CB	4.2	4.0	6.3	17	393	149	493	3.6	2.0	1.4	0.2	6.6	0.5
s7	PR,TR	PR,NC	3.7	3.9	6.4	512	132	630	264	4.3	1.3	0.8	0.0	2.4	0.1
s8	PR,TR	PR,TR	3.7	3.9	6.1	515	152	637	290	4.0	1.2	1.3	0.0	2.1	0.2
s9	PR,TR	PR,CB	4.0	3.9	6.2	92	406	236	499	2.7	2.0	1.9	0.2	6.4	0.4
s10	PR,TR	SEL,NC	3.7	3.9	6.4	528	113	645	252	4.5	1.1	1.0	0.0	2.9	0.2
s11	PR,TR	SEL,TR	3.7	3.9	6.1	527	117	645	259	4.2	1.0	1.4	0.0	3.0	0.2
s12	PR,TR	SEL,CB	4.1	3.9	6.0	53	387	184	489	3.2	2.0	1.3	0.2	6.8	0.5
s13	PR,CB	PR,NC	3.8	3.9	6.1	360	196	483	322	4.0	1.4	1.0	0.0	5.2	0.3
s14	PR,CB	PR,TR	3.8	3.9	6.1	354	229	484	356	3.7	1.3	1.0	0.1	5.2	0.3
s15	PR,CB	PR,CB	3.9	3.9	5.9	89	347	229	419	2.3	2.2	1.6	0.2	7.9	0.6
s16	PR,CB	SEL,NC	3.7	3.9	6.2	398	183	518	314	4.1	1.3	1.0	0.0	5.8	0.3
s17	PR,CB	SEL,TR	3.8	3.9	6.1	392	181	514	315	4.0	1.2	0.8	0.0	5.8	0.3
s18	PR,CB	SEL,CB	3.9	3.9	5.9	84	358	220	439	2.5	1.9	1.8	0.1	8.3	0.6
s19	SEL,NC	PR,NC	3.8	3.9	5.0	123	595	257	696	2.4	1.6	1.0	0.1	2.1	0.0
s20	SEL,NC	PR,TR	3.8	3.9	5.0	123	595	257	696	2.4	1.6	1.0	0.1	2.1	0.0
s21	SEL,NC	PR,CB	3.9	3.9	5.0	-152	662	13	761	1.3	1.8	2.2	0.4	6.0	0.5
s22	SEL,NC	SEL,NC	3.8	3.9	5.0	123	595	257	696	2.4	1.6	1.0	0.1	2.1	0.0
s23	SEL,NC	SEL,TR	3.8	3.9	5.0	123	595	257	696	2.4	1.6	1.0	0.1	2.1	0.0
s24	SEL,NC	SEL,CB	4.0	3.9	5.0	-165	693	2	793	1.0	2.0	2.0	0.5	5.8	0.4
s25	SEL,TR	PR,NC	3.8	3.9	5.0	123	595	257	696	2.4	1.6	1.0	0.1	2.1	0.0
s26	SEL,TR	PR,TR	3.8	3.9	5.0	139	588	273	688	2.6	1.4	1.1	0.1	2.1	0.0
s27	SEL,TR	PR,CB	3.9	3.9	5.0	-27	578	141	678	1.5	1.8	2.1	0.4	5.5	0.5
s28	SEL,TR	SEL,NC	3.8	3.9	5.0	123	595	257	696	2.4	1.6	1.0	0.1	2.1	0.0
s29	SEL,TR	SEL,TR	3.8	3.9	5.0	137	590	270	692	2.6	1.4	1.0	0.1	2.2	0.0
s30	SEL,TR	SEL,CB	3.9	3.9	5.0	-35	588	124	690	1.5	1.9	1.7	0.4	6.3	0.5
s31	SEL,CB	PR,NC	3.7	3.8	5.0	235	478	365	579	3.0	1.6	0.4	0.0	2.5	0.1
s32	SEL,CB	PR,TR	3.7	3.9	5.0	247	467	378	568	2.9	1.5	0.7	0.0	2.6	0.1
s33	SEL,CB	PR,CB	3.8	3.9	5.0	28	484	187	575	2.1	1.8	2.2	0.2	5.8	0.5
s34	SEL,CB	SEL,NC	3.7	3.8	5.0	235	478	365	579	3.0	1.6	0.4	0.0	2.5	0.1
s35	SEL,CB	SEL,TR	3.7	3.9	5.0	246	469	377	571	2.9	1.5	0.6	0.0	2.5	0.1
s36	SEL,CB	SEL,CB	3.8	3.9	5.0	11	498	163	586	2.2	1.8	1.9	0.2	6.8	0.5

Table 5: Deviation-3 (GEO)