GENESIS: An agent-based model of interdomain network formation, traffic flow and economics

Aemen Lodhi
School of Computer Science
Georgia Institute of Technology
aemen.lodhi@gatech.edu

Amogh Dhamdhere
CAIDA
University of California San Diego
amogh@caida.org

Constantine Dovrolis
School of Computer Science
Georgia Institute of Technology
constantine@gatech.edu

***Abstract*—We propose an agent-based network formation model for the Internet at the Autonomous System (AS) level. The proposed model, called *GENESIS*, is based on realistic provider and peering strategies, with ASes acting in a myopic and decentralized manner to optimize a cost-related fitness function. GENESIS captures key factors that affect the network formation dynamics: highly skewed traffic matrix, policy-based routing, geographic co-location constraints, and the costs of transit/peering agreements. As opposed to analytical game-theoretic models, which focus on proving the existence of equilibria, GENESIS is a computational model that simulates the network formation process and allows us to actually compute distinct equilibria (i.e., networks) and to also examine the behavior of sample paths that do not converge. We find that such oscillatory sample paths occur in about 10% of the runs, and they always involve tier-1 ASes, resembling the tier-1 peering disputes often seen in practice. GENESIS results in many distinct equilibria that are highly sensitive to initial conditions and the order in which ASes (agents) act. This implies that we cannot predict the properties of an individual AS in the Internet. However, certain properties of the global network or of certain classes of ASes are predictable. We also examine whether the underlying game is zero-sum, and identify three sufficient conditions for that property. Finally, we apply GENESIS in a specific "what-if" question, asking how the openness towards peering affects the resulting network in terms of topology, traffic flow and economics. Interestingly, we find that the peering openness that maximizes the fitness of different network classes (tier-1, tier-2 and tier-3 providers) closely matches that seen in real-world peering policies.**

I. Introduction

Tens of thousands of Autonomous Systems (ASes) interconnect in a complex and dynamic manner to form the Internet. These ASes belong to different categories e.g. enterprise networks, content sources, access providers, transit providers, or various combinations of the aforementioned categories. Additionally, they differ in geographic size ("expanse") and economic parameters (e.g., transit prices). ASes connect with one another mostly through two types of relations: customer-provider (or "transit") links and settlement-free peering links. Most interactions between ASes are local (unilateral in the case of provider selection and bilateral in the case of peering), decentralized and dynamic. These local connectivity decisions, however, can have a global impact on the economic viability of all ASes and the structure of the Internet.

This research was financially supported by Google and the National Science Foundation (grant CNS-1017139). Its contents are solely the responsibility of the authors and do not necessarily represent the official views of Google or NSF.

The Internet remains in a persistent state of flux subject to changes in various exogenous factors. How will the Internet change due to consolidation of content [1], large penetration of video streaming, falling transit prices [2], expanding geographic footprint of content providers [3], cheap local availability of peering infrastructure at IXPs [4]? We propose a computational agent-based network formation model, called "GENESIS", as a tool to study such questions. GENESIS is modular and easily extensible, allowing researchers to experiment with different peering or provider selection strategies, traffic matrix characteristics, routing policies, cost parameters, etc.

In this paper, we first introduce the main features of GENESIS and then use simple yet realistic provider selection and peering strategies, as proof-of-concept, to study the properties of the resulting networks. Starting from a population of ASes and initial topology, the model executes an iterative process in which each AS acts asynchronously to optimize its set of provider and peering links. We focus on the dynamic behavior of the model, identify the cause of some oscillatory sample paths (about 10% of the runs), and analyze the variability of the resulting equilibria. We show that under certain conditions about transit and peering costs, the set of transit providers engage in an (approximately) zero-sum game. We finally apply GENESIS in a specific "what-if" question, examining the impact of a parameter, related to the openness of transit providers towards peering, on the topology, traffic flow, and economics of the resulting network.

II. Related Work

Most of the previous related work has focused on characterizing and modeling the AS-level Internet topology. Various graph theoretic models aim to reproduce observed Internet topological properties (e.g., power-law degree distribution [5]) [6], [7], [8]. Another class of models take a bottom-up approach [9], modeling the optimization objectives and constraints of individual ASes, to create graphs that have the same topological properties [9], [10], [11], [12]. GENESIS is different than that line of research because it does not aim to be a topology generation model or tool; instead, it captures interdomain traffic flow, geographic constraints and economics to model the network formation process.

A large body of work on game-theoretic network formation models exists in the computer science and economics literature. We refer the interested reader to two recent books [13], [14]. Those models capture the objectives and potential strategies of each node, as players in a non-cooperative game, and they focus on proving the existence of (typically many) equilibria. The need for mathematical tractability, however, requires significant simplifications (such as lack of geographic constraints, simple cost functions, or uniform traffic flow between nodes. Consequently, the resulting networks are typically simple graphs, such as rings, trees or other regular structures. We avoid such simplifications in GENESIS. Additionally, few approaches in that line of work have focused on the dynamics of the network formation process and on the selection between multiple possible equilibria.

Our work is most related to agent-based computational models that simulate the dynamics of the network formation process, capturing the asynchronous and decentralized process through which nodes adjust their connectivity. The model by Holme et al. [15] is similar to GENESIS but it does not include peering and realistic routing. The model of Chang et al. [16] uses hard-wired strategies for provider and peer selection, among other differences. The model of Dhamdhere et al. [17] is more similar to GENESIS but it assigns a specific function to each node (e.g., content provider) and it focuses on the differences between two Internet instances (hierarchical versus "flat").

III. Model Description

We consider a population of N nodes (representing Autonomous Systems) which interconnect through two types of links: customer-provider and peering. Each node has the following attributes: a set of locations in which it is present, an amount of traffic it sends to and receives from every other node, and certain economic parameters, such as the transit price it would charge to its potential customers at a given location. We do not assign an a priori business function to nodes; a node may end up acting as tier-1 transit provider in one equilibrium and as a content provider in another. We next describe each component of GENESIS in more detail. A complete description of the model, with a longer justification, is available at a technical report [18]. The source code for the simulator that executes GENESIS is available at the same URL.

Geographical presence: There are G_M locations, and a node x is present at a set $G(x)$ of locations. These locations are roughly analogous to Internet Exchange Points (IXPs). Two nodes *overlap* if they share at least one location. For node x, $O(x)$ denotes the set of networks that overlap with x. A link between two nodes can be formed if they overlap.

Traffic matrix and transit traffic: The element T_{xy} of the N-by-N interdomain traffic matrix T denotes the average traffic volume generated by node x and consumed by node y. Overall, x generates traffic $V_G(x)$ and consumes $V_C(x)$. T_{xx}=0, i.e., we do not capture the local traffic within a network. The *transit traffic* $V_T(x)$ of x is the traffic volume that is neither generated nor consumed by x – it only passes through x enroute to its destination. The transit traffic of a node depends on the underlying network topology and the routing algorithm. Even if the interdomain traffic matrix T is constant, the transit traffic of a node may change as the underlying topology changes. The total traffic volume of a node $V(x)$ is given by the following expression:

$$V(x) = V_C(x) + V_G(x) + V_T(x) \tag{1}$$

Economic attributes: The economic attributes of a node include its transit prices (one for each location it is present at).

Transit cost: Let x be a transit customer of y, and let $P_y(x)$ be the *lowest* transit price of y across all regions in which x and y overlap. If $V_P(x)$ is the traffic exchanged between x and y, then the transit payment from x to y is:

$$TC(x) = P_y(x) \times V_P(x)^{\tau} \tag{2}$$

where τ is a transit traffic exponent that captures the *economies of scale* observed in practice. The transit revenue $TR(y)$ of y is simply the sum of transit costs paid by all customers of y.

Peering cost: Nodes engaging in settlement-free peering relations share the underlying peering costs. There are two primary mechanisms for peering — private and public — that have different cost structures. If the traffic exchanged between two peers is less than a threshold Ψ they peer publicly at an IXP, otherwise they peer privately. In public peering, nodes exchange traffic over a common switching fabric at an IXP, aggregating traffic from different peering sessions through the same port, whereas in private peering they set up a direct link with each peer.

Private peering cost: Let $V_{PP}(x, y)$ be the traffic exchanged between node x and its peer y over a private peering link. The total cost of private peering for x is given by:

$$PC_{prv}(x) = \alpha \times \sum_{y} V_{PP}(x, y)^{\beta} \tag{3}$$

where α is the peering cost per Mbps and β is the peering traffic exponent that accounts for the corresponding economies of scale.

Public peering cost: Let $V_{PP}(x, z)$ be the traffic exchanged between node x and its peer z over the public peering infrastructure. As x aggregates all its public peering traffic over the same port, the corresponding cost is:

$$PC_{pub}(x) = \alpha \times (\sum_{z} V_{PP}(x, z))^{\beta} \tag{4}$$

Fitness: The fitness of a node represents its net profit,

$$F(x) = TR(x) - TC(x) - PC_{pub}(x) - PC_{prv}(x) \tag{5}$$

If a node is a stub, i.e., a node without any customers, the first term is zero and the node's fitness will be negative.

Peering: Two nodes x and y are potential peers if they satisfy two *peering criteria* — x and y overlap geographically, and they do not have an existing customer-provider relationship. Additionally, a node x uses a *peering strategy* $S(x)$ to determine which of its potential peers it wants to peer with. Unlike provider selection, where a customer unilaterally chooses its provider, peering is a bilateral decision process. Thus, two potential peers x and y can peer iff the constraints of both $S(x)$ and $S(y)$ are satisfied. Depeering, however, is a unilateral decision by one of the peers. In this paper, we consider three peering strategies described in Section IV. We emphasize that GENESIS in not limited to those particular strategies; we have implemented several other strategies, including paid-peering, cost-benefit analysis, etc.

Provider selection: A node must have a transit provider if it cannot reach all other nodes in the network via its peers and customers. Node x selects a provider y if: *(a)* x overlaps with y, *(b)* y is "larger" than x (explained next), *(c)* y is not a peer of x, *(d)* y is not a customer of an existing peer of x, and *(e)* y is the least expensive among all nodes that satisfy the previous constraints. We say that a node y is larger than a node x if y is present in at least as many locations as x, and it carries more transit traffic than x. If a node x cannot find a provider that fulfills the previous criteria, then x becomes a $tier-1$ (T1) AS. In order to ensure a connected network, $T1$ nodes form a clique using peering links, even if they do not overlap. A more detailed analysis and justification of this provider selection model has appeared in [19].

Routing: In GENESIS, interdomain routing follows the shortest path subject to two common policy constraints: "prefer customer over peer over provider links" and satisfy the "valley-free" routing property.

Initial topology: To create the initial network topology, we select nodes sequentially and at random. For a selected node x, we determine its provider randomly from the set of nodes that *(a)* overlap with x, *(b)* are not in the customer tree of x, and *(c)* have greater geographic expanse than x. If we cannot find a provider for a node, then it joins the clique of tier-1 networks.

Network formation process: An execution, or sample path, of GENESIS proceeds in discrete time units called *iterations*. In each iteration, every node *plays* asynchronously once. The order in which nodes play during an iteration is determined at the start of the sample path, and it remains the same throughout that sample path. Every time a node plays, it carries out the following actions: *(a)* Examine depeering with existing peers, *(b)* Examine peering with new peers, *(c)* Provider selection, and *(d)* Peering strategy update. At the end of each iteration, we recompute the fitness of each node. If none of the nodes has adjusted its connectivity and peering strategy in that iteration, then it is easy to show that there will be no changes in subsequent iterations, and we say that GENESIS has reached an *equilibrium*.

The *state* of the network at any point in time can be defined based on the connections and peering strategy of all nodes. Two states **A** and **B** are distinct if they differ in terms of the underlying network topology or the peering strategy of one or more nodes. Even if we start with the same population of nodes and the same initial topology, two sample paths can result in two distinct equilibria as a result of different playing orders.

IV. Default Model & Validation

In Section III, we described the various components of GENESIS, without mentioning any parameter values or specific peering strategies. Here, we present a certain instance of GENESIS that we refer to as the "Default model" and that we use in the rest of the paper. The values of the Default model parameters are shown in Table I, together with a brief explanation for each parameter.

The peering strategies that we consider in this paper are the following three:

1) *Restrictive*: A node that uses this strategy does not peer with any other node unless if that is mandatory to maintain global reachability ("peering-by-necessity"). This peering strategy is followed only by T1 nodes; those nodes form a clique to keep the network connected.
2) *Open*: A node that uses this strategy agrees to peer with any other node that it overlaps with (except direct customers). In the default GENESIS model, the Open peering strategy is followed by stubs because those nodes aim to reduce their transit costs by peering with as many other nodes as possible.
3) *Selective*: A node x that uses this strategy agrees to peer with node y if $\frac{V_x}{V_y} \leq \sigma$ ($\sigma > 0$). In practice, there is a wide range of Selective peering strategies with several additional constraints and parameters (e.g., a minimum link capacity or a minimum number of points-of-presence) [20]. Our Selective strategy is only a model that aims to capture the essence of those requirements through a simple formula and single parameter σ. The Selective peering strategy in the default GENESIS model is followed by non-T1 transit providers.

A. Validation

The validation of models such as GENESIS is inherently difficult because there is a lack of available data about the economics of ISPs and about the interdomain traffic matrix and traffic flow. Even the AS-level Internet topology is not accurately known because only a small fraction of peering links are visible through BGP route monitors. Consequently, we have carried out a best-effort approach to validate GENESIS, examining whether it produces certain well-known quantitative or qualitative properties of the Internet. This effort should be viewed as a sanity check, rather than a comprehensive validation process, in view of the previous difficulties.

Average Path Length: The average path length of the AS-level Internet has remained almost constant at 4 AS-hops, at least during the last twelve years [21]. GENESIS produces networks with approximately the same average path length (3.7 AS-hops) with 500 nodes. Additionally, the average

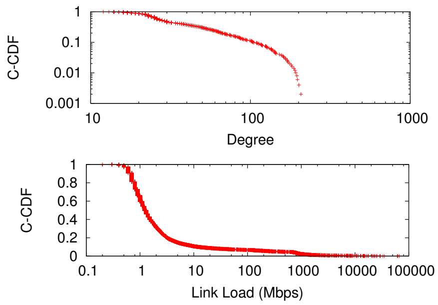

Fig. 1. Degree & Link load distribution

path length does not increase significantly as we increase the number of nodes from 500 to 1,000.

Degree Distribution: Figure 1 shows the complementary CDF (C-CDF) of the degree distribution that results from the default GENESIS model. Though not a strict power law (the tail of the distribution is truncated due to the limited number of nodes), it is clear that the degree distribution is highly skewed, with few nodes having a much larger degree than most other nodes.

Link Load Distribution: Figure 1 also shows the distribution of link traffic loads at equilibrium. Most links in the network are peering links that carry small traffic volumes. There are, however, few links that carry several orders of magnitude more traffic; those are mostly transit and peering links between large transit providers, or between transit providers and major content producers. Akella et al. [22] have observed a similarly skewed distribution of link loads in the Internet.

Fraction of transit providers: In the Default model, about 10% of the nodes end up as transit providers at equilibrium, which is the similar to the corresponding fraction reported in the measurement study by Dhamdhere et al. [21].

V. Stability and Oscillations

The network formation process in GENESIS is deterministic (the playing order does not change during a sample path), and so it can have one of two possible outcomes: either convergence to an equilibrium or a *limit cycle* in which the network moves repeatedly through the same sequence of states. An equilibrium results when none of the nodes changes its connections or peering strategy during an iteration. A limit cycle, or oscillation, on the other hand, results when the state of the network $S(t_k)$ in the k'th iteration is identical to the state $S(t_{k-m})$ in the (k-m)'th iteration with $m > 1$. The length of the limit cycle in that case is m. Note that there is no possibility for chaotic trajectories because the system has a finite number of states, and so if it does not converge to an equilibrium, it will eventually return to a previously visited state.

How often does the network formation process results in oscillations in practice? We ran many simulations with distinct node populations and observed that oscillations take place in about 10% of the sample paths. Even though this is not a large percentage, it is also not negligible and it has an important implication: if GENESIS is a good model for the Internet, we should not expect that the the interdomain connectivity in the latter will be stable. We return to this point at the end of this section. We also examined whether the frequency of oscillations increases with the number of nodes, but we did not observe any significant difference.

A. *Why to study equilibria?*

The Internet is in a constant state of flux, as new ASes and IXPs appear, and the traffic matrix, prices and costs vary with time. The reader may wonder, why to study equilibria in a system that is constantly evolving?

The fact that GENESIS leads to an equilibrium, at least in 90% of the sample paths, means that as long as those exogenous parameters (set of nodes, locations, traffic matrix, costs, etc) remain constant, the network formation process converges to a stable point. If this convergence takes place quickly, compared to the time scales in which those exogenous parameters change, the evolution of the Internet can be thought of as a trajectory through a sequence of equilibria, with a new equilibrium resulting every time there is a change in one of the previous exogenous parameters.

B. *Oscillations*

Causes of oscillations: We analyzed many oscillatory sample paths, attempting to identify a common cause behind all of them. *Every oscillation involves at least one transit provider that moves in and out of the tier-1 clique.* We explain this pattern with a simple example shown in Figure 2.

Let x, y and z be three nodes, with x and y being transit providers. The status of z does not affect the oscillation. x and z are initially transit customers of y. Consider the following sequence of actions.

1) The non-T1 provider x, using the Selective strategy, peers with z, as shown in Figure 2 (a). This reduces the transit volume of y (as traffic $x \leftrightarrow z$ now bypasses y) to the point that x cannot find a provider that has larger transit volume than itself.
2) Then, x acquires tier-1 status and switches to the Restrictive peering strategy. x and y peer because they are both tier-1 providers.
3) x depeers z, because z does not qualify as peer under the Restrictive strategy.
4) As a result of step-3, V_T of y is restored (as traffic $x \leftrightarrow z$ has to now go through y).
5) y depeers x, since it now has larger transit volume and it can become the provider of x.
6) As a result of step-5, x becomes a tier-2 provider and it switches back to the Selective peering strategy, as shown in Figure 2 (c).
7) x peers again with z, which returns us to the initial state.

The previous example is only an abstraction of more complex oscillatory sample paths that occur in practice, involving more nodes and more steps.

TABLE I
INPUT PARAMETERS

Parameter, Symbol, Description	Value	Explanation
Number of ASes N	500	Simulation time constraints
Number of geographic locations G_{Max}	50	Based on approximate ratio of IXPs to peering networks in the Internet. PeeringDB ratio 10.27. GENESIS ratio 10.0 [23]
Geographic expanse distribution	Zipf(1.6)	Based on data about number of participants at each IXP collected from PeeringDB [23]. $G(x)$ assigned randomly to each node
Maximum expanse for an AS	15	
Generated traffic distribution	Zipf(1.2)	Produces a heavy-tailed distribution of outgoing traffic. With this distribution, 0.1% of the ASes generate nearly 28% of the total traffic. This is consistent with the behavior reported in [24], [25] & [1], which show that the traffic produced by high-ranking ISPs and content providers follows a Zipf distribution. $V_G(x)$ assigned randomly to each node
Consumed traffic distribution	Zipf(0.8)	Produces heavy-tailed distribution of incoming traffic, similar to measured traffic distribution at Georgia Tech. $V_C(x) \propto \|G(x)\|$, rationale being that a node with large expanse will also have a large number of access customers
Mean consumed traffic	500 Mbps	
Private peering threshold Ω	50 Mbps	Survey of peering strategies [20]
Transit cost multiplier range $P(x)$	$[35,45]/Mbps per iteration	Parameterized based on IP transit prices advertised by VoxNet [26]. $P(x)$ assigned randomly
Transit cost exponent τ	0.75	Parameterized based on data from [16] and [20]
Peering cost multiplier α	$20/Mbps per iteration	
Peering cost exponent β	0.40	
Selective peering ratio σ	2.0	

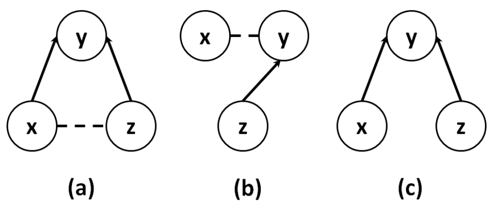

Fig. 2. Oscillation example

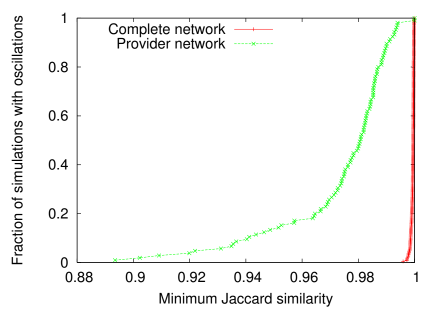

Fig. 3. CDF of minimum Jaccard similarity between network states during an oscillation

Network-wide effects of oscillations: We next study the effect of oscillations on the entire network as well as the affected nodes. We use the *Jaccard similarity* metric to quantify the similarity between two network states. If L_A and L_B are the sets of links in states **A** and **B** , respectively, then the Jaccard similarity between the two states is defined as:

$$J_{AB} = \frac{|L_A \bigcap L_B|}{|L_A \bigcup L_B|}$$

Note that a link is defined by both its two end-points as well as its type (transit versus peering). The value of $J_{AB} = 0$ indicates that the two networks do not have any link in common, whereas $J_{AB} = 1$ indicates that the two states are exactly the same. In addition to the complete network, we also examine the sub-network composed only of transit providers, excluding stubs and their links. We refer to this sub-network as the *provider network*.

For each oscillatory sample path, we compute the minimum Jaccard similarity across all pairs of states during an oscillation. Figure 3 shows the distribution of these minimum similarity values, across all oscillatory sample paths. When we examine the complete network, the minimum Jaccard similarity is almost equal to one in all oscillations, i.e., the oscillation affects the connectivity of only few nodes. We observed that in 95% of oscillations the number of nodes that undergo any change (connectivity and/or strategy) during the oscillation is less than 5% of the total network population. Note that all nodes involved in oscillations are not transit providers; some nodes are stubs that act as passive participants in the oscillation. When we only examine the provider network, the corresponding minimum Jaccard similarity values are lower but still higher than 0.90. In other words, oscillations cause more fluctuations in the connectivity of transit providers but those fluctuations are still mostly local in scope.

Length of oscillations: We measured that 95% of the oscillatiory sample paths have a limit cycle length that is less than 9 iterations. This length is positively correlated with the number of affected ASes (Pearson correlation coefficient: 0.92). Thus, the more nodes change connectivity or peering strategy during an oscillation, the longer it will take for the network to repeat its limit cycle.

How are oscillations relevant to the Internet? Peering agreements and negotiations are typically shrouded in secrecy, yet peering disputes between high-profile nodes over traffic ratios, content types, and peering conditions often generate attention. For example, Cogent and Sprint-Nextel terminated their peering relationship in 2008 after they failed to resolve a peering dispute. However, the peering relationship between them was restored after some time. Similar disputes have been reported between Level-3 and Cogent (2005), Cogent and Telia (2008), and Level-3 and Comcast (2010). Peering disputes between large providers often have significant effects, as single-homed customers of those nodes are unable to reach each other.

The fact that GENESIS can produce oscillatory sample paths should not be considered an artifact of the model. As described earlier, oscillations in GENESIS occur because some providers switch between tier-1 and non-tier-1 status, which resembles the peering disputes between tier-1 providers seen in practice. In the real world, such disputes do not persist and some exogenous mechanisms (such as revised contracts, negotiations, regulation or legal actions) are applied to break the impasse. GENESIS does not capture those exogenous mechanisms, but it captures that the endogenous dynamics of the network formation process cannot always manage to produce a stable network, and so some additional mechanisms should be designed to resolve such oscillations when they appear.

VI. Variability Across Equilibria

In this section we focus on those sample paths that converge to an equilibrium, and examine the differences between distinct equilibria. Specifically, given a population of nodes and an initial topology, does GENESIS always result in the same equilibrium? If not, how different are the resulting equilibria? What are the causes of these differences? Can we predict the properties of an individual node at equilibrium? Are there any global properties that are predictable with statistical significance? And who are the most predictable, or least predictable, nodes in an network, in terms of fitness?

Distinct equilibria: The network that will emerge from a GENESIS sample path highly depends on both the initial topology and the playing order. Specifically, given the same population of nodes and the same initial topology, 85% of the sample paths that converge to an equilibrium produce a distinct equilibrium. If the initial topology also varies across sample paths, this percentage increases to 90%. So, if GENESIS is a good model of the Internet, we understand that it is very hard in practice to predict the evolution of the Internet, as even minor changes in the order in which ASes act can have global effects.

Differences between equilibria: To quantify the difference between distinct equilibria, we again use the Jaccard similarity metric between the corresponding network states. Figure 4 shows the CDF of that metric for all pairs of distinct equilibria resulting from 100 sample paths. Note that the Jaccard similarity is typically higher than 0.90 when we only vary the playing order. When we also vary the initial topology, the similarity metric is lower (between 0.75-0.9) but still high in absolute terms. So, even though there are many distinct equilibria, they are quite similar in terms of topology.

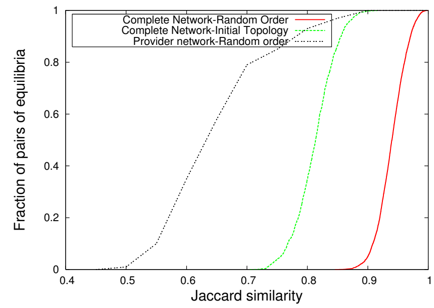

Fig. 4. CDF of Jaccard similarity between equilibria

Figure 4 also shows the distribution of Jaccard similarity for the subset of the network that includes only transit providers (the "provider network"). In this case, the Jaccard similarity is significantly lower, with 50% of the equilibrium pairs being less than 65% similar. In other words, even though stubs often end up with the same transit provider across different equilibria, the hierarchy of transit providers and the peering links between them are much less predictable.

Variability of fitness distribution: We next focus on the distribution of fitness values in an network. How different is this fitness distribution across different equilibria? We applied the two-sample Kolmogorov-Smirnov (KS) test on the fitness distribution across all pairs of equilibria. We did so both for the entire network, and for the provider network. We could *not* reject the null hypothesis that the distributions are identical for any pair of distributions at the 99% significance level. In other words, even though there are many distinct equilibria, the statistical distribution of the nodes' fitness remains invariant. However, this does not mean that the fitness of an individual node would remain the same across different equilibria; this is the subject of the next paragraph.

Variability of individual node fitness: Next, we look at the variability of the fitness of individual nodes across different equilibria. We compute the Coefficient of Variation (CoV) of the fitness of each node across 100 equilibria resulting from different playing orders. Figure 5 shows the CDF of the fitness CoV for the entire population of nodes. The CoV is close to zero for 90% of nodes indicating that most nodes see almost identical fitness in all equilibria. However, 5% of the nodes have CoV greater than 1.0 and 1% of the nodes have CoV as high as 3.0. Who are these "most unpredictable" nodes and what causes the large variability in their fitness?

We classify nodes in two classes. *Class-1* nodes are those that are either stubs at all equilibria, or that are transit providers at all equilibria. *Class-2* nodes, on the other hand, are those that are transit providers in some equilibria and stubs in others. Figure 5 also shows the CoV distribution for these two classes of nodes. Note that the CoV of Class-1 nodes is close to zero for almost all nodes, while the CoV of Class-

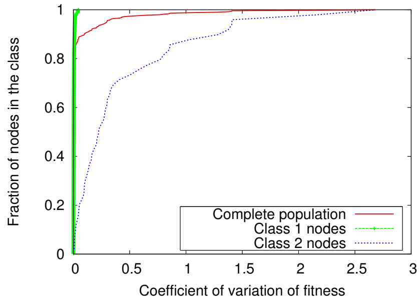

Fig. 5. CDF of fitness CoV for different network classes

2 nodes is significant for most of them. In the following, we discuss the properties of the nodes in each class in more detail.

Properties of Class-1 nodes: The largest fraction of nodes in this class are stubs (including nodes that were transit providers initially but are stubs at equilibrium). It is not surprising that the fitness of stubs does not vary significantly across equilibria. These nodes do not have revenues and their provider is always the cheapest node at the locations they are present at. Additionally, they use the Open peering strategy, and so their set of peers consists mostly of other stubs that are co-located with them. Additionally, however, Class-1 includes some large, in terms of geographical expanse, transit providers (12~15 locations). These providers have a large set of customers and so significant transit traffic. Their large transit volume and expanse restricts their set of providers and peers, reducing the variability in their connectivity and fitness across equilibria. Note that tier-1 providers always belong to Class-1.

Properties of Class-2 nodes: We observe that most nodes in this class have mid-range expanse (6~10 locations) and relatively high prices. Thus, while they are mostly unattractive as transit providers for stubs, smaller transit providers may choose them as provider, depending on their expanse and transit traffic volume. The latter, however, changes dynamically as a node gains or loses customers during a sample path. This aspect of the model creates a positive feedback effect where the larger the transit volume of a node is, the more it qualifies for being chosen as provider by other providers. However, who are the nodes that will benefit from this positive feedack strongly depends on the order in which the nodes act. In other words, the nodes in Class-2 are mostly nodes that end up as transit providers in some equilibria simply because they were "lucky" in those sample paths to get some customers early in the network formation process, accumulating transit volume and then attracting even more customers. It is these nodes that exhibit the largest variation in fitness and that act as the biggest source of variability in the resulting networks.

VII. Is it a Zero-Sum Game?

In the previous section, we showed that GENESIS can produce many distinct equilibria. While the fitness distribution does not change significantly across equilibria, the fitness of certain individual nodes varies widely. An important question is whether the gain of one node results in a corresponding loss for another node (or nodes) at the same equilibrium. In other words, *are the nodes of the network engaged in a zero-sum game, where the total fitness in the network remains constant across equilibria?* If so, under which conditions is this true? If not, when would the total fitness increase or decrease?

Consider a network that consists of a set $\mathcal{N}$ of nodes; a subset $\mathcal{P}$ of these nodes are transit providers while the remaining nodes (subset $\mathcal{S}$) are stubs. Let T_i, R_i and F_i represent the transit costs, peering costs, and fitness of network i, respectively. Consider a transit link between customer i and provider j, and let t_{ij} be the transit payment from i to j. The fitness of i includes the term $-t_{ij}$, while the fitness of j includes the opposite term t_{ij}. As a result, if we compute the total fitness across all nodes, the transit payments between customer-provider pairs cancel out. On the other hand, the cost that two nodes x and y incur for a peering link between them appears as a negative term in their respective fitness functions. Consequently, the total fitness of all nodes can be written as

$$\sum_{\mathcal{N}} F_i = -\sum_{\mathcal{N}} R_i \tag{6}$$

The previous equation shows that the set of nodes in $\mathcal{N}$ would participate in a zero-sum game if the sum of all peering costs is constant across equilibria. For the default parametrization in GENESIS, we observe that while the sum of all peering costs was not constant it did not vary significantly either. For instance, we measured the CoV of $\sum_{\mathcal{N}} F_i$ in 100 equilibria (resulting from different playing orders), and it was both low (0.04) and almost equal to the CoV of $\sum_{\mathcal{N}} R_i$. Hence, when we consider the entire population of nodes, the total fitness in the network does not vary significantly across equilibria and the underlying game is *approximately* zero-sum. This implies that that the resulting equilibria are approximately Pareto optimal, i.e., it is not possible to increase the fitness of a node without necessarily decreasing the fitness of any other node.

Let us now consider the subset of transit providers, and ask whether their total fitness remains constant across equilibria. Consider a transit link between customer i and provider j, with a transit payment t_{ij}. If i is a stub, then the term t_{ij} appears in the fitness of j. Since we compute the total fitness only of providers, however, the opposite term $-t_{ij}$ (which appears in the fitness of i) is not included in the summation. Consequently, the total fitness of all transit providers is

$$\sum_{\mathcal{P}} F_i = \sum_{\mathcal{S}} T_i - \sum_{\mathcal{P}} R_i \tag{7}$$

where the first term is the total amount of transit costs paid by stubs to providers, and the second term is the total costs for peering links only between providers. For the same 100 equilibria considered above for the entire population, the CoV of $\sum_{\mathcal{P}} F_i$ was lower (0.02) than the entire population.

So, the set of transit providers participate in a zero-sum game if the following three conditions are true across all equilibria: *(a)* the set of providers $\mathcal{P}$ (and so the set of stubs $\mathcal{S}$) remains the same, *(b)* the total transit payments from stubs remain constant, and *(c)* the sum of provider peering costs remain constant. Interestingly, we observed that the three conditions are typically true in GENESIS for the default parametrization, at least as an approximation and hence we can expect that the resulting equilibria are Pareto optimal for the population of transit providers.

However, an increase in σ, which allows transit providers to peer more openly, the peering costs show more variation resulting in higher CoV for $\sum_{\mathcal{N}} R_i$ and $\sum_{\mathcal{P}} F_i$. For example we observed the CoV of $\sum_{\mathcal{N}} R_i$ to be 0.08 and 0.27 and the CoV of for $\sum_{\mathcal{P}} F_i$ to be 0.07 and 0.73 for σ=10 and σ=1000 respectively. In an open peering environment the set of providers undergoes greater variation when subjected to different playing order. Thus, as peering increases in the network the notion of Pareto optimality for providers at equilibria becomes weaker.

VIII. Case Study: Peering Openness

Nodes engage in settlement-free peering mainly to reduce upstream transit costs and to connect directly to the sources and destinations of their traffic. However, nodes should $selectively$ decide which nodes to peer with because peering has its own costs and it increases the monitoring and management overhead. In GENESIS, a single parameter σ determines the peering openness of all nodes; a larger value of σ indicates an increased openness to peering. In this section, our goal is to study how this peering openness, captured by the parameter σ, affects the properties of the resulting network. In particular, we are interested in the value(s) of σ that results in the maximum fitness for different classes of nodes.

We focus on how the peering openness in the system affects the fitness, revenue and costs for three types of transit provider: tier-1 or T1 providers (they do not have a provider), tier-2 or T2 providers (their provider is a T1 node), and tier-3 or T3 providers (all other providers).

We simulate 20 different populations across 30 playing orders and examine a range of σ values which span the spectrum from “Restrictive” to “Open” peering. We then measure the fitness, revenue, and costs of each transit provider at different values of σ. We normalize the fitness of each node x in a population by its maximum absolute fitness. We similarly normalize the revenue and cost components of each node. A node x in a population is classified as belonging to a particular class of providers (T1, T2 or T3) if it belongs to that class in at least 80% simulations.

Figure 6 shows the average normalized fitness of all nodes in each class of providers as we vary σ. Figure 7 shows the average normalized transit costs, transit revenues and peering costs for tier-2 and tier-3 transit providers.

Fitness of T1 providers: As σ increases, nodes in the customer tree of T1 nodes peer increasingly with each other, and less traffic flows through T1 nodes. Consequently, T1 fitness shows a monotonic decrease as σ increases.

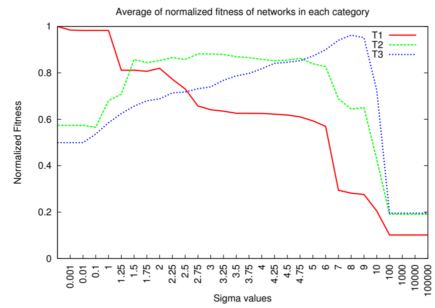

Fig. 6. Fitness variations of three transit provider classes versus σ

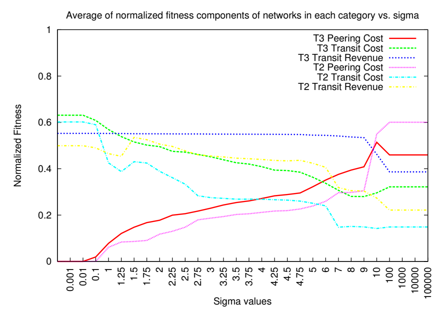

Fig. 7. Fitness components of tier-2 and tier-3 providers versus σ

Fitness of T2 providers: T2 providers in general have large transit volumes, due to their large customer base and expanse. For σ between 1 and 5, T2 providers are able to peer with other nodes that have similar traffic volumes (other T2 providers and major content providers), thus reducing their transit costs. At the same time, there is no major reduction in T2 transit revenues in this range of σ, as peering is not yet pervasive in the customer trees of T2 nodes. Additionally, T2 peering costs are low, as they are able to aggregate peering traffic on a few peering links. Consequently, their fitness is maximized in this range of σ.

Fitness of T3 providers: T3 providers and their customers generally have low traffic volumes and limited expanse, which limits their opportunities to peer. As σ increases, T3 providers are able to reduce their transit costs by peering. Simultaneously, their peering costs increase, but this increase is smaller than the decrease in transit costs. T3 nodes are less susceptible to losing transit revenue, as their customers have limited peering opportunities until σ increases to very large values. Figure 7 shows that the average transit revenue of T3 ASes remains almost constant until σ=10.0. The net effect is that T3 providers show a monotonic, albeit more gradual, increase in fitness as compared to T2 providers.

Open peering regime: Figure 6 shows a decrease in fitness of all provider categories beyond $\sigma = 10$. As we increase σ, we see an increase in peering between providers and stubs.

The more providers peer with stubs, the more they reduce the transit traffic (and hence revenues) of other providers. This effect, which we call *transit stealing*, becomes more pronounced beyond $\sigma = 10$. Thus, the threshold $\sigma \gtrapprox 10.0$ acts as the approximate edge of open peering in the network. The dynamics of the open peering regime are interesting. In this range of σ, increased peering reduces the transit traffic volume of most providers. A large value of σ implies that such providers agree to peer with stubs. Peering between providers and stubs reduces the transit traffic volume of the providers of those stubs, which further increases the willingness of those providers to peer with stubs. The result is a *positive feedback loop* where the more providers peer with stubs, higher is their collective openness to peering.

Global network properties: Finally, we study how the topological properties of the network change as we vary σ. We define *link density* as the ratio of the number of existing links to the number of possible links, given geographic constraints. We observed that link density increases from 70% for the default model ($\sigma = 2$), to 97% with *Open* peering ($\sigma \geq 10$). We also observed that the average path length decreases to 2.53 AS hops with *Open* peering, as opposed to 3.7 for the default model. An interesting aspect of increasing openness is that the the number of transit providers decreases from 10% of the population in the default case, to 7% with *Open* peering. As more ASes are able to reach their destinations via peering links, some transit providers lose all their transit traffic and become stubs.

IX. Conclusions

We proposed GENESIS, an agent-based network formation model that captures interdomain traffic flow, policy-based routing, geographic constraints, and the economics of transit and peering relations. GENESIS is a flexible and extensible tool that can be used to study the network formation dynamics and equilibria (or oscillations) that result under different conditions, parameters and provider/peer selection strategies. We examined the convergence properties of GENESIS, and found that GENESIS results in a stable network in most cases. The observed cases of instability occur due to transitions of a few nodes to and from the T1 clique, resembling real-world peering disputes between T1 providers. GENESIS shows both path dependence and dependence on initial conditions, i.e., it can produce different equilibria depending on the initial topology and playing order of nodes which we believe is also true of the Internet. We found that while these equilibria are distinct in terms of network topology, it is possible to predict certain properties of the network or of certain classes of nodes with statistical significance. We also showed that formation of the entire network can be thought of as an approximate zero-sum game under three conditions. As an application of GENESIS, we studied the properties of the network and the fitness of three classes of transit providers as we vary the level of peering openness.

In ongoing work, we use GENESIS to understand what happens when each node dynamically selects from a set of peering strategies the strategy that maximizes its fitness.

References

[1] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet Inter-domain Traffic," in *Proceedings of ACM SIGCOMM*, 2010.

[2] "Internet Transit Prices: Historical and Projected," http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php.

[3] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse?" in *Proceedings of the Passive and Active Measurement Conference (PAM)*, 2008.

[4] B. Augustin, B. Krishnamurthy, and W. Willinger, "IXPs: Mapped?" in *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, 2009.

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

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

[7] S.-H. Yook, H. Jeong, and A.-L. Barabasi, "Modeling the Internet's Large-scale Topology," *Proceedings of the National Academy of Sciences*, vol. 99, no. 21, pp. 13 382–13 386, 2002.

[8] X. Wang and D. Loguinov, "Understanding and Modeling the Internet Topology: Economics and Evolution Perspective," *IEEE/ACM Transactions on Networking*, 2010.

[9] L. Li, D. Alderson, J. C. Doyle, and W. Willinger, "Towards a Theory of Scale-free Graphs: Definition, Properties, and Implications," *Internet Mathematics*, 2005.

[10] 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.

[11] H. Chang, S. Jamin, and W. Willinger, "Internet Connectivity at the AS-level: An Optimization-driven Modeling Approach," in *Proceedings of the ACM SIGCOMM workshop on Models, methods and tools for reproducible network research (MoMeTools)*, 2003.

[12] J. Corbo, S. Jain, M. Mitzenmacher, and D. C. Parkes, "An Economically-Principled Generative Model of AS Graph Connectivity," in *Proceedings of the IEEE INFOCOM Mini-Conference*, 2009.

[13] S. Goyal, *Connections: An Introduction to the Economics of Networks*. Princeton University Press, 2007.

[14] M. O. Jackson, *Social and Economic Networks*. Princeton University Press, 2008.

[15] P. Holme, J. Karlin, and S. Forrest, "An Integrated Model of Traffic, Geography and Economy in the Internet," *ACM SIGCOMM Computer Communications Review (CCR)*, 2008.

[16] H. Chang, S. Jamin, and W. Willinger, "To Peer or Not to Peer: Modeling the Evolution of the Internets AS-level Topology," in *Proceedings of IEEE INFOCOM*, 2006.

[17] A. Dhamdhere and C. Dovrolis, "The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh," in *Proceedings of ACM CoNEXT*, 2010.

[18] "GENESIS," http://www.cc.gatech.edu/~dovrolis/Papers/genesis.pdf.

[19] A. Lodhi and C. Dovrolis, "A Network Formation Model for Internet Transit Relations," in *Proceedings of the 2010 Workshop on Economics of Networks, Systems, and Computation (NetEcon)*, 2010.

[20] "Peering Strategy Survey," http://drpeering.net/white-papers/Peering-Policies/A-Study-of-28-Peering-Policies.html.

[21] A. Dhamdhere and C. Dovrolis, "Twelve Years in the Evolution of the Internet Ecosystem," *IEEE/ACM Transactions on Networking*, 2011.

[22] A. Akella, S. Seshan, and A. Shaikh, "An Empirical Evaluation of Wide-area Internet Bottlenecks," in *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, 2003.

[23] "PeeringDB," http://www.peeringdb.com, September 2010.

[24] H. Chang, S. Jamin, Z. Mao, and W. Willinger, "An Empirical Approach to Modeling Inter-AS Traffic Matrices," in *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, 2005.

[25] A. Feldmann, N. Kammenhuber, O. Maennel, B. Maggs, R. De Prisco, and R. Sundaram, "A Methodology for Estimating Interdomain Web Traffic Demand," in *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, 2004.

[26] "VoxNet IP Services," http://www.voxel.net/ip-services.