

RESEARCH

Open Access

GENESIS-CBA: an agent-based model of peer evaluation and selection in the internet interdomain network

Aemen Lodhi^{1*}, Amogh Dhamdhere² and Constantine Dovrolis¹

*Correspondence:

aemen.lodhi@gatech.edu

¹School of Computer Science,
Georgia Institute of Technology,
Atlanta, USA

Full list of author information is
available at the end of the article

Abstract

Purpose: We propose an agent-based model for peer selection in the Internet at the Autonomous System (AS) level. The proposed model, *GENESIS-CBA*, is based on realistic constraints and provider selection mechanism, with ASes acting in a myopic and decentralized manner to optimize a cost-related fitness function. We introduce a new peering scheme, *Cost-Benefit-Analysis*, which, unlike existing peering strategies, gives ASes the ability to analyze the impact of each peering link on their economic fitness. Using this analysis, ASes engage in only those peering relations that can have a positive impact on their fitness.

Methods: Our proposed model captures the key factors that affect 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-CBA* 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.

Results: We find that such oscillatory sample paths occur in about 7% of the runs, and they always involve tier-1 ASes. *GENESIS-CBA* results in many distinct equilibria that are highly sensitive to initial conditions and the order in which ASes (agents) act.

Conclusions: Our results imply 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.

Keywords: Internet interdomain network; Peering strategies; Cost-benefit-analysis; Internet autonomous systems

Background

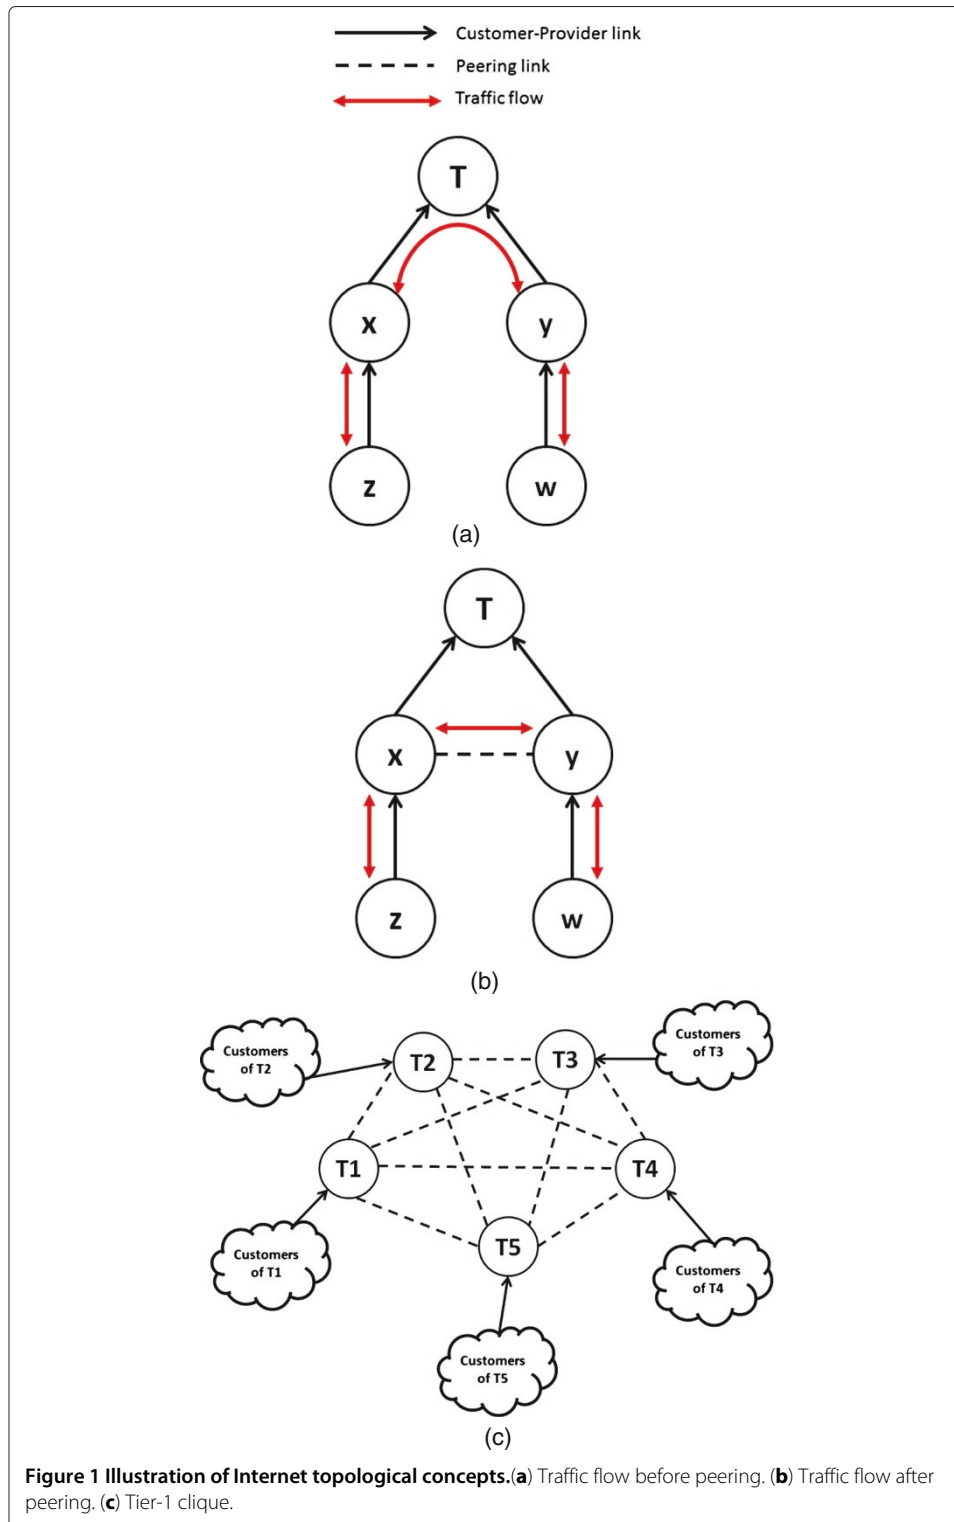
Tens of thousands of Autonomous Systems (ASes) interconnect in a complex and dynamic manner to form the Internet. Each AS is an independent entity under single administrative control, and is identified through a unique AS number (ASN), which is assigned by a central authority, the Internet Assigned Numbers Authority (IANA) (Internet Assigned Numbers Authority 2012). Each AS manages its own internal network topology and traffic routing (intradomain activity) and connects to the rest of the world

by forming links with other geographically co-located ASes (interdomain activity). These ASes generate traffic destined for other ASes in the Internet, consume traffic destined for them, and some of them transit traffic, i.e., they carry traffic on behalf of other ASes to transport it to its destination.

ASes belong to different categories: enterprise networks (or stubs), e.g., university campuses, banks etc.; content sources, which generate much more traffic than they consume, e.g., Google, Facebook, YouTube etc.; access providers, which consume much more traffic than they generate, e.g., residential ISPs, Comcast, France Telecom; transit providers, which carry traffic on behalf of other ASes to its destination, e.g., Level 3 Communications, Cogent, etc. Various combinations of the aforementioned categories also exist, e.g., ASes that provide both transit and access services such as AT&T. Thus, the Internet is a network of ASes where traffic flows from one AS to another en route to its destination. A traffic flow may traverse multiple transit providers before reaching its destination AS. Additionally, ASes 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. Under the customer-provider relationship, the provider is responsible for connecting the customer with the rest of the Internet. The provider may itself have a transit provider and so on. Settlement-free peering (subsequently referred to as “peering” for brevity) is more subtle than a customer-provider relationship. In peering, two ASes (or peers) agree to exchange traffic between themselves and their customers over a shared link. Peers do not pay each other for this exchange; however, they share the cost of the link.¹ For example, ASes peering at an Internet Exchange Point (IXP), use the IXP’s infrastructure for peering. Thus, the IXP’s infrastructure forms the shared link between the peers. While they do not pay one another for the traffic that they exchange among themselves, they pay the IXP for usage of its infrastructure.

Peering is typically carried out with the objective of saving transit costs by diverting traffic from customer-provider links to peering links. ASes generally follow broad rules, referred to as “peering strategies” to determine their peering links. Examples of such rules-of-thumb include *Open* (peer with all co-located ASes except customers), *Selective* (peer with only those ASes that have similar traffic volume), *Restrictive* (do not peer with anyone unless the peering relationship is necessary to maintain global reachability). The highest level of the Internet consists of a clique of approximately 15 ASes fully meshed among themselves through peering links. This clique is known as the “Tier-1” clique and these ASes are known as “Tier-1 transit providers”. Tier-1 providers do not require a transit provider as they can reach the entire Internet through their customers and peers. Figure 1 provides an illustration of these topological concepts. In Figures 1a and 1b x , y , z , w and T are ASes; T is a Tier-1 provider, x and y are transit providers, and w and z are stubs. Figure 1a shows the flow of traffic between transit providers and stubs through the Tier-1 provider. In Figure 1b, the introduction of a peering link between x and y causes the traffic to flow over a different path, avoiding the Tier-1 provider altogether. This results in reduced transit costs for x and y . Figure 1c illustrates the concept of the Tier-1 clique in the Internet. The ASes shown are Tier-1 providers, which form a complete mesh among themselves through peering links. Traffic between their customer trees is routed through these 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. The Internet remains in a persistent state of flux, subject to

changes in various exogenous factors, e.g., geographic expansion of ASes, new entrants in the connectivity market, removal of nodes from the network owing to bankruptcies, changes in technology, etc. How will the Internet change due to consolidation of content (Labovitz et al. 2010), large penetration of video streaming, falling transit prices (Internet Transit Prices: Historical and Projected 2010), expanding geographic footprint of content providers (Gill et al. 2008), cheap local availability of peering infrastructure at IXPs (Augustin et al. 2009)? We proposed a computational agent-based network formation model, called “GENESIS” (Lodhi et al. 2012b), 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 the context of the interdomain connectivity, ASes face an important question: which other ASes should they peer with? GENESIS does not empower agents to choose their peering strategy; instead, a peering strategy is assigned to each agent based on its hierarchical status: Tier-1 providers use *Restrictive*, all other transit providers use *Selective* and stubs use *Open* peering strategy. In this paper, we take the concept of choosing peers much further than just choosing a peering strategy and then peering with all those ASes that satisfy the chosen strategy. We modify the peering mechanism in GENESIS by introducing a *Cost-Benefit-Analysis* (CBA) scheme, wherein transit providers analyze the impact of each peering link on their economic fitness and then decide whether to peer or not. We refer to this modified version of GENESIS as GENESIS-CBA

CBA is not just another peering strategy like *Selective*, *Restrictive* or *Open*; it is a completely different way of deciding which ASes to peer with. Other peering strategies are rules-of-thumb and are used to approximate cost-benefit-analysis. They are used because they are easier to deploy. However, our work shows that such informal analysis can lead to situations where ASes end up choosing peering strategies detrimental to their long-term economic fitness (Lodhi et al. 2012c). With CBA, agents can engage in precise analysis of the impact of each peering link on their fitness.

In this paper, we first introduce the main features of GENESIS and then introduce our cost-benefit-analysis scheme, as proof-of-concept, to study the properties of the resulting networks. Starting from a population of ASes and an 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.

Related work

Most 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 (Faloutsos et al. 1999)) (Wang and Loguinov 2006; Yook et al. 2002; Wang X and Loguinov 2010). Another class of models takes a bottom-up approach (Li et al. 2005), modeling the optimization objectives and constraints of individual ASes, to create graphs that have the same topological properties (Li et al. 2005; Fabrikant et al. 2002; Chang et al. 2003; Corbo et al. 2009).

The GENESIS model was first introduced by Lodhi et al. (2012b). GENESIS differs from previously mentioned related work 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. The main difference between GENESIS (Lodhi et al. 2012b) and this work is the mechanism used by transit providers to determine their peers. In GENESIS, transit providers use rule-based peering strategies: Tier-1 transit providers are assigned *Restrictive* while all other transit providers are assigned *Selective* strategy. A transit provider using Restrictive strategy does not peer with any other agent unless the absence of a peering relationship results in a disconnected network; a transit provider using Selective strategy peers with agents which have similar traffic volume (it uses the ratio of total traffic volumes to determine the similarity of other agents to itself). Such rule-based peering strategies may lead to extraneous costs from redundant peering links and may prevent beneficial peering links from being formed as shown by Lodhi et al. in (2012a). GENESIS-CBA does away with such rule-based peering strategies and introduces a new *Cost-Benefit-Analysis* module to GENESIS. Additionally, GENESIS-CBA does not assign peering strategies to transit providers based on their hierarchical status; it empowers them to analyze the impact of each peering link on their fitness and then choose only those peers that can improve their fitness.

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 (Goyal 2007; Jackson 2008). Those models capture the objectives and potential strategies of each node as players in a non-cooperative game, and 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-CBA. 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. Holme et al. propose an agent-based network growth model (Holme et al. 2008). However, they treat link formation as a random process, and do not capture link rewiring, settlement-free peering, and use hot-potato routing instead of realistic policy based shortest-path Internet routing. The model of Chang et al. (2006) uses hard-wired strategies for provider and peer selection, among other differences. The ITER model of Dhamdhere et al. (2010) is more similar to GENESIS-CBA, but it assigns a specific function to each agent (e.g., content provider, enterprise customer etc.) Additionally, peering strategies are pre-assigned to the agents. In our work both the hierarchical position of the agent and its peering strategies are an outcome of the model. Their goal is to study the evolutionary transition of the Internet from a *hierarchical* to a *flat* structure, while our focus is on the peer selection mechanism adopted by transit providers.

Methods

The GENESIS model

We consider a population of N agents (representing Autonomous Systems) which interconnect through two types of links: customer-provider and peering. Each agent has the following attributes: a set of geographic locations in which it is present (corresponding to IXPs), a certain traffic volume that it sends to and receives from every other agent, and certain economic parameters, such as the transit price it would charge to its potential customers at a given location. We do not assign a business function to agents *a priori*; an agent may act as a transit provider, content/access provider, enterprise organization, or any combination of these functions depending on its generated/consumed traffic, transit pricing, or geographical presence. An inter-network traffic matrix gives the average traffic volume sent from each agent to every other agent. The model only captures the interconnections between distinct agents and does not consider intradomain activity and connectivity. The geographic locations, transit prices and traffic flow of each agent remain constant during the timescales of interest in this model. We focus on the impact of the dynamics of provider and peer selection: how does the internetwork and the economic fitness of agents change with time as agents choose the best transit provider and peers for themselves? Table 1 provides a list of all notations used in the model.

Geographical presence

Each agent is placed in one or more geographic locations. These locations are roughly analogous to the IXPs in the real world. The *world* consists of G_M locations and an agent x is present at a set $G(x)$ of randomly assigned locations. Two agents are *co-located* if they share at least one location. For agent x , $O(x)$ denotes the set of agents that are co-located with x . A link between two agents can be formed only if they are co-located. An exception to this rule exists if x and y are non-colocated *Tier-1* agents. This exception is introduced to ensure that the network is always connected.

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 agent x and consumed by agent y . Overall, x generates traffic $V_G(x)$ and when aggregated across all consumers of traffic it is given by:

$$V_G(x) = \sum_{i=1}^N T_{xi} \quad (1)$$

Overall, x consumes traffic $V_C(x)$ and when aggregated across all producers of traffic it is given by:

$$V_C(x) = \sum_{i=1}^N T_{ix} \quad (2)$$

$T_{xx}=0$, i.e., we do not capture exchange of traffic locally within the agent. 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 an agent depends on the underlying network topology and the routing algorithm. Even if the interdomain traffic matrix T is constant, the transit traffic

Table 1 List of abbreviations & notations

Notation/Abbreviation	Explanation
AS	Autonomous System
IXP	Internet Exchange Point
CBA	Cost Benefit Analysis
$O(x)$	Set of agents geographically co-located with x
$G(x)$	Set of geographic locations where x is present
$C(x)$	Set of agents which are customers of x
$S(x)$	Peering strategy of x
$R(x)$	Transit provider of x
T_{xi}	Traffic originating in x and destined for i
$V_G(x)$	Traffic generated by x
$V_C(x)$	Traffic consumed by x
$V_T(x)$	Traffic transiting x
$V_P(x)$	Traffic exchanged between x and its provider
$V(x)$	Total traffic managed by x
V_{xy}	Total traffic exchanged between x and y in then network
$TC(x)$	Transit cost of x
$TR(x)$	Transit revenue of x
$PC_{priv}(x)$	Private peering cost of x
$PC_{pub}(x)$	Public peering cost of x
$F(x)$	Fitness of x
$PP_{pub}(x)$	Public peers of x
$PP_{priv}(x)$	Private peers of x
$PP(x)$	All peers of x
$N(t)$	State of network at time t . The state of the network is defined by the topology of the network and the peering strategies of all agents
α	Peering cost per Mbps
β	Economies of scale exponent for peering traffic
τ	Economies of scale exponent for transit traffic
Ψ	Private peering threshold traffic volume (Mbps)

of an agent may change as the underlying topology changes. The total traffic volume of an agent $V(x)$ is given by the following expression:

$$V(x) = V_C(x) + V_G(x) + V_T(x) \quad (3)$$

Economic attributes

The economic attributes of an agent include its transit price, transit and peering costs, transit revenue and *fitness*.

Prices Each agent x maintains a vector of prices $P(x)$. Each element of the vector corresponds to the price the agent has in each location in which it is present. ASes generally maintain different prices at different geographic locations to reflect the underlying differences in competition. The price of agent x in each location is determined randomly, independent of its prices in other locations and the prices of co-located agents. The random prices are chosen from a specific range.

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 are co-located. If $V_P(x)$ is the traffic exchanged between x and y , then the transit payment from x to y is:

$$TC(x) = P_y(x) \times V_P(x)^\tau \quad (4)$$

where τ is a transit traffic exponent that captures the *economies-of-scale* observed in practice. The transit revenue of y , $TR(y)$ is the aggregate of transit costs incurred by all transit customers of y . Let $C(y)$ denote the set of all transit customers of y . Then its transit revenue is given by:

$$TR(y) = \sum_{\forall x \in C(y)} TC(x) \quad (5)$$

Peering cost Agents 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 Mbps they peer publicly at an IXP, otherwise they peer privately. In public peering, agents 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 (Norton 2011). (DE-CIX) provides an example of Internet peering costs at a European IXP.

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

$$PC_{prv}(x) = \alpha \times \sum_y V_{prv}(x, y)^\beta \quad (6)$$

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

Public peering cost Let $V_{pub}(x, z)$ be the traffic exchanged between agent 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:

$$PC_{pub}(x) = \alpha \left(\sum_z V_{pub}(x, z) \right)^\beta \quad (7)$$

Fitness The fitness of an agent represents its net profit,

$$F(x) = TR(x) - TC(x) - PC_{pub}(x) - PC_{prv}(x) \quad (8)$$

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

Peering

Each agent x has a peering strategy $S(x)$. Two agents x and y are potential peers if (a) x and y are co-located (b) x and y do not have a customer-provider relationship (c) the traffic exchanged between x and y over the proposed peering link, $V_{pp}(x, y)$, will be greater than a certain threshold Ω Mbps. The above conditions define *universal* peering constraints

for all pairs of agents. However, potential peers x and y can establish a peering relationship if and only if they satisfy the constraints of each other's peering strategies.

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 if and only if the constraints of both $S(x)$ and $S(y)$ are satisfied. Depeering, however, is a unilateral decision by one of the peers.

The set of all public and private peers of x is denoted by $PP_{pub}(x)$ and $PP_{prv}(x)$ respectively. $PP(x)$ denotes the set of all peers of x .

$$PP(x) = PP_{pub}(x) \cup PP_{prv}(x) \quad (9)$$

Let $y \in S(x)$ denote that y satisfies the constraints of peering strategy of x . Algorithm 1 shows the peering acquisition algorithm for node x . Algorithm 2 shows depeering decision by node x .

Algorithm 1 Peer acquisition

```

Construct  $PP_{potential}(x) = \{y : y \in O(x), y \notin C(x), x \notin C(y), y \notin PP(x)\}$ 
For  $\forall y \in PP_{potential}(x)$ 
  If  $y \models S(x)$ 
    If  $V_{xy} \geq \Psi$ 
      Do  $PP_{prv}(x) = PP_{prv}(x) + y$ 
         $PP_{prv}(y) = PP_{prv}(y) + x$ 
    Else
      Do  $PP_{pub}(x) = PP_{pub}(x) + y$ 
         $PP_{pub}(y) = PP_{pub}(y) + x$ 
  End
Do  $PP(x) = PP_{prv}(x) \cup PP_{pub}(x)$ 

```

Algorithm 2 Depeering

```

For  $\forall y \in PP(x)$ 
  If  $y \not\models S(x)$ 
    If  $y \in PP_{prv}(x)$ 
      Do  $PP_{prv}(x) = PP_{prv}(x) - y$ 
    Else
      Do  $PP_{pub}(x) = PP_{pub}(x) - y$ 
  End
Do  $PP(x) = PP_{prv}(x) \cup PP_{pub}(x)$ 

```

In this paper, we consider two peering strategies: *Open* and *Cost-Benefit-Analysis* (described later). We emphasize that GENESIS is not limited to these particular strategies. In our previous work (Lodhi et al. 2012b) we have implemented and analyzed *Selective* and *Restrictive* strategies.

Provider selection

An agent must have a transit provider if it cannot reach all other agents in the network via its peers and customers. Agent x selects a provider y if: (a) x is co-located 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 agents that satisfy the previous

constraints. We say that an agent y is larger than an agent x if y is present in at least as many locations as x , and it carries more transit traffic than x . Formally, the provider of node x denoted by $R(x)$ is given by:

$$R(x) = \arg \min_{y \in O(x)} P(y) : V_T(y) > V_T(x), |G(y)| \geq |G(x)|, y \notin PP(y), y \notin C(PP(x)) \quad (10)$$

If an agent 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 agents 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 (Lodhi and Dovrolis 2010).

Routing

In GENESIS, interdomain traffic is routed as in the real-world Internet. Traffic follows the shortest path subject to two common policy constraints: “prefer customer over peer over provider links” and satisfy the “valley-free” routing property. This traffic routing scheme is illustrated in Figures 1a and Figures 1b.

Initial topology

The model is initialized using an initial network topology. To create the initial network topology, we select agents sequentially and at random. For a selected agent x , we determine its provider randomly from the set of agents 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 an agent, then it joins the clique of tier-1 agents.

Network formation process

An execution, or sample path, of GENESIS proceeds in discrete time units called *iterations*. In each iteration, every agent *plays* asynchronously once. The order in which agents play during an iteration is determined at the start of the sample path, and it remains the same throughout that sample path. Every time an agent plays, it carries out the following actions in order: (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 agent. If none of the agents 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 agents. 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 agents and the same initial topology, two sample paths can result in two distinct equilibria as a result of the different order in which nodes may play.

The network formation process in GENESIS is deterministic 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. We describe the stability

of the model in Section ‘Stability of GENESIS-CBA’. The network formation process is formally described in Algorithm 3.

Algorithm 3 Network formation process

Input:
 $X \leftarrow$ Ordered set of agents $\{x_0, x_1, x_2, \dots, x_N\}$

Begin
 $t \leftarrow 0$

Start Iteration $i \leftarrow 0$

Play agent x_i

Examine depeering for agent x_i

If x_i depeered any peer in examination above

Do Compute $V_T(x_k) \forall x_k \in X$

Examine peering for agent x_i

If x_i acquired any new peer in examination above

Do Compute $V_T(x_k) \forall x_k \in X$

Examine provider selection for agent x_i

If x_i changed its provider in examination above

Do Compute $V_T(x_k) \forall x_k \in X$

If $i == N$

Do Compute $F(x_k) \forall x_k \in X$

Do Record $S(t)$

If $t > 0$ && $N(t) == N(t - 1)$

End *Equilibrium*

Else If $t > 1$ && $N(t) == N(t') : \exists t' < t - 1$

End *Oscillation*

Else

$t \leftarrow t + 1]$

Go to **Start Iteration**

Else

$i \leftarrow i + 1$

Go to **Play**

Peering strategies in GENESIS-CBA

The economic objective of peering is to offload as much upstream transit traffic on provider links to settlement-free peering links thus saving transit costs. We consider the following two peering strategies in this paper:

Cost benefit analysis

When using Cost Benefit Analysis (CBA) peering strategy, an agent x separately evaluates the impact of each peering link (existing or potential) on its fitness, as opposed to broad rule-based peering strategies, e.g., Selective, Open, etc. For example, under Open strategy, x would peer with all potential peers regardless of the impact of each peering relationship on its fitness. Under CBA, x agrees to establish (or retain) only those peering links that can have a positive impact on its fitness. In GENESIS-CBA, only transit providers employ the CBA strategy.

In order to establish a new peering link using CBA, x evaluates each potential peer y by hypothetically establishing a peering link with it. Each new link in the network causes some traffic to be re-routed along different paths based on the interdomain routing policy. x computes the impact of the hypothetical peering link by updating its traffic flow, costs and revenue. x agrees to peer with y if and only if its fitness increases as a result of the hypothetical link. In the real world, such link evaluations are carried out through *peering trials*. Under a peering trial, two ASes establish a peering relationship for a short time period and evaluate the impact of the temporary link on their economic fitness. The link is retained if both peers find the relationship beneficial for their fitness, otherwise it is terminated. Similar evaluations are carried out for existing peers when using CBA. Each existing peer is hypothetically depeered and the impact of depeering on fitness is computed. The link is terminated if depeering improves fitness, otherwise it is retained.

Scalability of CBA We assume that an agent x can choose any number of peers from $PP(x)$ (from 0 to $n-1$). Traffic in the interdomain network is routed based on the shortest path prefer-customer-then-peer-then-provider routing policy. The addition (or removal) of a peering link may affect the flow of traffic on other peering links. Thus, for x , peering with agent y_i may impact the evaluation of agent y_j . We explain this by means of the following example. Let y_j be a customer of y_i . Let V_{xy_i} and V_{xy_j} denote the traffic exchanged between x and y_i and y_j respectively. If x peers with y_i , then traffic from both y_i and y_j flows over the x – y_i peering link as shown in Figure 2. The peering cost for x with the single peering link is given by Equation 11. If x peers with both y_i and y_j then traffic from both ASes traverses separate peering links as shown in Figure 3. The peering cost for x with both peering links is given by Equation 12 while Equation 13 shows that the cost under the two scenarios is different.

$$PC(x) = \alpha \times (V_{xy_i} + V_{xy_j})^\beta \quad (11)$$

$$PC(x) = \alpha \times (V_{xy_i}^\beta + V_{xy_j}^\beta) \quad (12)$$


$$\because (V_{xy_i} + V_{xy_j})^\beta < V_{xy_i}^\beta + V_{xy_j}^\beta \quad (13)$$

Hence the question is: how should x go about finding its peers from among potential peers? In order to find the optimal set of peers, x would have to evaluate each combination of peers. The number of such combinations is given by:

$$\sum_{k=1}^n \binom{n}{k} \quad (14)$$


Note that each combination would involve evaluation of k peering links by both agents involved. Given the large co-location density in many regions of the interdomain network, the number of evaluations required to determine the optimal set of peers may be infeasible.

CBA heuristic CBA involves computation of all traffic flows, costs and revenue by x . Hence, employing this strategy to evaluate all potential and existing peers is not scalable.



We propose a simple heuristic to alleviate the scalability issues with CBA. Although our proposed heuristic does not compute the optimal set of peers, yet it allows agents to prioritize peers based on peering traffic aggregation and avoid redundant evaluations.

Our heuristic exploits a key feature of the *prefer-customer-then-peer-then-provider* routing policy employed in the Internet. It is based on the notion that if an AS x can reach an AS y through an existing peering link L_{xz} with another AS z , then the peering link L_{xy} would be redundant. Such a link would not save x any transit costs and would only add to its peering costs. Thus, if x finds that it has a peering path to y , it can refuse to peer with y



without actually evaluating it. Note that if x is able to reach y through an existing peering link, then it implies that y is in the customer tree of an existing peer of x .

This heuristic is explained in Figure 4. The figure shows x and z as peers while y is in the customer tree of z . In Figure 4a, x is not a peer of y . In this network, traffic from x to y flows over the peering link L_{xz} . Thus, x is able to reach y over an existing peering link. If V_{xz} and V_{xy} is the total traffic volume exchanged between x and z , and x and y respectively, then the peering cost of x is given by:

$$PC(x) = \alpha \times (V_{xz} + V_{xy})^\beta \quad (15)$$

In Figure 4b, x peers with z . Thus, traffic between x and z flows over the direct peering link and now the peering cost of x is given by:

$$PC'(x) = \alpha \times (V_{xz})^\beta + \alpha \times (V_{xy})^\beta \quad (16)$$

Note that the peering cost of x is increased because of the additional peering link.

$$PC(x) < PC'(x) \because (V_{xz} + V_{xy})^\beta < (V_{xz})^\beta + (V_{xy})^\beta \quad (17)$$

Thus, if an agent x can reach another agent y through an existing peer higher up in the network hierarchy, it does not need to establish a redundant peering link with y . The redundant peering link does not reduce upstream transit costs; on the other hand it adds to the cost of peering by reducing the advantage of economies-of-scale. Therefore, x can decide not to peer with y without actually doing cost-benefit-analysis. Algorithms 4 and 5 give the procedures for depeering and peering respectively using CBA.

Algorithm 4 Depeering by agent i using CBA

Input $PP(i) = \{x_1, x_2, \dots, x_n\}$
Sort $PP(i)$ hierarchically $PP(i) = \{x_1, x_2, \dots, x_t, x_{t+1}, \dots, x_z\} : x_t \notin C(x_{t+1})$
Do $PP(i) \leftarrow PP(i) - x_k : x_k \in C(x_j), j < k$ {Remove from $PP(i)$ all peers which are in customer trees of other existing peers}
Do CBA for $\forall x_j \in PP(i)$
 Do $F_{original} \leftarrow F(i)$
 Do Remove peering link between i and x_j
 Do Route traffic

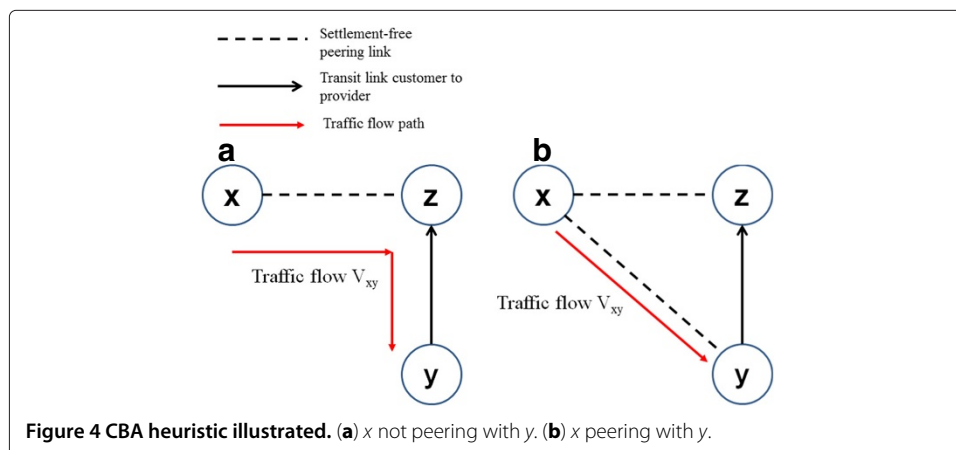


Figure 4 CBA heuristic illustrated. (a) x not peering with y . (b) x peering with y .

```

Do Compute  $F(i)$ 
     $F_{update} \leftarrow F(i)$ 
If  $F_{update} > F_{original}$ 
    Do  $PP(i) \leftarrow PP(i) - x_j$ 

```

Algorithm 5 Peering by agent i using CBA

```

Do  $PP'(i) \leftarrow j \forall j \notin C(k) \forall k \in PP(i)$ 
Do  $PP'(i) = PP'_p(i) \cup PP'_s(i) : C(j) = \emptyset \forall j \in PP'_p(i), C(j) \neq \emptyset \forall j \in PP'_s(i),$ 
     $PP_p(i) \cap PP'_s(i) = \emptyset$ 
Sort  $PP'_p(i)$  hierarchically  $PP'_p(i) = \{x_1, x_2, \dots, x_t, x_{t+1}, \dots, x_z\} : x_t \notin C(x_{t+1})$ 
Do CBA  $\forall x_j \in PP'_p(i)$ 
    Do  $F_{original} \leftarrow F(i)$ 
    Do Establish peering link between  $i$  and  $x_j$ 
    Do Route traffic
    Do Compute  $F(i)$ 
         $F_{update} \leftarrow F(i)$ 
    If  $F_{update} > F_{original}$ 
        Do  $PP(i) \leftarrow PP(i) + x_j$ 
        Do  $PP'_s(i) \leftarrow PP'_s(i) - x_k \forall x_k \in C(x_j), x_k \in PP_p(i)$ 
Do CBA  $\forall x_j \in PP'_s(i)$ 
    Do  $F_{original} \leftarrow F(i)$ 
    Do Establish peering link between  $i$  and  $x_j$ 
    Do Route traffic
    Do Compute  $F(i)$ 
         $F_{update} \leftarrow F(i)$ 
    If  $F_{update} > F_{original}$ 
        Do  $PP(i) \leftarrow PP(i) + x_j$ 

```

Open

An agent that uses this strategy agrees to peer with any other agent that is co-located with it, with the exception of its direct customers. In GENESIS-CBA, the Open peering strategy is followed by stubs. For stubs, Open peering does not always lead to the same network connectivity as CBA. However, since most stubs have limited geographic expanse, i.e., they are co-located with only few agents and thus have limited connectivity options, the difference in connectivity using either peering strategy is often small. Additionally, stubs aim to reduce their transit costs by peering with as many other agents as possible as they are not concerned with transit revenue. Thus, we improve the scalability of our model by limiting CBA adoption to transit providers only.

Parameterization of GENESIS-CBA

In the previous section we described the various components of GENESIS-CBA, without mentioning any parameter values. Next, we parameterize GENESIS-CBA to replicate the current state of the Internet as closely as possible. The parameter values

of GENESIS-CBA are shown in Table 2, along with a brief explanation for each parameter. We refer to the *instance* of GENESIS-CBA with these parameter values as the *Default model*.

Geographic distribution

The *world* in the Default model consists of 50 locations. The maximum number of locations in which an agent can be present is 15. We parameterize the Default model such that the number of locations at which agents are present follows a heavy-tailed distribution. Only a few agents are present in a large number of locations; most agents are limited to one location only. This parameterization is based on data collected from PeeringDB (2012), an online database where ASes voluntarily contribute information about themselves.

Traffic matrix

The Default model employs heavy-tailed distributions for generated and consumed traffic as seen in the real world. We parameterize the traffic matrix based on real-world

Table 2 Input parameters for default model

Parameter, symbol, description	Value	Explanation
Number of ASes N	500	Limited by 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 (PeeringDB 2012)
Geographic expanse distribution	Zipf(1.6)	Based on data about number of participants at each IXP collected from PeeringDB (PeeringDB 2012). $G(x)$ 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 (Chang et al. 2005; Feldmann et al. 2004) & Labovitz et al. (2010), 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.
Mean consumed traffic	500 Mbps	$V_C(x) \propto G(x) $, rationale being that a node with large expanse will also have a large number of access customers
Private peering threshold Ω	50 Mbps	Survey of peering strategies (Peering Strategy Survey 2011)
Transit cost multiplier range $P(x)$	[\$35,45]/Mbps per iteration	Parameterized based on IP transit prices recorded by Telegeography (Telegeography 2012). $P(x)$ assigned randomly
Transit cost exponent τ	0.75	Parameterized based on data from Chang et al. (2006) and Peering Strategy Survey (2011)
Peering cost multiplier α	\$20/Mbps per iteration	
Peering cost exponent β	0.40	
Selective peering ratio σ	2.0	

observations reported in (Chang et al. 2005; Feldmann et al. 2004; Labovitz et al. 2010). With this parameterization, 0.1% of all agents generate 28% of all the network traffic.

Economic components

We parameterize economic components of the model, e.g., transit cost per Mbps, peering cost per Mbps etc. based on real-world data. We use transit pricing data from Telegeography (2012) and peering cost per Mbps data from Amsterdam IXP (2012).

Results and discussion

We implemented GENESIS-CBA as a stand-alone single-threaded tool using C++. Our implementation did not use any existing network simulation or modeling tool. The results presented subsequently all correspond to this implementation. However, as stated earlier, GENESIS-CBA is modular and extensible and can easily incorporate new methods and policies.

Validation

The validation of models such as GENESIS-CBA is inherently difficult due to the lack of available data about the economics of ISPs, the interdomain traffic matrix, and the resulting 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-CBA, examining whether it produces certain well-known 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

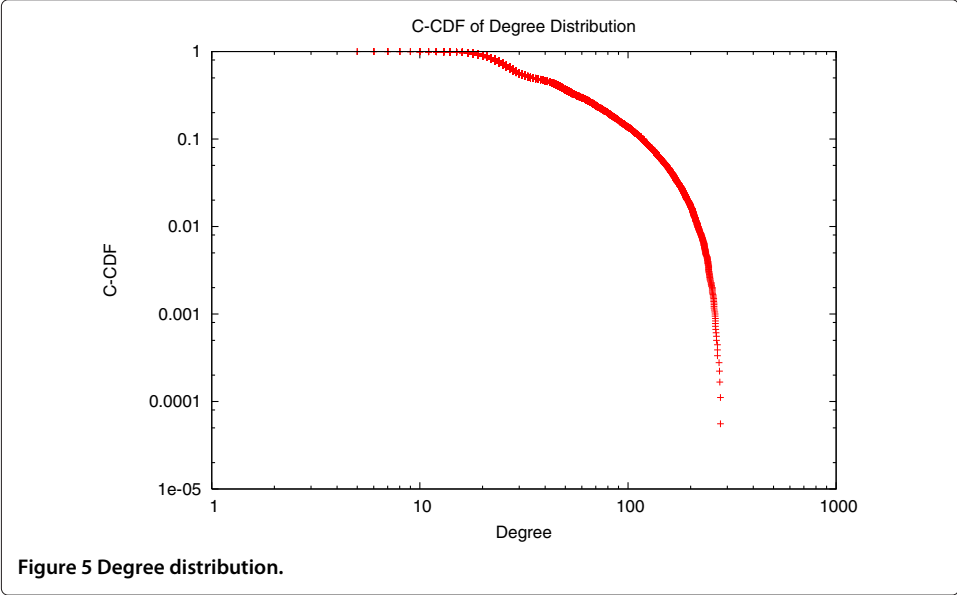
The average path length of the AS-level Internet has remained almost constant at 4 AS-hops, at least during the last twelve years (Dhamdhere and Dovrolis 2011). GENESIS-CBA produces networks with approximately the same average path length (4.4 AS-hops) with 500 agents.

Degree distribution

Figure 5 shows the complementary CDF (C-CDF) of the degree distribution that results from the default GENESIS-CBA model. Though not a strict power law (the tail of the distribution is truncated due to the limited number of agents), it is clear that the degree distribution is highly skewed, with few agents having a much larger degree than most other agents.

Link Load Distribution

Figure 6 shows the distribution of link traffic loads at equilibrium. Most links in the network are peering links (among stubs) 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. (2003) have observed a similarly skewed distribution of link loads in the Internet.



Fraction of transit providers

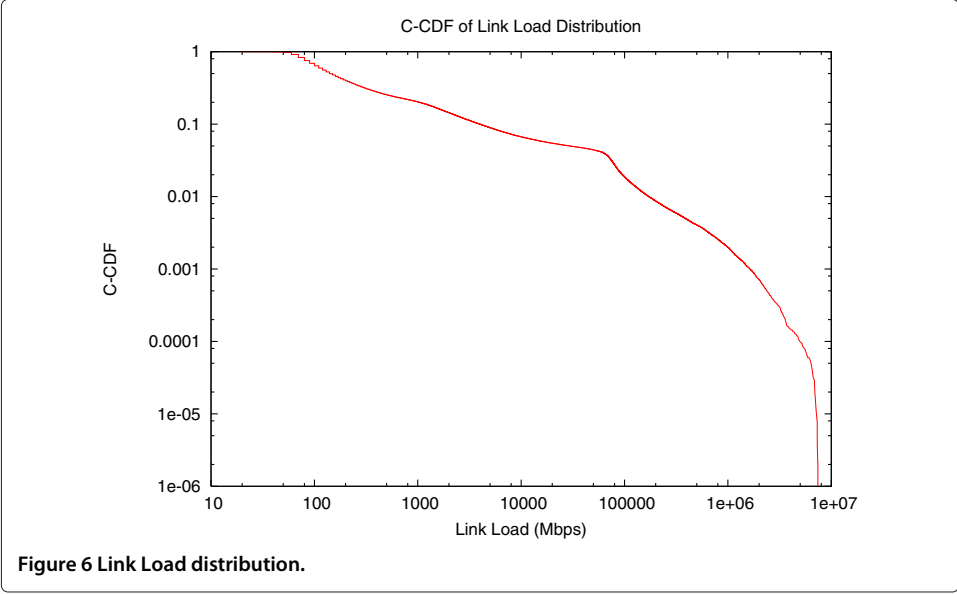
In the Default model, about 10% of the agents end up as transit providers at equilibrium, which is similar to the corresponding fraction reported in the measurement study by Dhamdhere et al. (2011).

Model characteristics

Convergence delay

We define the convergence delay as the number of iterations it takes for the sample path to converge to equilibrium.

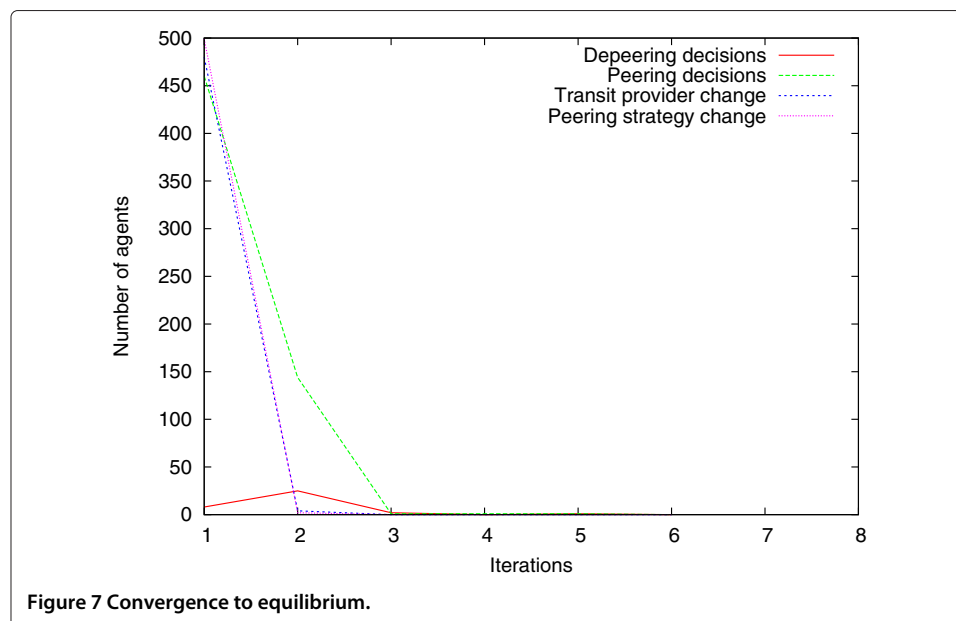
We find that the convergence delay is less than 10 iterations for all sample paths that converge to equilibrium. We investigate why the network converges to equilibrium




so rapidly despite the fact that agents act asynchronously. Figure 7 shows the number of agents that acquire new peers, remove existing peers, change transit providers, or change peering strategies over time for a typical sample path of the model. The figure shows a rapid decline in the number of agents that switch transit providers and peering strategies (between CBA and Open). We find that the network converges to stability in a small number of iterations due to the rapid convergence of stubs to their preferred transit providers. As agents select the cheapest eligible transit provider, the number of transit providers in the network reduces significantly in the first few iterations. Most agents become “stubs” (agents without transit customers), and converge to the cheapest transit provider and peering strategy. The dynamics of transit provider selection in the model dictate that stubs never change their transit providers and cannot themselves become transit providers at a later time. As a result, their peering strategies do not change over time. The few remaining transit providers accumulate large transit volumes and compete for the remaining customers. As a result, some existing transit providers also become stubs. This effect reduces the number of transit providers, enhancing stability in the network.

Effect of population size

We find that the behavior of the model does not change significantly as the population size increases over a fixed number of geographic locations. Figure 8 shows this effect. Initially, with a small number of agents, the population density across all geographic locations is low. At this point, however, more transit providers are required to maintain connectivity between all agents across all locations. As the population size increases, geographic locations become saturated with agents providing transit at the same price, causing the number of transit providers to stabilize (agents always choose the cheapest eligible transit provider). The stability in the number of transit providers leads to stable network characteristics, e.g., the average path length. Similarly, for reasons described above, the

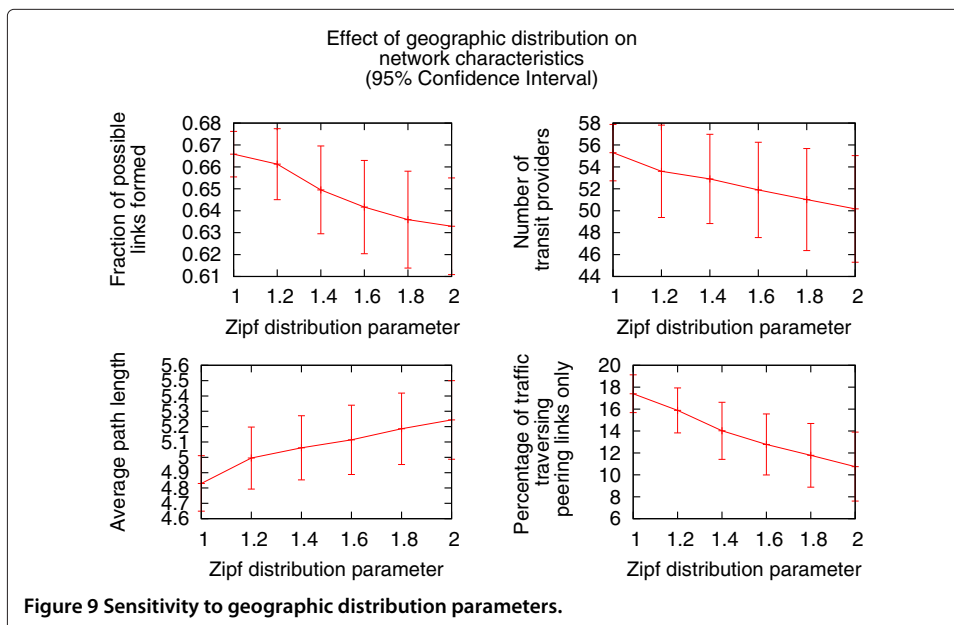




convergence behavior of the model also does not change significantly with increasing population size.

Effect of geographic and traffic distribution

Our investigation reveals that co-location among agents has the most significant effect on the resulting networks. We vary the Zipf parameter of the geographic distribution to determine the effect on the resulting networks. An increase in the parameter value makes the geographic distribution more skewed. Figure 9 shows the effect of geographic distribution on network characteristics. We find that a less skewed distribution leads to increased co-location density (i.e., more agents are co-located with one another) resulting



in a larger number of transit providers, greater connectivity, and reduced average path lengths. We also find that as a result of increased connectivity more traffic bypasses transit links and instead reaches its destination by traversing only peering links.

Stability of GENESIS-CBA

The network formation process in GENESIS-CBA is deterministic (the parameters of the population and the playing order does not change during a sample path), and so it can have one of two possible outcomes: 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 agents changes its connections or peering strategy during an iteration. On the other hand, a limit cycle (or oscillation) 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.

We ran many simulations with distinct node populations to evaluate how often the network formation process results in oscillations in practice. We observed that oscillations take place in about 7% 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-CBA is a good model for the Internet, we should not expect that the interdomain connectivity in the latter will be stable. We return to this point at the end of this section.

Why 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 study equilibria in a system that is constantly evolving?

The fact that GENESIS-CBA leads to an equilibrium, at least in 93% 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.

The study of equilibria in the context of cost-benefit-analysis enables us to study the situation where each agent, by using cost-benefit analysis, has reached a situation where it has no further incentive to change its connectivity. Additionally, the study of equilibria gives us insight into the behavior of the model: what factors, e.g., initial topology, playing order of agents etc., influence the formation of the network for a given population setting; which factors have a greater impact on the network formation process; are the networks formed for different populations with the same statistical features qualitatively the same or different etc.

Oscillations

Causes of oscillations We analyzed many oscillatory sample paths, attempting to identify a common cause behind all of them. We find that oscillations occur primarily owing

to the inability of agents using CBA to forecast with complete accuracy the network-wide effects of their actions. Agents using CBA can predict the immediate effect of formation (or termination) of a peering link. However, they cannot predict how other agents in the network will react to their actions. During the course of network formation a certain agent may be forced to revise its decisions multiple times, given the reactions of other agents. Thus, GENESIS-CBA produces a complex co-evolutionary network where limited foresight of agents may lead to complex oscillatory patterns.

We find that oscillations are limited to the tier-1 clique owing to the CBA heuristic conditions of not peering with a peer's customers. We explain this pattern with a simple example shown in Figure 10.

In Figure 10a, x and y are two tier-1 transit providers peering with each other using CBA. x and y have customers a and b respectively. x and y do not peer with each other's customers as the CBA scheme prohibits such links. Consider the following sequence of actions:

1. x plays and evaluates its peering link with y using CBA. It does so by hypothetically depeering y and recomputing its fitness.
2. Assume that x satisfies the transit provider selection criteria for y . y chooses x as its transit provider after being depeered as shown in Figure 10b.
3. x recomputes its fitness after depeering y and finds that its fitness has increased. As y switched from peer of x to its customer, x 's transit revenue increased and its peering cost decreased. x therefore decides to depeer y .
4. Since y is no longer a peer of x , it can peer with the customers of x . y evaluates peering with a . Since y can save transit costs by peering with a , its fitness increases. Thus, y decides to peer with a as shown in Figure 10c.
5. Note that the y — a peering link decreases the transit traffic of x , as traffic $a \longleftrightarrow b$ and $a \longleftrightarrow y$ does not transit through x any more. As a result of the decreased transit traffic, x no longer satisfies the provider selection criteria for y . Since GENESIS-CBA does not allow a partitioned network, a peering-by-necessity relationship is established between x and y . Once y peers with x , it no longer needs to peer with a , a customer of x . This brings us back to the initial network of Figure 10a.

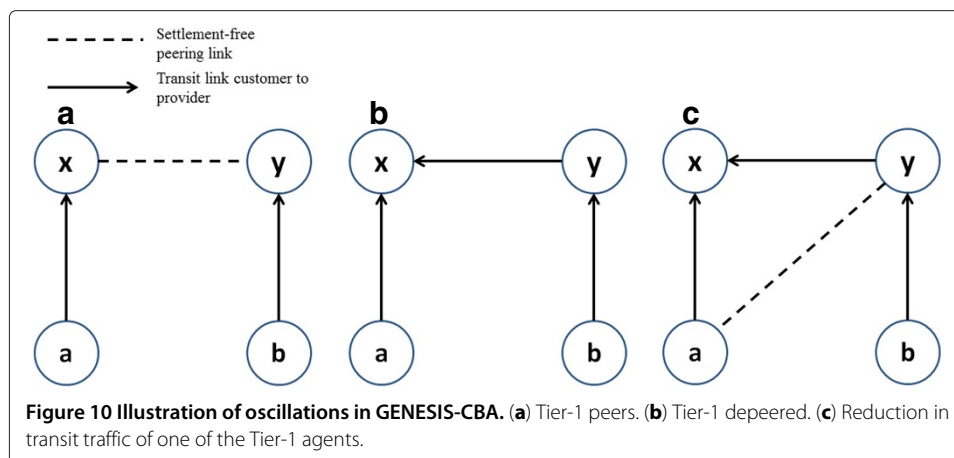


Figure 10 Illustration of oscillations in GENESIS-CBA. (a) Tier-1 peers. (b) Tier-1 depeered. (c) Reduction in transit traffic of one of the Tier-1 agents.


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

Network-wide effects of oscillations We next study the effect of oscillations on the entire network as well as the affected agents. 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:

$$J_{AB} = \frac{|L_A \cap L_B|}{|L_A \cup L_B|}$$

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 11 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 agents that undergo any change (connectivity and/or strategy) during the oscillation is less than 5% of the total network population. Note that all agents involved in oscillations are not transit providers; some agents 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 oscillatory sample paths have a limit cycle length that is less than 15 iterations. This length is positively correlated with the number of affected ASes (Pearson correlation coefficient: 0.94). Thus, the more agents 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 one another.

The fact that GENESIS-CBA can produce oscillatory sample paths should not be considered an artifact of the model. As described earlier, oscillations in GENESIS-CBA 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-CBA 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.

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 agents and an initial topology, does GENESIS-CBA 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 agent at equilibrium? Are there any global properties that are predictable with statistical significance? And who are the most predictable, or least predictable, agents in a network, in terms of fitness?

Distinct equilibria

The network that emerges from a GENESIS-CBA sample path highly depends on both the initial topology and the playing order. Specifically, given the same population of agents and the same initial topology, 90% 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 95%. So, if GENESIS-CBA 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 12 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 0.8 when we only vary the playing order. When we also vary the initial topology, the similarity metric is lower (between 0.65–0.8) but still high in absolute terms. So, even though there are many distinct equilibria, they are quite similar in terms of topology.

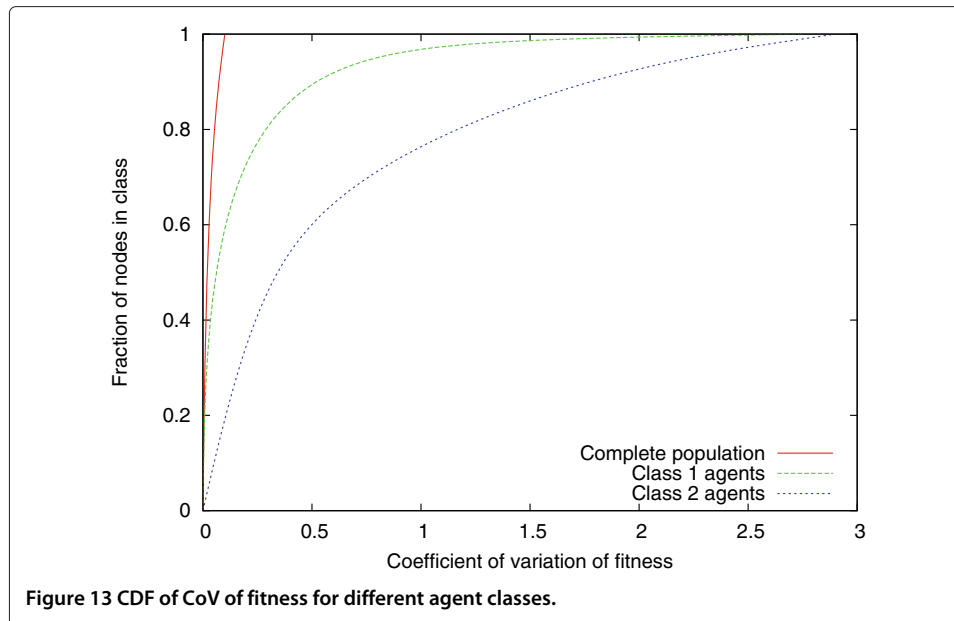
Figure 12 also shows the distribution of Jaccard similarity for the subset of the network that includes only transit providers (the “provider network”). In this case, the Jaccard similarity is significantly lower, with 45% of the equilibrium pairs being less than 55% similar. In other words, even though 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 a network. To determine the similarity of fitness distributions across different equilibria, we applied the two-sample Kolmogorov-Smirnov (KS) test on fitness distributions 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 agents’ fitness remains invariant. However, this does not mean that the fitness of an individual agent 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 agents across different equilibria. We compute the Coefficient of Variation (CoV) of the fitness of each agent across 100 equilibria resulting from different playing orders. Figure 13 shows the CDF of CoV of fitness for different agent classes. The CoV is close to zero for 90% of agents



indicating that most agents see almost identical fitness in all equilibria. However, 5% of the agents have CoV greater than 1.0 and 1% of the agents have CoV as high as 3.0. Who are these “most unpredictable” agents and what causes the large variability in their fitness?

We classify agents in two classes. *Class-1* agents are those that are either stubs at all equilibria, or that are transit providers at all equilibria. *Class-2* agents, on the other hand, are those that are transit providers in some equilibria and stubs in others. Figure 13 also shows the CoV distribution for these two classes of agents. Note that the CoV of Class-1 agents is close to zero for almost all agents, while the CoV of Class-2 agents is significant for most of them. In the following sections, we discuss the properties of the agents in each class in more detail.

Properties of Class-1 nodes The largest fraction of agents in this class are stubs (including agents 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 agents do not have revenues and their provider is always the cheapest transit provider 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 agents 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 an agent 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, the identity of the agents that will strongly benefit from this positive feedback depends on the order in which the agents play. In other words, the agents in Class-2 are mostly those 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 agents that exhibit the largest variation in fitness and that act as the biggest source of variability in the resulting networks.

Conclusions

We proposed GENESIS-CBA, 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-CBA implements a new peering scheme that can replace existing peering strategies. The new scheme allows agents to carry out economic analysis for each peering link. Agents can use this analysis to form only those peering links that improve their economic fitness. We examined the convergence properties of GENESIS-CBA, and found that GENESIS results in a stable network in most cases. The observed cases of instability occur due to connectivity constraints in the Tier-1 clique. GENESIS-CBA 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. In ongoing work, we use GENESIS-CBA to understand the effect of the scheme on the fitness distribution of different classes of autonomous systems in the Internet.

In future, we intend to extend GENESIS-CBA to turn it into a practical tool capable of being deployed by the operators managing the ASes in the interdomain network. Our objective is to align the information requirements of CBA with the information available to ASes in the real world. CBA uses the following information for analysis: topological information, e.g., customer-provider relationships, traffic flows between ASes and geographic co-location. While ASes can form estimates of the previously mentioned information from data available to them, their accuracy in the context of CBA is yet to be ascertained, and this is one of the goals of our future work. We envision GENESIS-CBA being deployed in the real world at a microscopic scale, i.e., at the AS level, to carry out what-if analysis of peering decisions. In this context, CBA strategy can be deployed by ASes in the Internet to aid in their analysis to determine which ASes to peer with. They can use CBA to analyze the effect of existing and potential peering links on their economic fitness.

Endnote

¹There is another type of peering which is known as “paid peering” in which one peer pays the other. However, “paid peering” is not the focus of this work.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

AD and CD supervised the design of the model and the study. The development work for the model and simulations were carried out by AL. All authors actively participated in data analysis and editing of the draft. All authors read and approved the final draft.

Acknowledgements

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.

Author details

¹School of Computer Science, Georgia Institute of Technology, Atlanta, USA. ²The Cooperative Association for Internet Data Analysis, University of California, San Diego, USA.

Received: 3 December 2012 Accepted: 22 March 2013

Published: 24 April 2013

References

- Akella, A, Seshan S, Shaikh A: **An empirical evaluation of wide-area internet bottlenecks.** In *Proceedings of ACM SIGCOMM IMC*; 2003.
- Amsterdam Internet Exchange Pricing** 2012. [<https://www.ams-ix.net/services-pricing/pricing>]
- Augustin, B, Krishnamurthy B, Willinger W: **IXPs: mapped?** In *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement Conference, IMC '09*. Chicago: ACM; 2009.
- Chang, H, Jamin S, Mao Z, Willinger W: **An empirical approach to modeling Inter-AS traffic matrices.** In *Proceedings of ACM SIGCOMM IMC*; 2005.
- Chang, H, Jamin S, Willinger W: **Internet connectivity at the AS-level: an optimization-driven modeling approach.** In *Proceedings of ACM SIGCOMM MoMeTools Workshop*; 2003.
- Chang, H, Jamin S, Willinger W: **To peer or not to peer: modeling the evolution of the internet's AS-level topology.** In *Proceedings of IEEE INFOCOM*; 2006.
- Corbo, J, Jain S, Mitzenmacher M, Parkes DC: **An Economically-Principled Generative Model of AS Graph Connectivity.** In *Proceedings of IEEE INFOCOM Mini-Conference*; 2009.
- Dhamdhere, A, Dovrolis C: **The internet is flat: modeling the transition from a transit hierarchy to a peering mesh.** In *Proceedings of ACM CoNEXT*; 2010.
- Dhamdhere, A, Dovrolis C: **Twelve years in the evolution of the internet ecosystem.** *IEEE/ACM Trans Netw* 2011, **19**(5):1420–1433.
- DE-CIX: **Frankfurt internet exchange point global peering pricing.** [<http://www.de-cix.net/products-services/de-cix-frankfurt/globepeer/>]
- Fabrikant, A, Koutsoupias E, Papadimitriou CH: **Heuristically optimized trade-offs: a new paradigm for power laws in the internet.** In *Proceedings of ICALP*; 2002.
- Faloutsos, M, Faloutsos P, Faloutsos C: **On power-law relationships of the internet topology.** In *Proceedings of ACM SIGCOMM*; 1999.
- Feldmann, A, Kammenhuber N, Maennel O, Maggs B, De Prisco R, Sundaram R: **A methodology for estimating interdomain web traffic demand.** In *Proceedings of ACM SIGCOMM IMC*; 2004.
- Gill, P, Arlitt M, Li Z, Mahanti A: **The flattening internet topology: natural evolution, unsightly barnacles or contrived collapse?** In *Proceedings of the Passive and Active Measurement Conference (PAM)*; 2008.
- Goyal, S: *Connections: An Introduction to the Economics of Networks*. New Jersey: Princeton University Press; 2007.
- Holme, P, Karlin J, Forrest S: **An integrated model of traffic, geography and economy in the internet.** *ACM SIGCOMM CCR* 2008.
- Internet Assigned Numbers Authority** 2012. [www.iana.org]
- Internet Transit Prices: Historical and Projected** 2010. [<http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php>]
- Jackson, MO: *Social and Economic Networks*. New Jersey: Princeton University Press; 2008.
- Labovitz, C, Iekel-Johnson S, McPherson D, Oberheide J, Jahanian F: **Internet inter-domain traffic.** In *Proceedings of ACM SIGCOMM*; 2010.
- Li, L, Alderson D, Doyle JC, Willinger W: **Towards a theory of scale-free graphs: definition, properties, and implications.** *Internet Math* 2005, **2**(4).
- Lodhi, A, Dovrolis C: **A network formation model for internet transit relations.** In *Proceedings of the Workshop on Economics of 2010 Networks, Systems, and Computation (NetEcon)*; 2010.
- Lodhi, A, Dhamdhere A, Dovrolis C: **Analysis of peering strategy adoption by transit providers in the internet.** In *Proceedings of the Workshop on the Economics of Networks, Systems, and Computation*; 2012a.
- Lodhi, A, Dhamdhere A, Dovrolis C: **GENESIS: An agent-based model of interdomain network formation, traffic flow and economics.** In *Proceedings of IEEE INFOCOM*; 2012b.
- Lodhi, A, Dhamdhere A, Dovrolis C: **Peering strategy adoption by transit providers in the internet: a game theoretic approach?** *SIGMETRICS Perform Eval Rev* 2012c, **40**(2):38–41.
- Norton, W: **The great (public vs. private) peering debate: peering at 10Gbps.** 2011.
- Peering Strategy Survey** 2011. [<http://drpeering.net/white-papers/Peering-Policies/A-Study-of-28-Peering-Policies.html>]
- PeeringDB** 2012. [<http://www.peeringdb.com>]

Telegeography: **IP tranist pricing service**. 2012. [http://www.telegeography.com]

Wang, X, Loguinov D: **Wealth-based evolution model for the internet AS-Level topology**. In *Proceedings of IEEE INFOCOM*; 2006.

Wang X, Loguinov D: **Understanding and modeling the internet topology economics and evolution perspective**. *IEEE/ACM Trans Netw* 2010, **18**(1):257–270.

Yook, SH, Jeong H, Barabasi AL: **Modeling the internet's large-scale topology**. *Proc Natl Acad Sci* 2002, **99**(21):13382–13386.

doi:10.1186/2194-3206-1-10

Cite this article as: Lodhi et al.: GENESIS-CBA: an agent-based model of peer evaluation and selection in the internet interdomain network. *Complex Adaptive Systems Modeling* 2013 **1**:10.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Immediate publication on acceptance
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ springeropen.com
