# Open Peering by Internet Transit Providers: Peer Preference or Peer Pressure?

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*—**Peering agreements between Autonomous Systems affect not only the flow of interdomain traffic but also the economics of the entire Internet ecosystem. The conventional wisdom is that transit providers are selective in choosing their settlement-free peers because they prefer to offer revenue-generating transit service to others. Surprisingly, however, a large percentage of transit providers use an *Open* peering strategy. What causes this large-scale adoption of *Open* peering, especially among transit providers? More importantly, what is the impact of this peering trend on the economic performance of the population of transit providers? We approach these questions through game-theoretic modeling and agent-based simulations, capturing the dynamics of peering strategy adoption, inter-network formation and inter-domain traffic flow. We explain why transit providers gravitate towards *Open* peering even though that move may be detrimental to their economic fitness. Finally, we examine the impact of an *Open* peering variant that requires some coordination among providers.**

*Keywords*: Internet, Autonomous System interconnections, settlement-free peering, peering strategies, economic utility

## I. INTRODUCTION

The Internet is a complex and dynamic interconnection of more than 50000 Autonomous Systems (ASes). A link in this context represents a business agreement between two ASes that can exchange traffic under various policy constraints. Some links represent *transit* relations, in which one AS acts as the *provider* of the other AS (the *customer*), offering the latter access to the entire Internet. Other links represent *settlement-free peering* (or simply "peering") relations, in which two ASes exchange their local and customer-originating traffic at no charge. These business relations have a heavy economic impact, affecting the profitability of the ASes.

The economic objective of peering is to reduce upstream transit costs. To a large degree, each AS X follows a *peering strategy* (or "peering policy"[1]) that is used to determine whether X will accept to peer with another AS Y. Even though they can vary widely in their details, most peering strategies can be grouped in three distinct classes: *Restrictive* (X peers only if necessary to avoid Internet partitioning; typically used by Tier-1 transit providers), *Selective* (X peers only with ASes

that are comparable with X, a notion that we will define more precisely in Section III), and *Open* (X is willing to peer with everyone, except its customers). The conventional wisdom is that transit providers use *Restrictive* or *Selective* peering, so that they can engage other ASes as their customers, thus increasing their transit revenues.

In the last few years, however, there is evidence that the Internet peering ecosystem is going through a major transformation from *more restrictive or selective* to *more open* interconnection. A recent study on peering interconnectivity at a large European Internet Exchange Point (IXP) reveals a "rich peering fabric", with 67% of all possible peering relationships established at the IXP [1]. Such interconnectivity characteristics can only arise if a large percentage of ASes present at the IXP engages in Open peering. We have analyzed data from PeeringDB [2], an online database where peering coordinators provide information about their ASes. This data shows that most transit and access providers (70-80%) use the Open peering strategy. This is counterintuitive, especially for transit providers, because if they peer openly how can they attract new customers or keep their existing ones? More surprisingly, this trend is not limited to small transit providers. The PeeringDB dataset reveals that 36% of transit providers with traffic volume greater than 100 Gbps (e.g., Hurricane Electric), 37% of providers with traffic volume between 50 and 100 Gbps (e.g., WIND Telecomunicazioni S.P.A.) and 66% of providers with global scope (e.g., DeltaTelecom) use *Open* peering. Furthermore, the findings of a recent survey show that 99.5% of peering relationships are formed without formal analysis or agreements [3].

These observations on peering behavior raise some important questions: *Why do so many transit providers use Open peering? What are the underlying interconnectivity dynamics that influence their peering decisions? What does this attraction towards Open peering imply for the economic performance of transit providers? Is the inclination towards Open peering uniform across all categories of transit providers?* In the absence of any formal analysis for most peering agreements, the network-wide economic impact of peering on transit providers is not well understood. Furthermore, Internet providers, independent of their type, are secretive about their economic objectives and operational data. Hence, we cannot address the previously mentioned questions empirically. Instead, we rely on analytical and computational modeling. Our work explores the peering behavior of transit providers and challenges the conventional notion that peering

[1]We use these two terms interchangeably.

is always economically beneficial. We show that myopic and selfish adoption of peering strategies, and lack of coordination among transit providers can push transit providers towards Open peering. We also explain why transit providers may find it difficult to break out of the stable, but suboptimal, equilibria formed through Open peering. Finally, we propose a simple coordination mechanism that transit providers can incorporate in their peering strategies to alleviate the economic loss associated with Open peering while retaining its benefits.

The rest of this paper is structured as follows. We briefly discuss related work in Section II. In Section III, we introduce a simplified internetwork model to analyze the gravitation towards Open peering. In Section IV, we use large-scale simulations with a computational agent-based model to analyze peering behavior in realistic settings. We present our proposed Open peering variant in section V and conclude in Section VI.

## II. RELATED WORK

Our work is most closely related to models of interdomain network formation that incorporate economic principles. Such models capture the decentralized and asynchronous processes through which ASes make peering decisions. Chang et al. propose a model in which ASes select transit providers and settlement-free peers based on economic factors and other constraints [4]. Unlike our work, their model exogenously assigns provider and peer selection strategies to each AS. Holme et al. propose an agent-based network growth model [5]. However, they treat link formation as a random process, and do not capture link rewiring, settlement-free peering, and realistic Internet routing. Corbo et al. model transit link formation and rewiring based on the economic utility that nodes derive from interconnection [6]. Their work focuses on creating well-performing internetworks, where the network's performance is the total utility of all nodes; they do not capture settlement-free peering. Dhamdhere et al. [7] propose the ITER network formation model to study the evolutionary transition of the Internet from a *hierarchical* to a *flat* structure. In ITER, node types and peering strategies are pre-assigned, whereas in our work they are an outcome of the model. Also, our goal is to study the peering policies adopted by Internet transit providers. Lodhi et al. introduced the GENESIS model for peering strategy adoption by transit providers [8]. Nodes in that model do not select their peering policy; the latter is determined by the node's hierarchical status. Further, the main focus of Lodhi et al. [8] was to examine the existence and heterogeneity of equilibria in GENESIS. An extensive line of research has focused on *topology generation*, i.e., models that create interdomain topologies matching certain statistical properties of the AS-level Internet. Rule-based generative models [9], [10], [11] aim to reproduce topological characteristics of the Internet, but do not capture the underlying economics. Another line of work [12], [13] takes the "bottom-up" approach, modeling the optimization objectives and constraints of ASes to produce realistic interdomain topologies. Our work does not focus on topology generation, but instead on the adoption of peering policies by transit providers, and the effect of those policies on the economics of transit providers. There is much prior work on game theoretic analysis of transit and peering relations [14], [15], [16], [17], [18], [19]. For reasons of mathematical tractability, however, these works often ignore
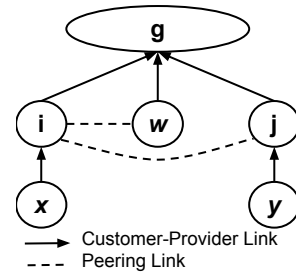


Fig. 1: Network model with the fixed customer-provider links and some potential peering links

some key real-world factors, e.g., transit vs peering links, policy-based routing, economies of scale, etc.

## III. ANALYTICAL RESULTS

In this section, we consider analytically a simple network model to investigate the conditions and reasons that lead transit providers to adopt Open peering. We show that Open peering results from myopic and uncoordinated decisions made by transit providers, and it can lead to a loss of economic utility depending on certain traffic conditions.

### A. Model Description

The network model, shown in Figure 1, consists of five nodes: $i$, $j$, $g$, $w$, $x$ and $y$. The nodes represent ASes that seek to optimize their utility by choosing the best peering strategy. $w$, $x$ and $y$ are *stubs* i.e., they do not have transit customers; they openly peer with any node that is willing to peer with them to minimize their transit costs. Node $g$ is a tier-1 transit provider that uses the Restrictive peering policy, and so it does not peer with any other node. Its presence in the model is only to ensure that there is at least one path between any two nodes even if there are no peering links. $i$ and $j$ are *transit providers* that choose dynamically the peering strategy that maximizes their economic utility; these two nodes are the only players in the following repeated game.

**Link formation and traffic routing:** Nodes interconnect through one of two types of links: (a) customer-provider or transit links, and (b) peering links. Players $i$ and $j$ are customers of $g$; $w$, $x$ and $y$ are customers of $g$, $i$ and $j$ respectively. The customer-provider links remain fixed during the following game. However, peering links change as $i$ and $j$ adopt different peering policies, as explained later. We denote the link between any two nodes $p$ and $q$ by $L_{pq}$.

Only geographically co-located nodes can form links. The stubs $w$, $x$ and $y$ are not co-located. All other players are co-located and can form links with each other.

**Traffic and routing:** The traffic flow (Mbps) sent from node $p$ to $q$ is denoted by $T_{pq}$. We denote the total traffic exchanged between $p$ and $q$ by $V_{pq}$,

$$V_{pq} = T_{pq} + T_{qp} \qquad (1)$$

The total traffic of a node $p$, denoted by $\hat{V}_p$, is the sum of all traffic that $p$ exchanges with other nodes.

| $s_p$ | Peers of $i$ | | Peers of $j$ | |
|---|---|---|---|---|
| | $I_\sigma(i,w)=1$ | $I_\sigma(i,w)=0$ | $I_\sigma(j,w)=1$ | $I_\sigma(j,w)=0$ |
| $R$ | $\emptyset$ | $\emptyset$ | $\emptyset$ | $\emptyset$ |
| $S$ | $j, w$ | $j$ | $i, w$ | $i$ |
| $O$ | $j, w, y$ | $j, w, y$ | $i, w, x$ | $i, w, x$ |

TABLE I: Peers of $i$ and $j$ under different peering strategies

Traffic follows the shortest path subject to two common policy constraints in the Internet [20]: "prefer customer over peer over provider links" and the "valley-free" routing property. If two nodes cannot exchange traffic directly over a peering link, they have to rely on an upstream transit provider to carry their traffic. For instance, in Figure 1, $x$ exchanges traffic with $w$ through its transit provider $i$; however, $i$ exchanges traffic with $w$ directly over their peering link instead of going through its transit provider $g$.

**Peering strategies:** Nodes form bilateral peering relationships based on their peering strategies. $s_p$ denotes the peering policy of a node $p$. We consider the following three peering strategies, based on the policies that are announced at PeeringDB [2] and that are widely discussed at NANOG [21] and peering surveys [22]:

1) *Restrictive (R)*: Node $p$ does not peer with any other node unless it is necessary to avoid network partitioning.

2) *Selective (S)*: Node $p$ only peers with nodes that have similar (or larger) size than itself. We use the total traffic of a node as a proxy for its size. The rationale is that if a node $q$ is much smaller than $p$, in terms of total traffic, then $p$ would prefer to become a transit provider of $q$ as opposed to a peer of $q$. $I_\sigma(p,q)=1$ denotes that $q$ satisfies the Selective peering constraint of $p$, stated as follows

$$I_\sigma(p,q)=1 \iff \frac{\hat{V}_p}{\hat{V}_q} \leq \sigma \qquad (2)$$

where $\sigma$ denotes the Selective traffic ratio constraint ($\sigma \geq 1$), Otherwise, $I_\sigma(p,q)=0$.

3) *Open (O)*: Node $p$ peers with all co-located nodes except its existing customers.

Due to space constraints, we further assume the following two conditions so that the analysis focuses on the more interesting instances of the game. First, transit providers $i$ and $j$ satisfy each other's Selective peering constraint, i.e., $I_\sigma(i,j)=1$ and $I_\sigma(j,i)=1$ for the given value of $\sigma$, and so they always peer with each other. If this was not true, the providers $i$ and $j$ would be attracted to Open peering even more. Second, the traffic of stubs $x$ and $y$ are such that $I_\sigma(i,y)=0$ and $I_\sigma(j,x)=0$, i.e., they do not qualify to become peers of $j$ and $i$, respectively. If this was not true, a provider would simply use Selective peering if $w$ qualifies to become its peer, and Open otherwise. Recall that stubs, $w$, $x$ and $y$ peer openly with any player willing to peer with them.

**Costs and revenues:** Each provider $v$ charges its transit customers a price of $P_v$ \$/Mbps. For instance, if player $i$ exchanges traffic $\mathcal{V}$ with its provider $g$, $i$ incurs a transit cost[2]

---

[2]We use a linear cost and revenue model in this section. The simulation model in Section IV uses a more realistic nonlinear function that captures economies of scale.

$C_i^T$,

$$C_i^T = P_g \times \mathcal{V} \qquad (3)$$

The transit revenue of $i$, denoted by $R_i$, is the transit cost incurred by its (only) customer $x$.

Peering also incurs a cost. If $\mathcal{V}'$ is the total peering traffic of player $i$, its peering cost is given by:

$$C_i^P = \alpha \times \mathcal{V}' \qquad (4)$$

where $\alpha$ is the peering cost factor (\$/Mbps).

For example, in the network of Figure 1, the transit and peering costs of player $j$ are given by:

$$C_j^T + C_j^P = P_g \times (V_{jw} + V_{wy}) + \alpha \times (V_{ij} + V_{jx} + V_{iy} + V_{xy}) \qquad (5)$$

Similarly, $j$'s revenue is given by:

$$R_j = P_j \times (V_{jy} + V_{wy} + V_{iy} + V_{xy}) \qquad (6)$$

In practice, peering is much cheaper than transit, meaning that $\alpha \ll P_v$ for any provider $v$.

**Utility:** The peering strategies $s_i$ and $s_j$ determine the interdomain topology of the underlying network (who is peering with whom), and thus the traffic flow on each transit and peering link. These traffic volumes then determine the utility of each player. So, we express the utility of each player $p$ as a function of the peering policies of $i$ and $j$: $u_p(s_i, s_j)$, $p \in \{i, j\}$. The strategy space of each player consists of the three previous strategies $(R, S, O)$ and is denoted by $\mathcal{S}$.

The utility of player $p$ is its revenue minus its transit and peering costs,

$$u_p(s_i, s_j) = R_p - C_p^T - C_p^P \qquad (7)$$

where the revenue and cost terms are calculated as previously described.

**Repeated game:** Network formation takes place in discrete time ("rounds"), with players $i$ and $j$ playing one after the other. Without loss of generality, the game starts with $i$'s action. The initial condition of the game is the pair of peering strategies $(s_i^0, s_j^0)$ at $t=0$, which also determines the initial topology. In each round, a player computes its own utility under each of the three candidate peering strategies, and selects the peering strategy that maximizes its utility at that time. In other words, each player acts myopically based on the information that it currently has, without trying to predict the actions of other players in future rounds ("best-response dynamics") and without trying to coordinate with other players.

Specifically, when player $p$ plays at time $t$, $p$ selects the policy $s_p^t$ that maximizes its own utility, given the policy selection of the other player $(-p)$ in the previous round. Thus, $p$ solves the optimization problem:

$$\max_{s_p^t \in \mathcal{S}} u_p^t(s_p^t, s_{-p}^{t-1}) \qquad (8)$$

We assume that if two strategies give the same utility, the player will choose the strategy that results in the minimum number of peering links (a more manageable configuration in practice). In other words, we break any ties in favor of first the Restrictive and then the Selective strategy.

The game terminates if the network reaches a Nash equilibrium, meaning that none of the two players can increase its utility by unilaterally switching to a different peering strategy. So, the strategy pair $(s_p^*, s_{-p}^*)$ is an equilibrium if

$$u_p(s_p^*, s_{-p}^*) \geq u_p(s_p, s_{-p}^*) \quad \forall s_p \in \mathcal{S} \qquad (9)$$

### B. Stability and Equilibria

The initial strategy vector $(s_i^0, s_j^0)$ can take nine possible values, given that $\mathcal{S}_p = \{R, S, O\}$ and $p \in \{i, j\}$. It is easy to see that the Restrictive strategy is always dominated by the Selective and Open strategies for both $i$ and $j$. Hence, we do not show the utility under the Restrictive strategy and focus on the remaining four initial strategy vectors. For each such vector we need to consider separately whether $w$ qualifies to be a peer of each provider under the Selective strategy, i.e., whether $I_\sigma(p, w) = 1$ or not for $p \in \{i, j\}$ (four cases). So, overall there are 16 possible cases that need to be analyzed.

We have confirmed that the previous repeated game converges to a Nash equilibrium in all cases. Further, as shown in Section III-C, there are only two possible equilibria: $(S, S)$ and $(O, O)$. In other words, eventually both providers use either the Selective or the Open strategy. It is possible to explain the gravitation towards Open peering through the simple one-shot game equilibria. However, that does not capture all the cases in which the network converges to Open peering.

### C. Best-Response Dynamics

In this section, we analyze the game in detail, focusing on the sequence of peering decisions by each provider, for the most interesting of the previous 16 cases.

#### 1) $(\mathbf{s_i^0}, \mathbf{s_j^0}) = (\mathbf{S}, \mathbf{S})$, $\mathbf{I_\sigma(i, w) = 1}$, $\mathbf{I_\sigma(j, w) = 1}$:

We first consider the case that $w$ qualifies to be a peer of $i$ and $j$ under Selective peering. Player $i$ plays first and evaluates all peering strategies. Switching to Open does not increase $i$'s utility because that provider can peer with both $w$ and $j$ under the Selective strategy, and so it can reach everyone (except $g$) through its peering links. Thus,

$$u_i(S, S) - u_i(O, S) = 0 \qquad (10)$$

$j$ carries out the same analysis and also decides to stay with the Selective strategy. So the resulting equilibrium is $(S, S)$.

An example of this case is when major content providers such as Google or Facebook ($w$ in the model) have sufficiently large traffic volume to peer with transit providers that use the Selective strategy. Non-tier-1 transit providers ($i$ and $j$ in the model) peer with such content providers, resulting in lower transit costs for both the content provider and the transit provider.

#### 2) $(\mathbf{s_i^0}, \mathbf{s_j^0}) = (\mathbf{S}, \mathbf{S})$, $\mathbf{I_\sigma(i, w) = 0}$, $\mathbf{I_\sigma(j, w) = 1}$:

We now consider the case that $w$ does *not* qualify to be a peer of $i$ when $s_i = S$, but it qualifies to be a peer of $j$ when $s_j = S$.[3] In this case, $i$ finds that Open dominates the Selective strategy,

$$u_i(O, S) - u_i(S, S) = (P_g - \alpha) \times (V_{iw} + V_{wx}) \geq 0 \quad (11)$$

[3]The case that $I_\sigma(i, w) = 1$ and $I_\sigma(j, w) = 0$ is symmetric.
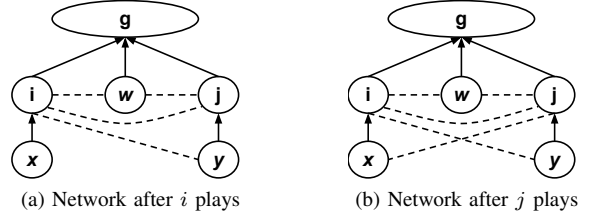


Fig. 2: Network formation in Case-2

because of the reduction in transit costs when the traffic that it exchanges with $w$ is routed through a peering link; recall that $P_g > \alpha$ (transit is more expensive than peering).

Under the Open strategy however, $i$ will also peer with $y$, the customer of $j$. The resulting network is shown in Figure 2a. Note that, as a result of the peering link between $i$ and $y$, provider $j$ now loses the revenue that was due to the traffic $V_{iy} + V_{xy}$; that traffic now bypasses $j$.

When it is $j$'s turn to play, it also finds that Open dominates Selective,

$$u_j(O, O) - u_j(O, S) = (P_j - \alpha) \times T_{xy} \geq 0 \qquad (12)$$

Note that the Open strategy is better for $j$ even though it was able to exchange traffic with $i$, $x$ and $w$ through peering links even under the Selective strategy. The reason is that if $j$ uses Open peering it can directly peer with $x$, and so the traffic from $x$ to $y$, $T_{xy}$, will be routed again through $j$, partially alleviating $j$'s earlier loss in revenue. On the other hand, the traffic flow from $y$ to $x$, $T_{yx}$ still bypasses $j$ through the peering link $L_{iy}$. Thus, $j$ also adopts the Open strategy, it peers with $x$, and the new strategy vector becomes $(O, O)$.

When the two players play again, they find that any unilateral deviation from the Open strategy would reduce their utility:

$$\begin{aligned} u_i(O, O) - u_i(S, O) &= (P_i - \alpha) \times T_{yx} &(13) \\ &+ (P_g - \alpha) \times (V_{iw} + V_{wx}) \geq 0 \end{aligned}$$

$$u_j(O, O) - u_j(O, S) = (P_j - \alpha) \times T_{xy} \geq 0 \qquad (14)$$

Thus, the network reaches the equilibrium $(O, O)$.

The net difference in $i$'s utility between the start of the game and this equilibrium is given by

$$\begin{aligned} u_i(S, S) - u_i(O, O) &= (P_i - \alpha) \times (V_{jx} + T_{xy}) &(15) \\ &- (P_g - \alpha) \times (V_{iw} + V_{wx}) \end{aligned}$$

$i$'s utility at equilibrium may be less than its initial utility. When $i$ adopts the Open strategy, it peers not only with $w$ (to reduce its upstream peering costs) but also with $j$'s customers, thereby diverting part of $j$'s customer traffic directly to $i$. Then, $j$ also adopts Open peering so that it can partially alleviate this loss in transit revenue. In doing so, $j$ diverts part of $i$'s customer traffic directly to $j$. Hence, while $i$ initially benefits from a reduction in transit costs, it finally also suffers a loss in transit revenue. The net effect on the utility of $i$ and $j$ depends on the relation between the transit prices of the involved providers and the relative size of the affected traffic flows.

*Effect of traffic volume:* Let us further discuss the role of the traffic that $i$ and its customers exchange with $w$ ($V_{iw} + V_{wx}$ in Equation 15), and the traffic that $i$'s customers exchange with $j$ and its customers ($V_{jx} + T_{xy}$ in Equation 15). The former is traffic for which $i$ saves transit costs by peering with $w$, while the latter is traffic for which $i$ loses revenue when $j$ adopts Open peering.

As a first-order approximation, we can assume that the transit price of all providers is roughly the same, perhaps due to competition (i.e., $P_i \approx P_j \approx P_g$). Then, according to Equation 15, Selective peering is preferred if $(V_{jx} + T_{xy}) \gg (V_{iw} + V_{wx})$. This situation would arise if $x$ is a much larger content provider than $w$. $i$ cannot afford to lose the transit revenue from its existing customer $x$, which is what would happen if both providers use Open peering.

If, however, $(V_{iw} + V_{wx}) \gg (V_{jx} + T_{xy})$, then switching to Open peering is preferred. Such a situation can arise if $w$ is a much larger content provider than $i$'s customers, and so $i$ would incur a large transit cost if not peering with $w$. Although $i$'s adoption of Open peering forces $j$ to adopt Open peering as well (resulting in a reduction of $i$'s revenue), this revenue loss is much smaller than the upstream transit cost that $i$ would have incurred by not peering with $w$.

Similar considerations apply to $j$.

*3)* $(\mathbf{s_i^0}, \mathbf{s_j^0}) = (\mathbf{S}, \mathbf{S})$, $\mathbf{I_\sigma(i, w) = 0}$, $\mathbf{I_\sigma(j, w) = 0}$:

This case also results in the $(O, O)$ equilibrium. $i$ initially adopts the Open strategy to reduce upstream transit costs. $j$ adopts Open peering to reduce its own upstream transit costs or to recover lost revenue, as explained in Section III-C2. The analysis of the rest of the game is the same as in Section III-C2.

*4)* $(\mathbf{s_i^0}, \mathbf{s_j^0}) = (\mathbf{O}, \mathbf{O})$, $\mathbf{I_\sigma(i, w) = 1}$, $\mathbf{I_\sigma(j, w) = 1}$:

In this case, $w$ qualifies to be a peer of both $i$ and $j$ under Selective peering. However, both players start from the Open strategy. The initial network for this game is shown in Figure 2b. $i$ can avoid any upstream transit cost with the Selective strategy since $I_\sigma(i, w) = 1$. However, if $i$ adopts Selective peering, it would need to depeer $y$ leading to a loss in revenue because of the flow $T_{yx}$, as shown in Equation 13. Consequently, $i$ stays with the Open strategy. $j$ acts in the same manner. Neither player changes its peering strategy and the equilibrium is $(O, O)$.

The equilibrium $(O, O)$ is suboptimal for both $i$ and $j$. In the case of $i$, for instance,

$$u_i(S, S) - u_i(O, O) = P_i \times V_{jx} + (P_i - \alpha) \times T_{xy} \geq 0 \quad (16)$$

Note that the peering constraints on $w$ in this case are the same as in Section III-C1, but the initial peering strategies are different. While the game of Section III-C1 reaches the optimal equilibrium $(S, S)$, this game stays at the suboptimal equilibrium $(O, O)$.

## D. Peer Preference or Peer Pressure?

We summarize and generalize the effect that was observed in the previous game as follows: a transit provider may decide to switch to Open peering to reduce its upstream transit costs.

By doing so, it can form peering links with at least some customers of its existing peers, diverting revenue-generating traffic from the latter. Those peers would then have the incentive to also switch to Open peering so that they can partially alleviate their lost revenue by attracting traffic that is destined to their customers through direct peering links with others. Thus, transit providers can end up in a state where the loss in transit revenues is larger than their savings in upstream transit costs.

Instead of adopting Open peering to reduce upstream transit costs ("peer preference") some transit providers may be forced to do so to partially alleviate the transit revenue loss caused by their peers ("peer pressure"). We showed that the resulting Open peering equlibrium may be suboptimal for transit providers.

The gravitation of transit providers towards Open peering occurs due to (a) myopic behavior and (b) lack of coordination among peering transit providers. For example, in case III-C3 $i$ myopically decides to keep the peering link $L_{iy}$; otherwise the removal of that link would have reduced its utility. Had $i$ switched to Selective peering, accepting a short-term loss, $j$ would find the contentious link $L_{jx}$ redundant and it would also switch to Selective. In other words, players $i$ and $j$ peer with $y$ and $x$, respectively, not out of preference but because of the pressure to avoid short-term loss.

One could argue that our assumptions of myopic behavior and lack of coordination among Internet transit providers are over-simplistic. It may be true that in a "tiny" Internet with just a handful of providers some form of coordination, or anticipation of the actions of other players, is feasible. Considering however that the Internet consists of thousands of providers, and that each of them only has limited information about the traffic, customers, peers, or even the strategic objectives of other providers, we believe that the previous two assumptions are not unrealistic in practice. In section IV, we show with large-scale simulations and under more realistic conditions that the observed gravitation towards Open peering is also evident there.

## IV. COMPUTATIONAL STUDY

The analytical model of the previous section only considers two transit providers and it does not incorporate several factors that are important in practice such as dynamic, location-dependent competitive pricing, economies-of-scale, public vs. private peering, heavily skewed traffic matrix, etc. In this section, we employ computational agent-based modeling involving a large number of agents in a realistic setting to validate and further investigate the analytical results of the previous section. The main objective is to examine whether a move of transit providers towards Open peering can be observed under more realistic conditions. Further, we want to explore which classes of providers are more affected by this move, and to determine the economic impact of Open peering on the AS population in a macro scale.

## A. Computational Model Description

We use an agent-based simulation framework, called GEN-ESIS, that we developed in earlier work [8]. GENESIS incorporates a large number of agents, non-linear pricing to reflect

economies-of-scale, complex geographic co-location patterns, public and private peering methods, a skewed interdomain traffic matrix that is parameterized to reflect real-world traffic patterns, realistic transit and peering costs, etc. We reproduce the parameterization from GENESIS [8] in Table II for completeness. A detailed presentation of the GENESIS model, validation, choice of parameters and its properties, e.g., scalability, convergence, sensitivity to different parameters, can be found in [8].

In addition, we introduce here two modifications to GENESIS for the purposes of this study:

1) In the original GENESIS model, transit providers are assigned a peering strategy based on their hierarchical status in the network (Tier-1 transit providers use Restrictive, non-Tier1 transit providers use Selective and stubs use Open). In order to reflect the peering strategy decision process described in section III-A, we modify GENESIS so that transit providers evaluate each of the three peering strategies and choose the one that maximizes their utility.

2) In the original GENESIS model, transit prices are randomly assigned to ASes. Further, these randomly assigned prices do not change over the course of the simulation. Here, we introduce competitive, location-dependent dynamic pricing. All transit providers at a given location have the same price. Thus, a transit provider which has presence in more than one locations may have a different price at different locations. The greater the number of transit providers at a location, the lower the transit price. The parameterization of this pricing model is based on data from TeleGeography [23].

### B. Simulation Results

We begin by comparing the following two *scenarios*: *Selective-Restrictive* (SR) and *Selective-Restrictive-Open* (SRO). Under the *SR* scenario, transit providers choose only between the Selective and Restrictive strategies. In the *SRO* scenario, transit providers choose between Selective, Restrictive and Open. Stubs always use Open peering in both scenarios. In each scenario, we run multiple simulations, each with a distinct population of agents, to generate 100 network equilibria. A comparison of the previous two scenarios allows us to investigate the effect of Open peering on both the providers that adopt it and on the Internet as a whole.

*1) Gravitation towards Open peering:* In this section, we compare the *SR* and *SRO* scenarios in terms of the peering strategies adopted by the population of transit providers.

The distribution of peering strategies in the *SR* scenario shows that 90% of providers use Selective peering, while the remaining 10% use Restrictive. In *SR*, most providers peer with agents that have similar hierarchical status. The peers of transit providers include large content providers and consumer stubs that are able to satisfy the Selective peering criteria by virtue of their large traffic volume. On the other hand, the *SRO* scenario results in a radical change in the peering strategy distribution – 79% of transit providers adopt Open peering, 20% adopt Selective peering, while only 1% adopt Restrictive.

We find that the attraction towards Open peering is not uniform across all transit providers. Figure 3 shows that the
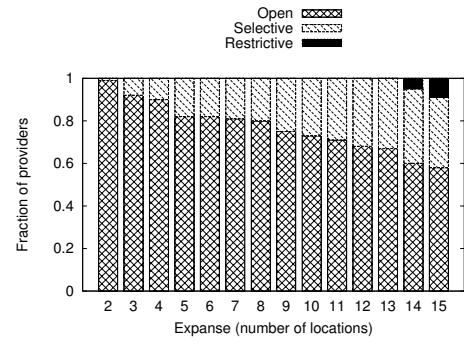


Fig. 3: Peering strategy distribution – classification by geographical expanse.



Fig. 4: Peering strategy distribution – classification by traffic volume.

fraction of transit providers adopting Open peering decreases as their geographical expanse[4] increases. Similarly, Figure 4 shows that the fraction of agents adopting Open peering decreases as their traffic volume increases. When classified based on network hierarchy we find that 6.5% of Tier-1 (agents without transit providers), 58% of Tier-2 (agents which are transit customers of Tier-1 providers ) and 87% of Tier-3 (all non-Tier1 and non-Tier2 agents) transit providers adopt Open peering.

What causes these differences in the adoption of Open peering? The geographical expanse and peering strategy of a provider are correlated for two reasons. First, the local traffic of an agent is a function of its geographical expanse; as expanse increases, so does its local traffic volume [8]. Second, an agent with large expanse can attract more customers, which increases its transit volume. A transit provider with large expanse (and hence large traffic volume) is able to peer with other large transit providers using the *Selective* or *Restrictive* policies, avoiding peering with smaller agents.

*2) Impact on economic utility:* The conventional wisdom is that Open peering is generally associated with a reduction of upstream transit costs and thus, increased utility. However, the complex interdependencies between providers that were discussed in Section-III often cause the opposite effect, i.e.,

---

[4]The geographical expanse of a provider is the number of locations in which it is present. Recall that two agents can interconnect only if they are co-located.

| Parameter, Symbol, Description | Value | Justification |
|---|---|---|
| Number of agents $N$ | 500 | Simulation time constraints. |
| Number of possible locations $G_{max}$ | 50 | Based on ratio of number of IXPs to ASes in the Internet [2]. |
| Geographic expanse distribution | Zipf(1.6) | Based on data about number of participants at each IXP, obtained from PeeringDB [2]. As a result, 50% of nodes are present to only 1 location. Only 0.6% nodes are present at 15 locations, which is the maximum possible expanse of an AS in the model. |
| Generated traffic distribution | Zipf(1.2) | Produces a heavy-tailed distribution of outgoing traffic [24], [25], [26]. |
| Consumed traffic distribution | Zipf(0.8) | Produces heavy-tailed distribution of incoming traffic. |
| Mean consumed traffic | 500 Mbps | $V_C(x) \propto |G(x)|$, a node with large expanse will also have a large number of access customers. |
| Private peering threshold $\Omega$ | 50 Mbps | Based on a survey of peering strategies [22]. |
| Transit cost multiplier $P(x)$ | \$[5,120] | Based on IP transit prices for ARIN region reported by TeleGeography [23]. Values scaled for small number of providers in the simulations. Based on these values the average cost for 1 Gbps transit traffic in our simulation model is \$11/Mbps. |
| Transit cost exponent $\tau$ | 0.75 | Based on data from [4] and [22]. Based on these values the cost for 1 Gbps peering traffic in our simulation model is \$0.31/Mbps. |
| Peering cost multiplier $\alpha$ | \$20 | |
| Peering cost exponent $\beta$ | 0.40 | |
| Selective peering ratio $\sigma$ | 2.0 | |

TABLE II: Input parameters for the GENESIS simulation model [8]

Open peering results in utility loss. In each of the 100 equilibria that we generated using GENESIS, the cumulative utility of all transit providers under the *SRO* scenario is lower than that of the *SR* scenario. As a whole, the population of transit providers does better without Open peering.

A decrease in the cumulative utility of the provider population under *SRO* does not imply, however, that *all* providers see lower utility. We find that 30% of the transit providers have higher utility in the *SRO* scenario. To understand this effect, recall that a provider's utility is a function of its transit costs, peering costs, and transit revenues. We classify transit providers into two classes based on whether their utility increases or decreases by more than 10% between the *SR* and *SRO* scenarios. Providers of both classes see an increase in peering costs and a decrease in transit costs under *SRO*. The difference between the two classes is due to changes in transit revenues. Practically all transit providers that have lower utility under *SRO*, see their transit revenues decrease by more than 20% in that scenario. Among the transit providers that experience a utility increase on the other hand, 70% do not see a significant variation (meaning, it stays within 20% of their utility under *SR*), and only 10% of them have a larger utility increase. Providers that gain significantly in the *SRO* scenario are typically those that a) have customers who cannot peer with other providers (as a result of co-location constraints) and b) they can peer with many other agents.

*3) Who gains and who loses from Open peering?:* We examine here in more detail the characteristics of providers that either gain or lose utility as a result of the gravitation towards *Open* peering. To better understand which providers gain or lose from *Open* peering, we classify them into 4 classes, based on their traffic volume and number of customers.

**Class-1:** Small traffic volume, few customers.
**Class-2:** Small traffic volume, many customers.
**Class-3:** Large traffic volume, few customers.
**Class-4:** Large traffic volume, many customers.

Class-1 includes players in the bottom 30% of providers by traffic volume and the bottom 30% of providers by number
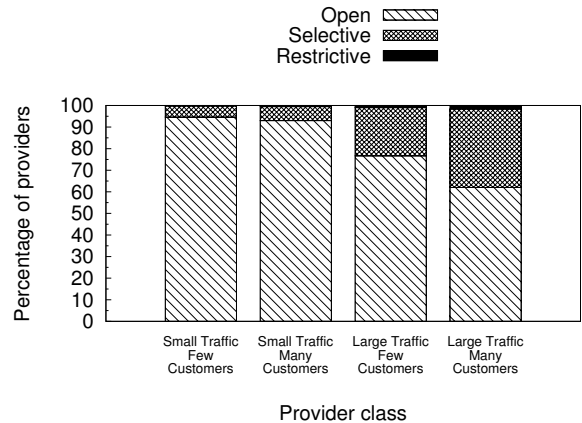


Fig. 5: Peering strategy adoption by 4 classes of transit providers.

of customers. We identify providers in the other three classes similarly. Based on this classification, 16% of providers are in Class-1, 9% in Class-2, 9% in Class-3 and 37% in Class-4. Figure 5 shows the peering strategy distribution in each class. We find that the affinity for *Open* peering decreases as the traffic volume and number of customers increase. In each class, we identify providers that use *Selective* peering under *SR* and *Open* peering under *SRO*. Figures 6(a) and 6(b) show, respectively, the utility of nodes in Class-1 and Class-4 when they switch from *Selective* under *SR* to *Open* under *SRO*.

**Class-1:** 90% of Class-1 players undergo an increase in utility. They reduce their upstream transit costs through Open peering as they are denied peering by larger providers due to their small size. Their customers do not have many peers since they are even more limited in geographic expanse. This makes Class-1 player less vulnerable to their transit traffic being diverted away from them.

to infer the customers of an existing peer $j$ by examining the BGP routes that $j$ sends to $i$ over the peering session.

### A. Utility under DCF

Consider, for instance, the network of Section III-C2. If the DCF constraint is adopted, $i$ will switch to Open peering so that it can directly peer with $w$ but $i$ will not peer with $y$. This allows $i$ to divert all its transit traffic to peering links $L_{ij}$ and $L_{iw}$. Since, no customer traffic is diverted away from $j$, the latter does not have any benefit to switch to Open peering and it stays with the Selective strategy. The equilibrium $(O, S)$ in this case is optimal for both players:

$$u_i(O,S) - u_i(O,O) = (P_i - \alpha) \times (V_{jx} + T_{xy}) \geq 0 \quad (17)$$

$$u_j(O,S) - u_j(O,O) = (P_j - \alpha) \times (V_{iy} + T_{yx}) \geq 0 \quad (18)$$

Recall that the network in Section III-C2 reached equilibrium $(O, O)$ where $i$ not only reduced its upstream transit costs but also suffered a loss in transit revenue. After the introduction of DCF, $i$ only reduces its upstream transit costs. Hence, the DCF variant improves $i$'s utility.

We have also investigated the impact of the DCF Open peering variant on different types of providers using GENESIS. We find that under *SRD*, the cumulative utility of transit providers is greater than under *SRO* and it approaches that with the *SR* scenario. Figure 7 shows the CDF of the cumulative utility of transit providers that use *Selective* peering under the *SR* scenario, but switch to *Open* under *SRO* and *DCF* under *SRD*. The CDF is computed across 100 equilibria.

The cumulative utility of the transit provider population improves with DCF. This is because (a) transit providers no longer "steal" revenue-generating traffic from their peers and (b) transit providers can aggregate peering traffic over few links, reducing peering costs due to the related economies-of-scale. Note, however, that DCF will not have a positive effect on all transit providers. Transit providers higher in the network hierarchy will be more disinclined to peer with those lower in the hierarchy. Hence, some transit providers lower in the hierarchy may have a negative impact on their utility under DCF. In Section IV-B3, we showed that, without DCF, 80% of Class-1 providers gain and 84% of Class-4 providers lose by the introduction of the Open peering strategy. When DCF is adopted, however, 53% of Class-1 providers gain and 68% of Class-4 providers lose utility. In other words, DCF tends to reduce the positive or negative economic impact of Open peering, even though there is still significant variability across different provider classes. Furthermore, if a few providers disregard DCF then their peers conforming to DCF will end up with worse economic utilities.

### B. DCF adoption dynamics

We now explore the behavior of providers if DCF and Open peering are both available as two distinct strategy choices. The key question is whether providers find DCF sufficiently attractive to adopt it without explicit coordination.

We consider the following scenarios:
**SRO:** Providers choose among: 1) *Selective*, 2) *Restrictive*, 3) *Open*.
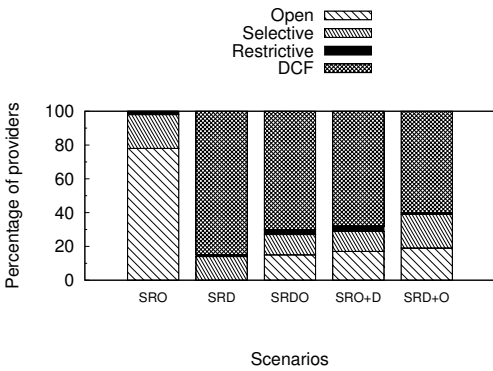**SRD:** Providers choose among: 1) *Selective*, 2) *Restrictive*, 3)



Fig. 6: Utility variations in two classes of transit providers.



Fig. 7: Cumulative utility distribution of transit providers under 3 scenarios

**Class-4:** Providers in Class-4 mostly lose by adopting *Open* peering; 84% of them show utility loss. Their large number of customers makes such providers more vulnerable to having their transit traffic being diverted away from them as their customers peer with other providers. Providers in Classes 2 and 3 have less predictable behavior due to their conflicting characteristics. 55% of providers in Class-2 and 25% in Class-3 increase their utility with Open peering.

## V. AN OPEN PEERING VARIANT

The combination of myopic peering strategy selection and lack of coordination between transit providers results in utility loss for many transit providers. The main underlying issue is that, under the Open strategy, providers peer with customers of their peers. Here, we consider a simple coordination scheme in which transit providers agree not to do that. The proposed scheme is referred to as **Direct Customer Forbiddance (DCF)**. With DCF, a provider $p$ can still peer openly but it should not peer with anyone that is a customer of $p$'s peers.

To illustrate this rule, we refer to the network of Figure 1. If $i$ and $j$ adopt the DCF policy, they cannot peer with $y$ and $x$ respectively. In practice, it is feasible for a transit provider $i$

Fig. 8: Strategy distribution of providers in five distinct scenarios

*DCF*.

**SRDO:** *DCF* and *Open* are both available from the beginning. Providers choose among: 1) *Selective*, 2) Restrictive, 3) Open, 4) *DCF*.

**Perturb *SRO* equilibrium with *SRO+D*:** When the network reaches equilibrium in the *SRO* scenario, we perturb it by adding the *DCF* strategy in the set of available policies, and force providers to re-evaluate their peering strategy.

**Perturb *SRD* equilibrium with *SRD+O*:** When the network reaches equilibrium in the *SRD* scenario, we perturb it by adding Open peering to the available strategies, and force all providers to re-evaluate their peering policy.

Figure 8 shows the strategy distribution for each of the previous five scenarios. We find that the providers are more attracted to DCF than to Open. Further, while DCF is able to significantly perturb the *SRO* equilibrium, attracting 68% of the transit providers, the Open policy is not able to significantly perturb the *SRD* equilibrium. These simulation results imply that the adoption of the *DCF* peering variant may be possible, at least between a large subset of transit providers, even if there is no explicit coordination between them.

## VI. Conclusions

We have investigated the surprising popularity of *Open* peering among transit providers, and the economic implications of this trend in both individual transit providers and their population. With a simple analytical model, we showed that myopic and uncoordinated adoption of peering strategies often leads to the adoption of Open peering by transit providers even when this move hurts their economic utility. With a realistic agent-based model we have evaluated these effects at a larger scale. The overall economic impact of the trend towards Open peering is that most providers lose utility. Certain classes of providers (especially those with small traffic volume and few customers) can gain from Open peering. On the other hand, most transit providers with many customers and large traffic volume lose significantly. Finally, we proposed a variation of Open peering (DCF) that is based on an implicit coordination scheme in which transit providers refrain from peering with customers of their peers. The DCF scheme can be adopted, at least by a large fraction of the transit providers, without explicit coordination and it has a significant positive effect on the cumulative utility of transit providers.

## References

[1] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a large European IXP," in *Proceedings of the ACM SIGCOMM 2012*.

[2] "PeeringDB," http://www.peeringdb.com, October 2011.

[3] "Survey of Characteristics of Internet Carrier Interconnection Agreements," http://www.pch.net/resources/papers//peering-survey/PCH-Peering-Survey-2011.pdf.

[4] 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.

[5] P. Holme, J. Karlin, and S. Forrest, "An Integrated Model of Traffic, Geography and Economy in the Internet," *ACM SIGCOMM CCR*, 2008.

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

[7] 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.

[8] A. Lodhi, A. Dhamdhere, and C. Dovrolis, "GENESIS: An Agent-Based Model of Interdomain Network Formation, Traffic Flow and Economics," in *Proceedings of IEEE INFOCOM*, 2012.

[9] S. Shakkottai, M. Fomenkov, R. Koga, D. Krioukov, and K. Claffy, "Evolution of the Internet AS-Level Ecosystem," *European Physical Journal B*, no. 2, pp. 271–278, Mar 2010.

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

[11] S.-H. Yook, H. Jeong, and A.-L. Barabasi, "Modeling the Internet's Large-scale Topology," *Proc Natl Acad Sci USA*, vol. 99, no. 21, pp. 13 382–6, 2002.

[12] 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.

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

[14] N. Badasyan and S. Chakrabarti, "A Simple Game Theoretic Analysis of Peering and Transit Contracting Among Internet Access Providers," in *Telecommunications Policy Research Conference*, 2005.

[15] S. Lippert and G. Spagnolo, "Internet Peering as a Network of Relations," *Telecommunications Policy*, vol. 32, pp. 33–49, February 2008.

[16] P. Baake and T. Wichmann, "On the Economics of Internet Peering," *Netnomics*, vol. 1, pp. 89–105, 1999.

[17] E. Anshelevich, B. Shepherd, and G. Wilfong, "Strategic Network Formation through Peering and Service Agreements," in *Proceedings of FOCS*, 2006.

[18] S. Shakkottai, "Economics of Network Pricing with Multiple ISPs," in *Proceedings of IEEE INFOCOM*, 2005.

[19] E. Meirom, S. Mannor, A. Orda *et al.*, "Formation Games and the Internet's Structure," *arXiv preprint arXiv:1307.4102*, 2013.

[20] F. Wang and L. Gao, "On Inferring and Characterizing Internet Routing Policies," in *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, 2003.

[21] "North American Network Operators' Group," http://www.nanog.org.

[22] "A Study of 28 Peering Policies," peering.net/white-papers/Peering-Policies/A-Study-of-28-Peering-Policies.html.

[23] Telegeography, "IP Tranist Pricing Service," http://www.telegeography.com, 2011.

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

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

[26] 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 ACM SIGCOMM IMC*, 2004.