An Economically-Principled Generative Model of AS Graph Connectivity

Citation

Corbo, Jacomo, Shaili Jain, Michael Mitzenmacher, and David C. Parkes. 2009. An economically-principled generative model of AS graph connectivity. IEEE Infocom 2009: 2941-2945.

Published Version

doi:10.1109/INFCOM.2009.5062263

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3967319>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

An Economically-Principled Generative Model of AS Graph Connectivity

Jacomo Corbo
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104
jacomo@wharton.upenn.edu

Shaili Jain Michael Mitzenmacher David C. Parkes
School of Engineering and Applied Science
Harvard University
Cambridge, MA 02138
{shailij, michaelm, parkes}@eecs.harvard.edu

I. INTRODUCTION

End-to-end packet delivery in the Internet is achieved through a system of interconnections between the network domains of independent entities called Autonomous Systems (ASes). Inter-domain connections are the result of a complex, dynamic process of negotiated business relationships between pairs of ASes. We present an economically-principled generative model for Autonomous System graph connectivity. While there is already a large literature devoted to understanding Internet connectivity at the AS level, many of these models are either static or based on generalized stochastics.

In a thoughtful critique of such models, Li, Alderson, Doyle and Willinger [10] show that while many generative models reproduce certain statistical features of the AS graph, they fail to capture the good performance of realistic networks [10]. In a study of the AS's intra-domain graph, Li, Alderson, Willinger and Doyle [11] define performance instead in terms of network throughput and show that it is very unlikely that randomized generative models will yield graphs that have the highly-optimized structure of real-world networks. The goal of this paper is to provide insight into the economic drivers that yield, over time, the rich and complex AS interconnection patterns that constitute today's Internet.

Notable features of our model include the assignment of AS business models with an asymmetric gravity model of inter-domain traffic demand [3], an explicit representation of AS utility that incorporates benefits for traffic routed, congestion costs, and payments between ASes, and a deterministic process for link revision that can cascade throughout the network. This is the first attempt at AS graph modeling that incorporates a diffusion process to capture how ASes respond to direct and indirect externalities from changes in the network structure, which brings it closer to an equilibrium model.

We validate our model against other generative models. To do this, we define the social planner's problem which is parameterized by the business models of the graph and provide a method to compare earlier generative models with our model by optimizing the placement of business models on the network. We find that our model yields graphs that are better performing as compared to other dynamic generative models. We also show that our model yields a structured placement of nodes endogenously, where this placement of nodes

generally reflects ASes' business models. This is some of the first evidence of the significance of the business competitive landscape in determining the structure of the AS graph.

A. Related Work

Previous economically-principled models [1], [7] have been formulated as game-theoretic models with static equilibria, which makes it difficult to understand the graph's evolution over time. Moreover, these models assume homogeneous or identical ASes, that edges have uniform costs or capacities, and hinge on fixed demand models.

Chang et al.'s [4] formulation of an AS's decision problem uses an empirically-motivated demand model previously introduced in [3], which we also employ. Our model differs from theirs in that our utility function is explicit about the economic tradeoffs at play and our model does not involve any randomization beyond the sampling of business models, which is tuned here to empirically-measured distributions. While Chang et al.'s model also allows for the revision of links, each AS revises its links when a periodic "timer" goes off. Their method of link revision does not cascade throughout the network as ASes react to their neighbors' link revisions, as is provided in our diffusion process.

II. THE AS GRAPH FORMATION MODEL

Our model contains the following components: each node is assigned a *business model*, traffic demands are computed using an asymmetric gravity model, a utility function that incorporates benefits for traffic routed, congestion costs, and payment transfers between ASes, and revision of links using a forest fire diffusion model.

In our model, AS nodes come into contact with other nodes and lay down links to maximize their economic benefits. Our model only considers customer-provider links and where the decision to establish a link is always initiated by the customer, who pays the provider for the link and essentially for access to the rest of the network. The joint actions of customers choosing their providers defines a directed graph reflecting the customer-provider business relationships between ASes, with edges from customers to providers. While link payments are one-sided, traffic flows in both directions since the customer pays for transit traffic to and from its providers.

Inter-domain traffic demand is tied to ASes' respective

customer bases. We capture this by a model given by Chang, Jamin, Mao and Willinger [3], assigning each node a business model according to their distribution from empirical data and defining traffic demands based on these. Generally speaking, a customer AS's utility for connecting to a provider AS's domain is a function of how that connection will affect its own customers' traffic demand, the link's impact on its links and network congestion, and finally on the balance of payments made and received by it for routing traffic along all of its adjacent customer-provider connections.

A. Overview of the Dynamic Process

As nodes arrive in the network, they are given a business model that is chosen randomly from a joint distribution [3]. A newly born node chooses to connect to the existing AS inter-domain graph by picking the single AS provider that myopically maximizes its utility function. The dynamic process unfolds in discrete periods as follows.

The Dynamic Process:

- 1) In period t , a single AS node is born with business attributes $\langle \alpha_i, \beta_i, \gamma_i \rangle$
- 2) The newly born AS proceeds to place a *single* link to maximize its utility, indicating who will be its provider
- 3) The new node's provider then has the occasion to revise its links. It can either lay down a single additional provider link, replace any number of its current provider links or delete provider links based on the action that maximizes its utility.
- 4) If the new node's provider decides to make a change, then each of its providers has the occasion to do the same. Likewise, for each of these providers that make a change, each of their providers have occasion to do the same also. This process burns through the graph in a depth first search manner, backtracking when revision stops on a particular branch.

In what follows, let N denote the set of nodes with $n = |N|$. The action of a node $i \in N$ is a vector $s_i \in \{0, 1\}^n$ indicating which nodes i has chosen as its providers. We let $s = s_1 \times \dots \times s_n$ be the joint action of all nodes.

B. The Graph

The joint action s defines a directed graph $G(s) = (N, E(s))$. An edge $e = (i, j) \in E(s)$ is established if and only if $s_i(j) = 1$ and designates that i is a customer of j , which is to say that i pays j for the link. Let $E_i(s_i) = \{(i, j) | s_i(j) = 1\}$ be the set of node i 's provider links, with $E(s) = \cup_{i \in N} E_i(s_i)$. Moreover, let $E_i^u(s) = \{(i, j) : s_i(j) = 1\} \cup \{(j, i) : s_j(i) = 1\}$ refer to *all* edges adjacent to i .

C. AS Business Models

An AS's business model reflects its utility for incoming and outgoing traffic, as well as its disutility for routing traffic through its domain. We follow the business model representation of Chang et al. [3]. Formally, each node $i \in N$ has a business model parameterized with coefficients $(\alpha_i, \beta_i, \gamma_i) \in (0, 1]^3$ where α_i reflects AS i 's demand for inbound traffic and β_i reflects its demand for outbound traffic. The parameter γ_i captures an AS's relative capacity to be an effective inter-domain access provider. A high value of

Fig. 1: Degree Distribution for our AS graph model compared to that of the preferential attachment graph model

γ_i suggests that i is an effective provider. We can think of the business model parameters $(\alpha_i, \beta_i, \gamma_i)$ as representing an AS's utility for providing residential access, web hosting, and business access services, respectively. The business model coefficients $(\alpha_i, \beta_i, \gamma_i)$ are chosen from the joint distribution $F(a, \Sigma)$, where a refers to the distribution of γ_i , which is currently drawn from an empirically-derived power law. Since business model coefficients are highly correlated in real-life, we use an empirically-derived pairwise correlation matrix Σ to compute α_i and β_i . Once business models have been assigned, we compute traffic demands based on an asymmetric gravity model (from Chang et al. [3]).

The traffic demand is represented by the matrix $B(G(s))$ and entry $b^{kl}(s)$ represents the total demand for traffic from k to l . We assume that all traffic is routed. The routing policy determines the value of $x_e^{kl}(s)$ for all $e \in E$ and all pairs of $k, l \in N$, where $x_e^{kl}(s)$ is defined as the flow of traffic originating from k and destined for l traveling along edge e and assume that no packets are dropped by ASes. We assume that routing policy designates a single path from source to sink along which to route traffic. We denote $P_{kl}(G(s))$ as the path that the routing policy designates to send traffic from k to l . Therefore, we have

$$x_{i,j}^{kl}(s) = \begin{cases} b^{kl}(s) & \text{if } (i, j) \in P_{kl}(G(s)) \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

The routing policy that we use is the “No Valley and Prefer Customer” Routing Algorithm [6], which is closer to the way traffic is routed on the real AS graph as opposed to shortest path routing.

D. The Utility Function

Given the set of nodes N , with corresponding business models, the graph $G(s)$ of inter-domain connections and traffic demand matrix $b^{kl}(s) \in B(s)$, the utility of an AS node i is as follows:

$$u_i(s) = \sum_{j \in N} b^{ij}(s) + \sum_{j \in N} b^{ji}(s) - \tau_i \cdot \sum_{e \in E_i(s)} \sum_{k, l \in N} x_e^{kl}(s) - \sum_{j: (i, j) \in E_i(s)} t_{ij}(s) + \sum_{j: (j, i) \in E_j(s)} t_{ji}(s)$$

The third term represents the congestion cost associated with traffic routed through node i , while the last two terms rep-

resent transfers between AS i and its customers and providers. Recall that $E_i^u(s)$ designates all edges adjacent to i . Note that this is a cost τ_i applied to all traffic flow through i , capturing i 's cost for routing inbound and outbound traffic as well as for all transit traffic.

ASes with lower transit routing costs in turn provide more affordable and more reliable service to customer domains, making them the preferred inter-domain access providers. The price charged to a customer AS, j , by the provider AS, i , is modeled as a function of the congestion associated with traffic being routed through i 's servers and sub-domains as well as a mark-up for the bilateral traffic flow along the purchased link. We assume that the mark-up on traffic from costs incurred is separable from price of flow on a link. That j should pay for traffic in both directions is how most customer-provider arrangements are made, reflecting the fact that j is paying i for access to the rest of the AS graph network.

We stress that λ_i and μ_i are a function of i 's identity, both terms relating to i 's effectiveness as an access provider, and that they are customer-anonymous, i.e. independent of j . In practice, j 's traffic demand along the proposed link does matter in these per-unit charges. Our assumption holds particularly well for small customers linking to much larger providers and for the rare cases where large customers link to relatively small providers.

An AS i 's costs for routing inter-domain traffic are tied to how its network is provisioned. Two important factors affecting an AS's transit costs are the length of inter-domain links and the inter-domain bandwidth capacity. Lower transit costs are associated with topologies with greater geographic coverage and topologies that are optimized for larger traffic volumes. In our model, the effectiveness of i as an access provider is captured by the coefficient γ_i of its business model. With this in mind, we choose parameters τ_i , λ_i and μ_i to vary super-linearly in γ_i to reflect the large variability among ASes' prices (and presumably costs) for customer traffic [4]. Precisely, for $\tau, \lambda, \mu > 0$ as model parameters, we have that

$$\tau_q = \tau \cdot e^{-\gamma_q}, \quad \lambda_q = \lambda \cdot e^{-\gamma_q}, \quad \mu_q = \mu \cdot e^{-\gamma_q}. \quad (2)$$

E. Revision of Links

Once a node lays down a link to its provider, this provider is given the opportunity to revise its links and this process continues recursively until no providers make a change. This process propagates upstream from customers to providers. Customers may add a link to a new provider, replace or delete links to existing providers, but a provider may not add, replace, or delete customer links. In order to make our revision process tractable and ensure that our revision process does not cycle, we perform a depth first search where branches die out once a provider decides not to make a change. The revision of links process in our model is much like the Forest Fire model of Leskovec et al. [9] although our revision process is deterministic rather than randomized.

At each node in the depth first search, a node computes its best response function for either adding a single provider

link given the current topology of the network, deleting any number of its provider links, or reconfiguring its current budget of provider links by repositioning or deleting its current links and adding up to a single additional link to its budget.

We can analogously define a revision process with limiting depth d , where the depth first search revision process terminates on a branch either when a node does not make a revision or when the distance from the original node along the branch exceeds d .

III. SIMULATION METHODOLOGY

A. Defining the Performance of a Graph

Our main interest is in evaluating the relative performance of generated graphs against those of other generative models. We use a measure of social welfare to compare the relative performance. In this, we follow Li et al. [10] who notice that rule-based and purely stochastic generative models may reproduce certain statistical features of the graph, such as power-law degree distributions, but fail to capture important structural features related to the performance of the graph in question.

We define the performance of a graph as the social welfare function $W(G) = \sum_{i \in N} u_i(G)$. Notice that all the payments cancel out, so this objective function is just the total demand met by the network discounted by the congestion cost experienced by all nodes.

$$W(G) = \sum_{i \in N} \left(\sum_{l \in N} b^{il} + \sum_{k \in N} b^{ki} - \tau_i \cdot \sum_{e \in E_i^u} \sum_{k, l \in N} x_e^{kl} \right) \quad (3)$$

This is a reasonable model of social welfare for a network of utility-maximizing ASes.

B. Comparing Network Performance

We compare graphs generated by our generative model (ASGM) against a number of graph topologies: preferential attachment (PA) graphs [2], copying model (CM) graphs [8], and graphs derived from the generalized random graph (GRG) model [5]. In the preferential attachment generative model, nodes are born one at a time. In each time step, with probability p , the new node connects to a node already in the network uniformly at random. With probability $1 - p$, the new node connects to a node already in the network in proportion to its total degree. In the copying model, we have an out-degree parameter k and a probability p . Much like the preferential attachment model, nodes are born one at a time. In each time step, the newly born node v forms d outlinks. In each time step, a prototype vertex v_p is chosen uniformly at random from the set of nodes already in the graph. Node v then forms d outlinks as follows. With probability p , node v 's i^{th} outlink is chosen uniformly at random from nodes already in the graph. With probability $1 - p$, node v 's i^{th} outlink is chosen to be node v_p 's i^{th} outlink. The GRG model generates random graphs with a given expected degree sequence $D = \{d_1, d_2, \dots, d_n\}$. In the GRG model, a link is formed between nodes i and j with probability proportional to $d_i d_j$.

Fig. 2: The (a) social welfare of our AS graphs is compared to copying model (CM), preferential attachment (PA) and generalized random graphs (GRG) as a function of the size N of these graphs. Our model graphs with business models optimally allocated are also shown (OPT(ASG)). We also explore the space of connected network graphs having the the same power law degree sequence in (b). The power law degree distribution against which all networks are compared is the average distribution achieved endogenously over 50 runs of our model. A pairwise rewiring procedure (described in [10]) is used to fit 50 individual instances of each type of network to this degree distribution, while preserving a simple, connected graph. Welfare values for all networks are shown, along with average performance of the original networks against the average s-metric of the rewired networks (see bold values). (c) Business Model Coefficient vs. Betweenness Centrality (d) Social Welfare and Diameter as a function of Burn Depth. Social welfare $W(G)$ and the s-metric value $S(G)$ as a function of the variance of transit provider types γ . Recall that the s-metric value is a measure of a graph's relative likelihood against a background set of graphs with the same degree distribution. Since the degree distribution obtained endogenously by our model varies with the distribution of business models, each s-metric value in (b) is normalized against the maximal s-metric graph for a given business model type distribution.

In Section IV, we compare graphs generated by our AS graph model to other graphs in terms of social welfare. We measure a graph's performance in terms of a traffic demand-based social welfare function and the computation of traffic demands is predicated on business models, which requires us to assign business models to preferential attachment, copying, and GRG random graphs. We do so optimally. We ensure that the graphs are of the same size and that the N nodes that comprise all graphs are made up of exactly the same set of business models. Given a set of business models and a graph, we assign a business model to each node in the graph to optimize the graph's particular objective function (i.e. welfare, demand, congestion), in the spirit of presenting all comparison graphs in the best possible light. The traffic demand model makes this a non-linear assignment problem which we solve by adaptive simulated annealing.

We also compare the structural features of the graphs yielded by these generative models in Section IV. We do this by way of the *s-metric* introduced by Li et al. [10]. The *s-metric* measures the extent to which high-degree nodes are connected to other high-degree nodes. A useful property of the s-metric is that it can distinguish between graphs that have the same degree sequence:

Definition III.1. For any graph $G = (V, E)$ with degree sequence $D = \{d_1, d_2, \dots, d_n\}$ the s-metric is defined as follows:

$$S(G) = \sum_{(i,j) \in E} d_i d_j$$

Since the use of the s-metric requires the same degree sequence, we use a randomized rewiring process as in Li et al. [10]. The s-metric also gives the likelihood of a graph, in that given a fixed degree sequence, higher $S(G)$ values correspond to graphs that more more likely and lower $S(G)$ values correspond to graphs that are less likely [10].

IV. SIMULATION RESULTS

In this section, we show a number of properties that our dynamically grown graphs satisfy. In particular, we show that our model can be parameterized so as to endogenously form networks with a power law node degree distribution. Significantly, in stark contrast to other randomized and rule-based topology generators, we find that power law graphs generated by our model bear the hallmarks of good engineering design in terms of good congestion, throughput, and total node welfare properties. We investigate the sensitivity of our results to heterogeneity in the type space, showing that the empirically-motivated choice of a yields graphs that have particular *good* features and performance. We restrict our business model distribution parametrization, in our choice of a and Σ , as per empirical measurements in [3].

A. Power Law Networks

With the traffic demand parameterized as per [3], we observe that our generated AS graphs satisfy power law degree

distributions (as we would expect [10]), shown by approximately linear behavior on a log-log scale. Though power law degree distributions are not unusual, they are still a key statistical property that AS graphs satisfy. Any valid AS model must generate graphs with power law degree distributions, however they should not be the only metric by which to judge AS graph generative models [10]. We observe that the customer degree distribution, the provider degree distribution and the overall degree distribution of our dynamically generated graphs satisfy a power law degree distribution. In Figure 1, we compare the degree distribution of graphs from our AS graph model with the degree distribution of graphs generated with preferential attachment model, appropriately parameterized to yield a close fit to the total degree distribution of our model-generated graphs, as well as to the degree distribution of preferential attachment graphs that have been randomly rewired so as to yield an even more precise fit to the degree distribution yielded by our model. We observe that the exponent of the customer degree distribution is smaller (i.e. more negative) than the exponent of the overall degree distribution, which is in turn smaller than the exponent of the provider degree distribution, which is what is found in practice [4].

B. Comparing Network Performance

In Figure 2(a), we find that our AS graph generative model yields graphs with higher social welfare than graphs generated by the preferential attachment and copying model. Note that this is even though business model placement is endogenous in our graphs.

The results in Figure 2(b) yield similar results to [10], in that graphs with high $\mathcal{S}(G)$ values tend to have low performance. These set of graphs show that our model-generated graphs fare well against all other graphs, even against copying model graphs, which intrinsically have many more edges. This speaks to the power of economic constraints on AS behavior to achieve good system-wide performance, even if AS actions are uncoordinated and myopic. That the AS graph model yields a narrow band of networks on the likelihood scale is evidence of a carefully engineered design.

Figure 2(b) also shows the value of the optimal business model allocation for the ASGM graphs and we find that the optimal business model allocation is close in performance to that of the ASGM graphs. We also observe that our model reproduces, endogenously, something near the optimal, social welfare-maximizing placement of business models in the graph. We judge this by labeling all nodes according to their dominant (highest value) business model coefficient and measuring a node's location in the graph by its *betweenness centrality*:

Definition IV.1. Given the graph $G = (V, E)$, the betweenness centrality $C_B(v)$ of a node $v \in V$ is $\sum_{s,t \in V} \frac{\sigma_{st}(v)}{\sigma_{st}}$, where σ_{st} refers to the number of shortest geodesic paths between s and t , and $\sigma_{st}(v)$, the number of shortest geodesic paths between s and t that pass through node v .

Figure 2(c) shows a result regarding the average value of

the dominant business model coefficient of nodes against their betweenness centrality. It compares results for graphs derived from our AS graph model with link revisions to our AS graph model without link revisions. The plot suggests that our model endogenously achieves an optimizing placement of business models in the graph. For our model-generated graphs, as the dominant coefficient grows, business access providers move toward the center of the graph quickest of all, whereas residential access providers are much more likely to be located on the fringes of the graph. The relative sensitivity of the allocation of business models in our model to the size of the transit provider coefficient is illustrated by the sample correlation coefficient between betweenness centrality and the three business model coefficients for all nodes in the graph. Considering over 50 network instances, we obtain correlations of **0.33**, **0.41**, and **0.67** for residential access, web hosting, and business provider coefficients, respectively.

We also consider the role of the link revision process in Figure 2(d). We plot social welfare as a function of the depth of forest fire link revision. The results of Figure 2(d) are twofold, in that link revision is crucial to obtaining graphs with high social welfare and that link revision is not necessary beyond small depths. This is exactly what one would expect since the AS graph is known to have a small diameter (between 5-7). To reinforce this point, we have also shown the diameter of the network as a function of the link revision depth. For all values of link revision depth, the diameter is fairly small, but it continues to decrease as we increase link revision depth.

ACKNOWLEDGEMENTS

We thank seminar participants at Harvard, NYU Stern, 2007 Joint Workshop on The Economics of Networked Systems and Incentive-Based Computing NetEcon+IBC. We also thank 4 anonymous reviewers for constructive comments.

REFERENCES

- [1] E. Anshelevich, B. Shepherd, and G. T. Wilfong. Strategic Network Formation through Peering and Service Agreements. In *FOCS*, pages 77–86, 2006.
- [2] A.-L. Barabasi and R. Albert. Emergence of Scaling in Random Networks. In *Science*, volume 286, pages 509 – 512, 1999.
- [3] H. Chang, S. Jamin, Z. Mao, and W. Willinger. An Empirical Approach to Modeling Inter-AS Traffic Matrices. In *Proc. of ACM Internet Measurement Conference*. ACM Press, 2005.
- [4] H. Chang, S. Jamin, and W. Willinger. To Peer or not to Peer: Modeling the Evolution of the Internet's AS-level Topology. In *INFOCOM*, 2006.
- [5] F. Chung and L. Lu. The Average Distance in a Random Graph with Given Expected Degrees. *Internet Mathematics*, 1:91–113, 2003.
- [6] L. Gao and F. Wang. The Extent of AS Path Inflation by Routing Policies. In *Proceedings of IEEE Global Internet Symposium*, 2002.
- [7] R. Johari, S. Mannor, and J. N. Tsitsiklis. A Contract-Based Model for Directed Network Formation. *Games and Economic Behavior*, 2006.
- [8] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. Stochastic Models for the Web Graph. In *Proc. of Foundations of Computer Science (FOCS)*, 2000.
- [9] J. Leskovec, J. M. Kleinberg, and C. Faloutsos. Graphs Over Time: Denisification Laws, Shrinking Diameters and Possible Explanations. In *KDD*, pages 177–187, 2005.
- [10] L. Li, D. Alderson, J. C. Doyle, and W. Willinger. Towards a Theory of Scale-Free Graphs: Definition, Properties, and Implications. *Internet Mathematics*, 4(2), 2006.
- [11] L. Li, D. Anderson, W. Willinger, and J. Doyle. A First-Principles Approach to Understanding the Internet's Router-level Topology. In *SIGCOMM*, 2004.