Evolution of the Internet AS-Level Ecosystem

Srinivas Shakkottai[1], Marina Fomenkov[2], Ryan Koga[2], Dmitri Krioukov[2], and Kc Claffy[2]

[1] Texas A&M University, College Station, USA
sshakkot@tamu.edu
[2] Cooperative Association for Internet Data Analysis, University of California, San Diego, USA

Abstract. We present an analytically tractable model of Internet evolution at the level of Autonomous Systems (ASs). We call our model the multiclass preferential attachment (MPA) model. As its name suggests, it is based on preferential attachment. All of its parameters are measurable from available Internet topology data. Given the estimated values of these parameters, our analytic results predict a definitive set of statistics characterizing the AS topology structure. These statistics are not part of model formulation. The MPA model thus closes the "measure-model-validate-predict" loop, and provides further evidence that preferential attachment is the main driving force behind Internet evolution.

Keywords: Preferential attachment, Internet evolution, AS-level topology, Internet measurement.

1 Introduction

In the past decade we have seen extraordinary and relentless growth of Internet connectivity around the world. This rapid development has led to a burgeoning of companies generating, carrying, and sinking content. Each company with its own routing domain is roughly represented as an Autonomous System (AS) in the global routing system. An AS might be a transit Internet service provider (ISP), a content provider (or sink), or a combination of these. Some ASs are highly interconnected, while others have only a few links. In 1999 Faloutsos *et al.* [1] observed that despite all this diversity, the distribution of AS degrees obeys a simple power law.

Many researchers have attempted to model the Internet as an evolving system [2,3,4,5,6,7,8,9,10,11]. However, questions regarding the main drivers behind Internet topology evolution remain [12]. In this paper our main objective is to create an evolutionary model of the AS-level Internet topology that simultaneously:

1. is realistic,
2. is parsimonious,
3. has *all* of its parameters measurable,
4. is analytically tractable, and
5. "closes the loop."

J. Zhou (Ed.): Complex 2009, Part II, LNICST 5, pp. 1605–1616, 2009.

Parsimony implies that the model should be as simple as possible, and, related to that, the number of its parameters should be as small as possible. The fifth requirement means that if we substitute measured values of these parameters into analytic expressions of the model, then these expressions will yield results matching empirical observations of the Internet. However, most critical is the third requirement [12]: as soon as a model has even a few unmeasurable parameters, one can freely tune them to match observations. Such parameter tweaking may create an illusion that the model "closes the loop," but in the end it inevitably diminishes the value of the model because there is no rigorous way to tell why one model of this sort is better than another since they all match observations. To the best of our knowledge, the multiclass preferential attachment (MPA) model that we propose and analyze in this paper is the first model satisfying all five requirements listed above.

A salient characteristic of our model is that we distinguish between two kinds of ASs: ISPs and non-ISPs. The main difference between these two types of ASs is that while both ISPs and non-ISPs can connect to ISPs, no new AS will connect to an existing non-ISP since the latter does not provide transit Internet connectivity. In Section 2 we analyze the effect of this distinction on the degree distribution. In Section 3 we account for other processes. ISPs can form peering links to exchange traffic bilaterally. They can also go bankrupt and be acquired by others. Finally, they can multihome, i.e., connect to multiple providers. Prior work has often focused on these processes as the driving forces behind Internet topology evolution. We show that in reality they have relatively little effect on the degree distribution. Using the best available Internet topology data, we measure the parameters reflecting all the process above, and analytically study how they affect the degree distribution.

However, the degree distribution alone does not fully capture the properties of the Internet AS graph [13]. The dK-series formalism introduced in [13] defines a systematic basis of higher-order degree distributions/correlations. The first-order ($1K$) degree distribution reduces to a traditional degree distribution. The second-order ($2K$) distribution is the joint degree distribution, i.e., the correlation of degrees of connected nodes. The distributions can be further extended to account for higher-order correlations [13], or for different types of nodes and links, called annotations [14]. In the economic AS Internet, there are two types of links: links connecting customer ASs to their providers (c2p links), and links connecting ISPs to their peers (p2p links). Reproducing the $2K$-annotated distribution of AS topologies suffices to accurately capture virtually all important topology metrics [13,14]. An important feature of the MPA model is that, by construction, it naturally annotates the links between ASs by their business relationships. In Section 4 we perform a $2K$-annotated test. We generate synthetic graphs using the MPA model, and find that these graphs exhibit a startling similarity to the observed AS topology according to almost all the $2K$-annotated statistics. This validation, in conjunction with observations in [13, 14], ensures that other important topology metrics also match well.

2 Two-Class Preferential Attachment Model

We first recall the original preferential attachment (PA) model [15]. In the PA model there is only one type of nodes. Suppose that nodes arrive in the system at the rate of one node per unit time. Let them be numbered $s = 1, 2, 3, ...$ as they arrive. Then the number of nodes in the system at time t is equal to t. A node entering the system brings a link with it. One end of the link is already connected to the entering node, while the other end is loose. We call such an un-associated end a *loose connection*. According to the PA model, nodes attach to existing ones with a probability proportional to their degrees. Thus, the probability that a node of degree k is selected as a target for the incoming loose connection, is its degree divided by the total number of existing connections in the system, which is $\frac{k}{2t}$. The original PA model yields a power-law degree distribution $P(k) \sim k^{-\gamma}$ with $\gamma = 3$, but the linear preference function can be modified by an additive term such that the model produces power laws with any $\gamma > 2$ [15, 16, 17].

There are two fundamentally different types of ASs—ISPs and non-ISPs—that differ in whether they provide traffic carriage between the ASs they connect or not. No new AS would connect to an existing non-ISP since it cannot provide Internet connectivity. No other work has attempted to model this observation, which is fundamental to understanding the evolution of Internet AS-level topology. Thus, our first modification of the PA model is to consider these two classes of nodes (Figure 1). New ISPs appear at a rate 1 and connect to other ISP-nodes with a linear preference. New non-ISPs appear at some rate ρ per unit time and attach themselves to existing ISP-nodes with a linear preference. However, *no further attachments to non-ISPs can occur*. Thus, with respect to degree distribution, only the ISP-nodes contribute to the tail of the power law as the non-ISP nodes will all have degree 1. Given that a link between an ISP node and a non-ISP node has only one end that contributes to the degree of ISP nodes, and since we look only at ISP nodes to find the degree distribution, the links that have an ISP node on one end and a non-ISP node on the other should be counted only once, i.e., they are counted as contributing 1 connection.

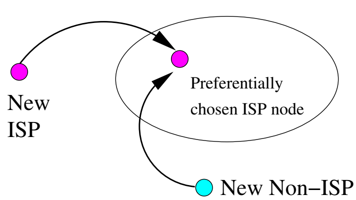

Fig. 1. Two-class preferential attachment model. There are ISP nodes and non-ISP nodes.

Since each ISP node contributes 2 connections to the network and a non-ISP node contributes only 1, the total number of existing connections in the network at time t is $(2 + \rho)t$, which implies that the probability of any loose connection connecting to an ISP-node of degree k is

$$\frac{k}{(2+\rho)t}. \tag{1}$$

We use the notation $p(k,s,t)$ to denote the probability that an ISP node s has degree k at time t. Then the average degree of an ISP node s at time t is

$$\bar{k}(s,t) = \sum_{k=1}^{\infty} kp(k,s,t) \tag{2}$$

Both entering ISPs and non-ISPs have one loose connection each, so the number of loose connections entering the system at time t is $1+\rho$. Then, from (1), the continuous-time model of the system is

$$\frac{\partial \bar{k}(s,t)}{\partial t} = \frac{1+\rho}{(2+\rho)t}\bar{k}(s,t), \tag{3}$$

with boundary condition $\bar{k}(t,t) = 1$ for $t > 1$. Solving this equation yields

$$\bar{k}(s,t) = \left(\frac{s}{t}\right)^{-\frac{1+\rho}{2+\rho}}. \tag{4}$$

This model represents a deterministic system in which if ISP node s has degree $\bar{k}$, then ISP nodes that arrived before s (in the interval $[0,s)$) have degree at least $\bar{k}$. Thus, s also represents the number of ISP nodes that have degree at least $\bar{k}$. It follows from (4) that the number of ISP nodes that have degree $\bar{k}$ or higher is

$$\frac{t}{\bar{k}^{\frac{2+\rho}{1+\rho}}}.$$

Note that the number of ISP nodes that arrive in $[0,t]$ is just t. Hence, the *fraction* of ISP nodes that have degree $\bar{k}$ or higher is

$$\frac{1}{\bar{k}^{\frac{2+\rho}{1+\rho}}}.$$

Since this fraction is essentially a complimentary cumulative distribution function, we differentiate it and multiply by -1 to obtain the density function

$$f(\bar{k}) = \frac{2+\rho}{1+\rho}\bar{k}^{-\left(2+\frac{1}{1+\rho}\right)}, \tag{5}$$

which corresponds to the probability distribution function

$$P(k) \sim k^{-\left(2+\frac{1}{1+\rho}\right)}. \tag{6}$$

Validation Against Observed Topology 1. *Dimitropoulos* et al. *[18] applied machine learning tools to WHOIS and BGP data to classify ASs into several different classes, such as Tier-1 ISPs, Tier-2 ISPs, IXPs, universities, customer ASs, and so on. They validated the resulting taxonomy by actual examination of a large* *number of ASs. We use their results and divide ASs into two classes based on whether they are ISPs or not. According to [18], the number of ISPs is about* 30% *of ASs, while non-ISPs make up the other* 70%. *The measured value* $\rho = 7/3$ *yields* $P(k) \sim k^{-2.3}$*, whose exponent is quite close to the observed value of about* -2.1 *to* -2.2 *[1, 19, 20].*

A major implication of this section is that the observed value of the power-law exponent finds a natural and simple explanation: it is due to preferential attachment and to a directly measured high proportion of non-ISP nodes to which newly appearing nodes cannot connect.

3 Multiclass Preferential Attachment: Peering, Bankruptcy, Multihoming and Geography

In this section we add further refinements to our model and show that, contrary to common beliefs, none of these refinements have a significant impact on the degree distribution shape.

Relationships between ASs change over time, as ASs pursue cost-saving measures. If the magnitude of traffic flow between two ISPs is similar in both directions, then reciprocal peering with each other allows each ISP to reduce its transit costs. Under the assumption that all customer ASs generate similar volumes of traffic, high degree ASs would exchange high traffic volume and rationally seek to establish reciprocal peering with other high degree ASs. We denote the rate at which peering links appear by c. The probability that a new peering link becomes attached to a pair of ISP-nodes of degree k_1 and k_2 is proportional to $k_1 k_2$.

When ISPs go bankrupt, their infrastructure is usually acquired by another ISP, which then either merges the ASs or forms a “sibling” relationship in which their routing domains appear independent but are controlled by one umbrella organization. Thus, in terms of the topology, bankruptcy means that a connection shifts from one ISP to another. Since high degree ISPs tend to be wealthier, they are more likely to be involved in such takeovers. We denote the rate of bankruptcy by μ per unit time.

A growing AS may decide to multihome, i.e., to connect to at least two Internet providers. One would expect that higher degree ISPs with a need for reliability would multihome to other higher degree ISPs. We model this phenomenon by assuming that multihoming links appear in the system at rate ν per unit time. The probability that a new link becomes attached to a pair of ISP-nodes of degree k_1 and k_2 is proportional to $k_1 k_2$. The links are directed from the customer to the provider, and we assume that the higher degree ISP is the provider. We also assume that non-ISPs multihome to an average of m providers each. The model is illustrated in Figure 2.

Under this complete MPA model, using techniques similar to that of Section 2, we can show that

$$\gamma = 2 + \frac{1-\mu}{1+2\nu+m\rho+2c+\mu}. \tag{7}$$

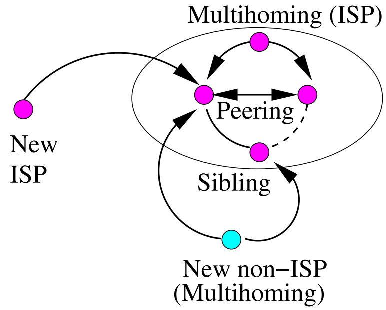

Fig. 2. The multiclass attraction model with multihoming

Validation Against Observed Topology 2. *We used the annotated Route Views data [21] from [22] in order to obtain the empirical distribution of number of ISPs to which ASs multihome. We find that the average number of providers that ISPs multihome to is 2 (i.e.,* $\nu = 2$*) and the average number of providers that non-ISPs multihome to is* 1.86 *(i.e.,* $m = 1.86$*). Dimitropoulos* et. al *[22] also showed that roughly* 90% *of links are of customer-provider type, i.e., these links pertain to transit relationships, with payments always going to the provider ISP. They find (a lower bound of)* 10% *of links are peering, i.e., these links correspond to bilateral traffic exchange without payment. In the model, customer links appear in the system at a rate of* $1 + m\rho$*. We thus calculate* $c = (1 + \nu + m\rho)/9 = 0.704$ *peering links per unit time. The authors of [22] also estimate that the fraction of sibling links is too small to measure accurately and we take* $\mu = 0$*. Substituting these values into the exponent expression (7) results in* $\gamma = 2.114$ *that matches the observed values lying between* 2.1 *and* 2.2 *[1, 19, 20]. However, the salient feature is that the ratio of non-ISPs to ISPs* ρ *is the dominating term and the others are relatively less significant.*

We skip for brevity the derivations of the following two additional predictions of the MPA model:

– The model yields the power-law distributions for the number of peers and customers that ISPs possess. The exponents of these power laws are identical to that of the overall degree distribution γ.
– The distribution of providers of ISPs is a random variable $1 + X$, where X is exponentially distributed with parameter ν.

Validation Against Observed Topology 3. *Using the same data from [22], the CCDF of the number of providers of an ISP (after subtracting the one initial provider) versus its degree is shown in Figure 3, which is plotted in the semi-log scale. The exponential curve fit to the initial part of the graph has a slope of* -0.7*, i.e., the average number of providers is* $1 + 1/0.7 = 2.4$*, which is close to our empirically measured mean value of* 2 *in Validation 2.*

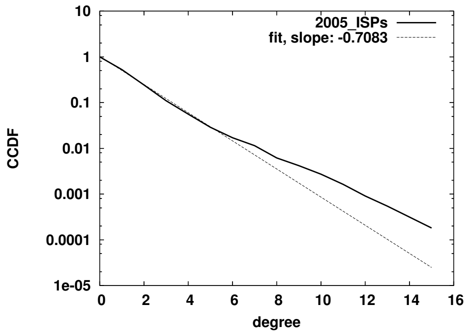

Fig. 3. Illustrating the fact that the provider distribution (after subtracting the one initial provider) for ISPs is close to exponential

We note that the purpose of the last validation is not to show that the distribution is exactly of the form $1 + X$ where $X \sim \exp(1/\nu)$, but to show that it is definitely not a power law.

We make one final observation regarding the MPA model enriched with geographic information. We could divide the world into different geographical regions, each growing at a different rate. Due to the self-similar nature of power-law topologies, the resulting graph would still bear identical properties to the MPA model as long as the parameters ρ, c, and μ are the same in all regions. Evidence supporting this hypothesis is available in [20, 23], where Chinese or European parts of the Internet are shown to have properties similar to the global AS topology.

4 Model Validation by Simulation

We have developed a model that describes the evolution of the AS-level topology and validated the analytical results using measured parameters. We now simulate the MPA model using all of the measured parameters.

The MPA model generates *annotated* graphs, with links being either customer-to-provider (c2p) or peer-to-peer (p2p). Dimitropoulos *et al.* [14] have constructed a definitive set of metrics that are sufficient to capture all the important properties of the annotated AS topology. These metrics are: (i) the degree distribution (DD): the traditional distribution of node degrees; (ii) the annotated distributions (ADs): the distributions of customers, providers, and peers that nodes have; (iii) the annotated degree distribution (ADD): the joint distribution of customers, providers, and peers of nodes, measuring the correlations among the three numbers "at a node;" and (iv) the joint degree distributions (JDDs): the JDDs measure the correlations of node degrees "across the links" of different types. We will compare all these metrics between the graphs that the MPA-model produces and the annotated data set available at [24], which is extracted from BGP tables [21].

The values of the parameters we use in our simulations are $\rho = 2.3$, $\nu = 2$, $c = 0.704$, and $m = 1.86$, which we recall are the ratio of the numbers of non-ISPs to ISPs, ratio of ISP multihoming links to ISPs, ratio of peering links to ISPs, and the average number of providers to which non-ISPs multihome, respectively. In the simulations we do not allow two p2p links between the same ISPs, nor do we allow a peering link between a customer and a provider. Similarly, we do not allow multiple c2p links between the same two ASs. We run the simulation with deterministic link arrivals based on their arrival rates, with a total of $7,200$ ISP nodes. This is the number of ISP nodes that we were able to identify and label from BGP data. We do not model bankruptcy since as mentioned earlier, the rate at which it occurs is too small to get an accurate estimate of our bankruptcy ratio μ.

Degree Distribution (DD): Figure 4(a) shows the DD of the graphs generated by the MPA model, and its comparison with the observed topology. As predicted in the previous section the MPA model produces a power law DD, and the exponent of the CCDF matches well with the BGP plot.

Annotated Distributions (AD): Figures 4(b)–4(d) show the ADs generated by the MPA model. We compare the customer, peer, and provider degree distributions of the simulated graph with that of the BGP tables. As predicted, the ADs of number of customers and peers are both power law graphs with the same exponent as the DD.

We plot the CCDF of the number of providers that ISPs multihome to (on linear x-axis and logarithmic y-axis) in Figure 4(d). They are approximately of form $1+X$, where X is exponentially distributed. The curves show a discrepancy in slope. We believe that it arises due to the fact that almost all the distribution mass is concentrated at small degrees, as the mean is 2, and the number of ISPs with high multihoming degree is small.

Annotated Degree Distribution (ADD): Each ISP has some numbers of providers, peers, and customers. The ADD is the joint distribution of these numbers across all the ISPs. We illustrate these correlations in Figures 4(e) and 4(f). To construct these plots, we first bin the ISP nodes by the number of providers that they have (the x-axis), and then compute the average number of customers or peers that the ISPs in each bin have (the y-axis). We observe that the MPA model approximately matches the BGP data against these metrics as well.

Joint Degree Distributions (JDDs): While the ADD contains information about the correlations between the numbers of different types of nodes connected to an ISP, it does not reveal information about the degree correlations between the parameters of different ISPs connected to each other, i.e., whether higher degree ISPs are more likely to peer with each other, etc. This information is contained in the average neighbor connectivity, which is a summary statistic of the joint degree distributions in Figures 4(g) and 4(h). Specifically, let the probability that a node of degree k has a c2p link to a node of degree k' be called $P_{c2p}(k'|k)$. Then the average degree of the provider ISPs of ISPs that have degree k is $\bar{k}_{c2p}(k) = \sum_{k'=1}^{k_{max}} k' P_{c2p}(k'|k)$. In a full mesh graph with n nodes and undirected links, since all nodes have degree $n-1$, the value of this coefficient is

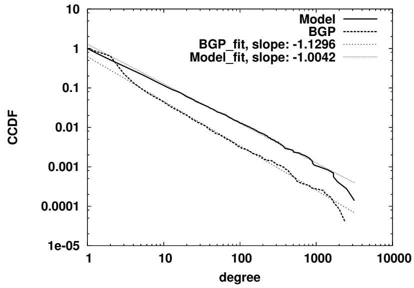

(a) Degree distribution of ASs (CCDF)

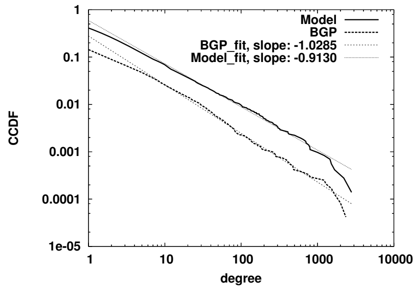

(b) Empirical customer distribution (CCDF)

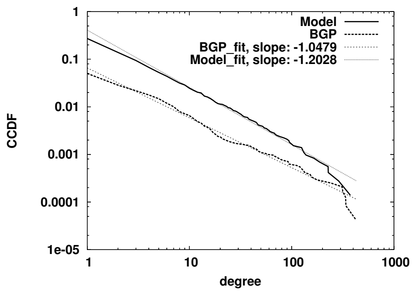

(c) Empirical peer distribution (CCDF)

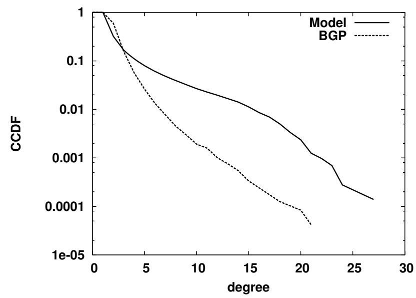

(d) Empirical provider distribution (CCDF)

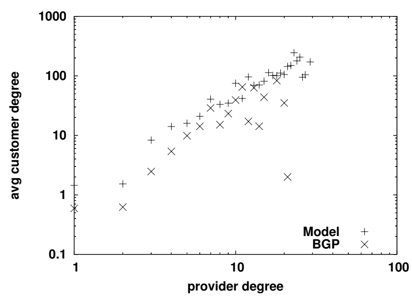

(e) ISP's provider vs. customer degrees

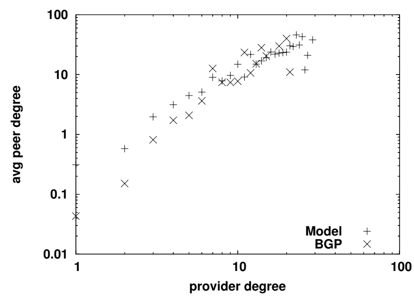

(f) ISP's provider vs. peer degrees

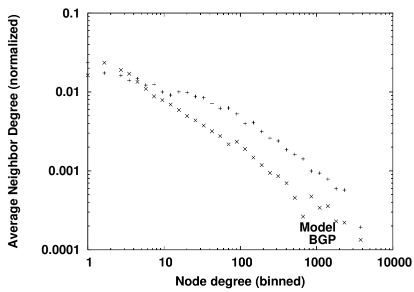

(g) ISP's degree vs. its customers' degree

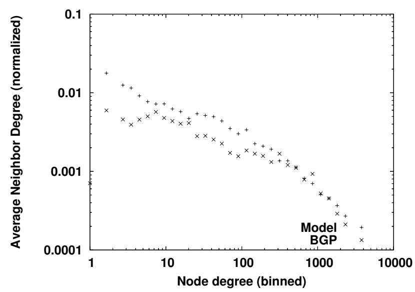

(h) ISP's degree vs. its peers' degree

Fig. 4. Validation of the MPA model by simulation

simply $n-1$. We show the normalized value $k_{c2p}(k)/(n-1)$ in Figure 4(g). The similarly normalized values of $\bar{k}_{p2p}(k)$ are shown in Figure 4(h). These functions exhibit similar behaviors for the MPA model and BGP data.

A good match of the metrics considered in this section, coupled with observations in [13, 14] that AS topologies are accurately captured by their $2K$-annotated distributions, ensures that the MPA model reproduces many other important topological properties of the Internet AS-level topology.

5 Conclusion

We constructed a realistic and analytically tractable model of the Internet AS topology evolution that we call the multiclass preferential attachment (MPA) model. The MPA model is explicitly based on preferential attachment, and we believe it uses the *minimum number of measurable parameters* altering the standard preferential attachment mechanism to produce topologies that are remarkably similar to the real Internet topology. Each model parameter reflects a realistic aspect of the AS dynamics. We measure all the model parameters using the best available AS topology data, substitute them in our derived analytic expressions for the model, and find that it produces topologies that match observed ones against a definitive set of network topology characteristics. These characteristics are derivatives of the second-order degree correlations, annotated with AS business relationships. Matching them ensures that synthetic AS topologies match the real one according to all other important metrics [13, 14].

The model parameter that has the most noticeable effect on the properties of generated topologies, reflects the ratio of ISP to non-ISP ASs. Contrary to common beliefs, other parameters, taking care of AS peering, bankruptcies, multihoming, etc., are significantly less important. No other parameters or complicated mechanisms appear to be needed to explain the Internet topology annotated with AS business relationships. In other words, preferential attachment alone, with the minimal MPA modifications, can almost fully explain the complexity of the economic AS Internet abstracted as an annotated graph. An interesting open question is about the origins of preferential attachment in the Internet. Given that the vast majority of AS links connect customer ASs to their providers [22], this question reduces to finding how customers select their providers. The popularity of providers, their "brand names," may be an important factor explaining the preferential attachment mechanism acting in the Internet.

We conclude with a remark that a majority of drastically different large-scale complex networks (biological, social, economic, WWW, etc.) have topologies similar to the Internet's, suggesting that there might exist unifying yet undiscovered laws driving evolutionary dynamics of all these networks [25, 26, 27, 28]. We believe that the first step toward such laws is development of models that are realistic, parsimonious, analytically tractable, use only measurable parameters, and validate against the best available observations. We hope our work inspires further activity in this direction.

Acknowledgments. We thank Tom Vest for useful discussions. This work was supported by NSF CNS-0434996 and CNS-0722070, DHS N66001-08-C-2029, and by Cisco Systems.

References

1. Faloutsos, M., Faloutsos, P., Faloutsos, C.: On power-law relationships of the Internet topology. In: SIGCOMM 1999: Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication, pp. 251–262. ACM Press, New York (1999)
2. Chang, H., Jamin, S., Willinger, W.: Internet connectivity at the AS-level: An optimization-driven modeling approach. In: MoMeTools 2003: Proceedings of the ACM SIGCOMM workshop on Models, methods and tools for reproducible network research, pp. 33–46. ACM Press, New York (2003)
3. Zhou, S., Mondragón, R.J.: Accurately modeling the Internet topology. Phys. Rev. E. 70, 066108 (2004)
4. Serrano, M.Á., Boguñá, M., Díaz-Guilera, A.: Competition and adaptation in an Internet evolution model. Phys. Rev. Lett. 94, 38701 (2005)
5. Serrano, M.Á., Boguñá, M., Díaz-Guilera, A.: Modeling the Internet. Eur. Phys. J. B 50, 249–254 (2006)
6. Chang, H., Jamin, S., Willinger, W.: To peer or not to peer: Modeling the evolution of the Internet's AS-level topology. In: INFOCOM 2006. 25th IEEE International Conference on Computer Communications, Barcelona, Catalunya, Spain, April 23-29, 2006. IEEE, Los Alamitos (2006)
7. Wang, X., Loguinov, D.: Wealth-based evolution model for the Internet AS-level topology. In: INFOCOM 2006. 25th IEEE International Conference on Computer Communications, Barcelona, Catalunya, Spain, April 23-29, 2006. IEEE, Los Alamitos (2006)
8. D'Souza, R., Borgs, C., Chayes, J., Berger, N., Kleinberg, R.: Emergence of tempered preferential attachment from optimization. Proc. Natl. Acad. Sci. USA 104, 6112–6117 (2007)
9. Corbo, J., Jain, S., Mitzenmacher, M., Parkes, D.: An economically-principled generative model of AS graph connectivity. In: NetEcon+IBC 2007: Proceedings of the ACM Joint Workshop on The Economics of Networked Systems and Incentive-Based Computing. ACM Press, New York (2007)
10. Bar, S., Gonen, M., Wool, A.: A geographic directed preferential Internet topology model. Comput. Netw. 51(14), 4174–4188 (2007)
11. Holme, P., Karlin, J., Forrest, S.: An integrated model of traffic, geography and economy in the Internet. Comput. Commun. Rev. 38(3) (2008)
12. Krioukov, D., Chung, F., claffy, k., Fomenkov, M., Vespignani, A., Willinger, W.: The workshop on Internet topology (WIT) report. Comput. Commun. Rev. 37(1) (2007)
13. Mahadevan, P., Krioukov, D., Fall, K., Vahdat, A.: Systematic topology analysis and generation using degree correlations. In: SIGCOMM 2006: Proceedings of the 2006 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 135–146. ACM Press, New York (2006)
14. Dimitropoulos, X., Krioukov, D., Riley, G., Vahdat, A.: Graph annotations in modeling complex network topologies. ACM Transactions on Modeling and Computer Simulation (to appear, 2009) arXiv:0708.3879

15. Barabási, A.-L., Albert, R.: Emergence of scaling in random networks. Science 286, 509–512 (1999)
16. Dorogovtsev, S.N., Mendes, J.F.F., Samukhin, A.N.: Structure of growing networks with preferential linking. Phys. Rev. Lett. 85(21), 4633–4636 (2000)
17. Krapivsky, P.L., Redner, S., Leyvraz, F.: Connectivity of growing random networks. Phys. Rev. Lett. 85(21), 4629–4632 (2000)
18. Dimitropoulos, X., Krioukov, D., Riley, G., claffy, k.: Revealing the Autonomous System taxonomy: The machine learning approach. In: PAM 2007: Proceedings of the 7th International Conference on Passive and Active Network Measurement, Adelaide, Australia, March 30-31 (2006)
19. Chang, H., Govindan, R., Jamin, S., Shenker, S.J., Willinger, W.: Towards capturing representative AS-level Internet topologies. Comput. Netw. 44(6), 737–755 (2004)
20. Mahadevan, P., Krioukov, D., Fomenkov, M., Huffaker, B., Dimitropoulos, X., claffy, k., Vahdat, A.: The Internet AS-level topology: Three data sources and one definitive metric. Comput. Commun. Rev. 36(1), 17–26 (2006)
21. University of Oregon RouteViews Project, `http://www.routeviews.org/`
22. Dimitropoulos, X., Krioukov, D., Fomenkov, M., Huffaker, B., Hyun, Y., claffy, k., Riley, G.: AS relationships: Inference and validation. Comput. Commun. Rev. 37(1) (2007)
23. Zhou, S., Zhang, G., Zhang, G.: The Chinese Internet AS-level topology. IET Commun. 1(2), 209–214 (2007)
24. CAIDA. AS Relationships Data. Research Project, `http://www.caida.org/data/active/as-relationships/`
25. Albert, R., Barabási, A.-L.: Statistical mechanics of complex networks. Rev. Mod. Phys. 74, 47–97 (2002)
26. Dorogovtsev, S.N., Mendes, J.F.F.: Evolution of Networks: From Biological Nets to the Internet and WWW. Oxford University Press, Oxford (2003)
27. Newman, M.E.J.: The structure and function of complex networks. SIAM Rev. 45(2), 167–256 (2003)
28. Boccaletti, S., Latora, V., Moreno, Y., Chavez, M., Hwanga, D.-U.: Complex networks: Structure and dynamics. Phys. Rep. 424, 175–308 (2006)