Policy-Aware Topologies for Efficient Inter-Domain Routing Evaluations

Yihua He — yhe@cs.ucr.edu
Michalis Faloutsos — michalis@cs.ucr.edu
Srikanth V. Krishnamurthy — krish@cs.ucr.edu
Marek Chrobak — marek@cs.ucr.edu
University of California, Riverside

***Abstract*—The Internet community has not reached a consensus on an appropriate topological model for evaluating the performance of inter-domain routing protocols. Using the current Internet topology is not realistic, since its size is prohibitively large for, say, a packet-level BGP simulation. Furthermore, routing policies, which play a critical role in inter-domain routing, are often ignored in many simulation studies. In this paper, we address this issue by designing an algorithm to generate small-scale, realistic, and policy-aware topologies. We propose HBR, a network sampling method, which produces topologies that preserve the fundamental properties of the Internet graph, including, in particular, its hierarchical structure. Our approach provides a long-term solution to the difficult problem of AS-level routing evaluations: it can be used to generate small realistic topologies in the future, starting from any newer or more complete Internet instance.**

I. Introduction

Inter-domain routing studies need to use simulations, since theoretical analysis and experimentation cannot be easily used for many types of BGP performance evaluations. BGP and the inter-domain interactions are too complex for theoretical analysis, especially when it comes to large scale phenomena. Experimentation, on the other hand, is very cumbersome: replicating a medium-size inter-domain network in a lab is not trivial, and experimenting on the Internet itself is not a welcomed proposition to network operators. As a result, simulations are widely used to test and validate new techniques for BGP improvements [29][28][6][8][34] and to study the behavior and performance of BGP with different parametric settings [17][25].

Our goal is to enable efficient and meaningful inter-domain routing studies, by providing a tool to produce topologies that are (a) sufficiently small, so that simulations can be conducted and repeated in "human" time (e.g. a few days), and, at the same time, (b) faithfully represent the Internet graph, so that the results provide reasonable estimates of the performance in the real world. Currently available methods for topology generation have not, in our view, yet been able to satisfy these two criteria at the same time.

Interestingly, routing policies are not considered in many previous and even recent studies [8][34]. These studies model BGP as a pure path vector routing protocol that chooses the shortest path and each AS always advertises the best (shortest) known route to all of its neighbors. However, this is not a realistic behavior due to routing policies. Not considering these routing policies may lead to inaccurate or unrealistic conclusions. Here, we use the term **policy-aware** to refer to a topolgy that represents routing policies, and use the term **no-policy** to refer to a topology without representing any routing policies. A policy-aware topology has annotated edges, which represent the type of relationship between the connected ASes. For example, directed edges are often used to indicate a provider-customer relationship. Note that a realistic policy-aware topology has to be **BGP-connected**: any two ASes must be able to communicate over a path that does not violate any routing policy.

A key problem is that no *reasonably sized* and *representative* inter-domain topologies are currently available for conducting simulations. The complete Internet topology (25,000 ASes) is too large for any packet-level simulators, such as SSFNET [3]. On top of this, BGP simulations are typically repeated a number of times to test different parameter settings and for ensuring reliability of the results; the total number of runs could exceed 200,000 [17]. For this reason, researchers either only use small canonical topologies [17], or rely on topology generators [21][26][7][24] and sampling approaches [22][30] to produce smaller scale Internet-like topologies. However, none of these approaches produces policy-aware topologies. A recent work [11] proposes an Internet topology generator with AS relationships, but the topology is not guaranteed to be BGP-connected. The lack of small, representative and BGP-connected topologies impedes our ability to evaluate inter-domain routing in an efficient manner.

In this paper, we propose **H**ierarchy-**B**ased **R**eduction (**HBR**), a novel approach to sample the complete Internet topology. A key novelty is that HBR uses and respects the hierarchical structure of the Internet, to produce small policy-aware BGP-connected[1] topologies.

HBR can reduce a topology successfully to 20% of the original size. According to our initial evaluation, HBR can successfully produce small realistic policy-aware topologies that are approximately 20% of the size of the original graph. We validate the realism of our topologies using: (a) an extensive list of graph metrics that characterize the Internet graph, and (b) an actual evaluation of BGP performance, such as BGP convergence time. (Due to space limitations, the latter evaluation is not included here, and it can be found in [18].)

Our work in perspective. The value of our work is not the generation of a *particular topology*, but the development of a *method* to generate topologies. Thus, as the Internet grows or as we measure it more accurately, our approach can always be used to sample the newer, more complete topology.

In addition, our work can be seen as a promising step towards developing a more realistic BGP topology model – a goal that is considered to be ambitious and non-trivial [27]. Further enrichment of the model is left for the future work. For example, since we model the AS-level topology at the level of ASes (i.e., each node represents an AS), intra-AS dynamics

[1]HBR can generate provably BGP-connected topologies, if the initial topology is BGP-connected, but this not discussed here due to space limitations.

or the interactions between iBGP and BGP [13] cannot be captured in our model.

II. Background and related Work

The Internet is composed of tens of thousands Autonomous Systems (ASes). The Border Gateway Protocol (BGP) is the *de facto* routing protocol used to exchange reachability information among these ASes and to interconnect them. Simulations have been widely used to study BGP parameters such as MRAI [17], and Route Flap Damping [25], and to evaluate new inter-domain protocols [29][28][6][8] [34]. However, simulations in all these studies only use no-policy topologies.

Routing policies are commonly implemented in today's Internet. Policies can be thought of as the rules with which an AS accepts, modifies, and advertises further route information (route updates) that it receives from its neighbors. Although an AS may have specific routing policies for each of its neighbor ASes, general policies are normally determined by the business relationships (say, provider-customer or peer-peer) with its neighbor ASes. One common policy for a multi-homed AS, for example, is that it will not advertise the routes learned from one of its providers to its other providers. This is because the multi-homed AS does not want to carry transit traffic between its providers. An AS could also prefer to use and advertise routes learned from its customers to routes learned from its providers, even if the paths via its customers are longer [35].

The business relationships between the ASs can be inferred from global routing tables[14][36][5][9]. Gao et al. [16][15] study and formalize the model of routing policies that is widely used now. Labovitz et al. [23] measure the impact of topology on BGP performance using data from about 200 ISPs, but their goal is not to provide a topology model as we do here.

The challenge of realistic BGP simulations: An Internet-scale BGP simulation is very resource consuming and often impossible. The required memory for detailed BGP simulators, such as [3][2][1], increases cubically with the size of the network [10]. In the most popular BGP simulator SSFNET [3], a simulation on a 1000-AS no-policy topology could consume 2GB memory even if each AS only announces one prefix. C-BGP [32] can perform large scale BGP simulations. However, it only implements the BGP decision process, and does not consider details of the protocol, such as timers and BGP messages. A recent simulator, simBGP [31], can perform large-scale simulations by ignoring the protocol stack below the application layer, but the number of prefixes in each simulation is very limited. Even if memory requirement were not an issue, large simulations would take months for a single run.

III. Sampling Methods

In this section, we tackle the problem of how to produce smaller-scale, while still representative, "Internet-like" topologies with policies.

A. Identifying Required Properties of the Topology

As discussed earlier, when policies are considered, there are two additional properties that a topological model needs to satisfy. First, **the topology must be "BGP-connected"**. The concept of connectivity is extended when routing policies are incorporated: nodes are not only required to be within a connected component, but are also required to reach each other without violating any routing policy. For example, if the only paths between a pair of ASes A and B are through some customers of these two ASes, then the topology is not BGP-connected, as such paths violate the export filtering policy, and there is no other path between A and B. Since any AS in the Internet should be able to reach any other AS, a policy-aware Internet topology should be "BGP-connected". Second, **the topology must be "relationship loop-free"**. For example, ASes A, B and C form a relationship loop if A is B's provider, B is C's provider, and C is A's provider. In reality, relationship loops should not occur; otherwise, BGP is not guaranteed to converge [15].

Obtaining a smaller scale topology is more challenging when we incorporate policies. Most previous efforts [21][26][22][30], in fact, do not consider routing policies. A recent work [11] attempted to generate a topology with AS relationships by enforcing the joint distribution of provider, customer and peer degrees. However, this approach is not quite satisfactory because the resulting topology is not guaranteed to be BGP-connected nor relationship loop-free. We examined an extension of this scheme that trims away BGP-disconnected nodes, but we discovered that such a practice greatly deteriorates the quality of the generated topologies. In fact, in some cases, trimming could make a topology completely degenerate into BGP-disconnected pieces. In addition, a careless assignment of AS relationships could introduce relationship loops.

B. Hierarchy-Based Reduction

We present Hierarchy-Based Reduction (*HBR*), a novel method to sample a large initial topology. We develop several variations of HBR. We first present the basic method (*HBR0*), the operation of which can be described in three stages:

Initialization stage: We first identify and select all ASes that have no providers in the initial topology. As we will see later, these ASes form a clique, which we call **top-clique**, with peer-to-peer links only, if the topology is BGP-connected.

Iterative stage: For each AS selected in the Initialization stage, we select its customers with probability p. This is step 1 in this stage. In the next step, we take the new ASes from step 1 and select their customers with probability p. We repeat the process step by step until an iteration does not select any new ASes.

Assembling stage: We construct the smaller topology by keeping all the links between the selected ASes including peer-to-peer links. Naturally, the relationship reflected by a link in the new graph is the same as that in the initial graph. Algorithm 1 provides a pseudo-code description of our method.

C. HBR variations and uniform reduction

We recognize that HBR0 is not the only way to reduce a large Internet AS topology. Thus, for each of the three reduction stages introduced in Section III-B, we consider possible alternatives, and compare the performance. Obviously this is

Algorithm 1 HBR0 algorithm: $G(V, E) \Rightarrow G_s(V_s, E_s)$

Input: original topology $G(V, E)$, sampling rate $0 < p \leq 1$
Output: smaller topology $G_s(V_s, E_s)$

$TopASes \Leftarrow$ get top clique ASes from $G(V, E)$
$V_s \Leftarrow TopASes$
$CurrentLayerASes \Leftarrow TopASes$
while $CurrentLayerASes$ not empty **do**
  $NextLayerASes \Leftarrow$ empty
  for all AS in $CurentLayerASes$ **do**
    for all $cust$ such that $cust$ is a customer of AS **do**
      if $0 \leq rand() < p$ and $cust$ not in V_s **then**
        add $cust$ into $NextLayerASes$
      end if
    end for
  end for
  $V_s \Leftarrow V_s \bigcup NextLayerASes$
  $CurrentLayerASes \Leftarrow NextLayerASes$
end while
$E_s \Leftarrow$ empty
for all $edge$ in E **do**
  if both nodes at the two ends of the $edge$ is in V_s **then**
    add $edge$ into E_s
  end if
end for
return $G_s(V_s, E_s)$

not an exhaustive list, and the analysis of other modification of HBR0 is left for future work.

HBR1: Here, in the *Initialization stage*, instead of selecting every top-clique AS in the initial topology, we only choose a subset of the tier-1 ASes. The number (s) of the chosen tier-1 ASes has a lower bound $MinS$. s is also related to the initial clique size ($InitS$) and the sampling rate p: $s = MinS + \lceil p * (InitS - MinS) \rceil$. We use $MinS = 1$ in this paper unless otherwise stated. The reason for considering this alternative method is that there are normally fewer tier-1 ASes in a smaller scale Internet instance, as shown from the history of Internet. HBR1 tries to match the corresponding number of tier-1 ASes when it reduces a large topology to a smaller one.

HBR2: Here, in the *Iterative stage*, customers of ASes from an upper tier are considered only once with probability p at a given step. As a result, in this scheme multi-homed ASes have a lower chance of being selected in HBR2 than in HBR0.

HBR3: Here, in the *Assembling stage*, instead of keeping all provider-customer edges among the selected ASes, we only keep the provider-customer edges along which the customer was selected. This variation reduces the number of edges, as well as the number of multi-homed ASes.

For comparison purposes, we present two additional, uniform reduction heuristics: DDRV and DDRE, which are variations of the method in [22] to sample no-policy networks.

DDRV: Directed Deletion of Random Vertex. Remove each AS, independently, with probability $1 - p$, and keep all edges between the remaining ASes. Finally, choose the largest BGP-connected component.

DDRE: Directed Deletion of Random Edges. Remove each edge, independently, with probability $1 - p$. In the end, choose the largest BGP-connected component.

Note that, in order to improve the performance of these approaches, we do not remove the top-clique ASes in DDRV or the edges of the top-clique in DDRE.

IV. Evaluation

There are at least two possible ways to assess the success of a reduced topology: one can either try to match the properties of real Internet instances in history, or try to match the properties of the initial unreduced instance. If the topological properties do not change with size, these two approaches converge. However, as shown in [12] and later in this section, no strict size-independent Internet topology property seems to exist so far. In fact, some properties *have to* change with size. For example, the average path length bewteen the nodes in a grid topology increases as the size of the network increases. Thus, we decide to compare the properties of the sampled topology with the real (historical) Internet topology that has approximately the same size.

We conduct our evaluation using the data from Oregon Routeviews [4]. This is the most frequently used route archival data to infer AS-level Internet topologies. Furthermore, it is the only data archive that has instances dated back to 1997. Although this is not a complete topology, we argue that this is less important in our case: (i) our comparison is consistent, since we start from a large instance of the same data set, and (ii) our sampling does not depend on the completeness of peer-to-peer edges, which are the ones mostly missing from the data [19]. We use snapshots of the Internet topology from Dec 1997 to Dec 2006, a 9-year span during which the size of Internet topology has grown 8-folds, from approximately 3,000 ASes in Dec 1997 to nearly 24,000 ASes in Dec 2006.

To infer the AS relationships of an AS topology, we use the algorithm described in [14]. Other methods are available, but they either require additional/seed data, which is not always available for the topology instances we have from 1997 to 2006 [36][9][33], or do not always work for our topology instances. A common problem with all inference algorithms is that they sometimes produce a topology that is not BGP-connected. We trim the ASes that are not connected to the top-clique of the topology. The number of the trimmed ASes is very small, typically less than 1% of ASes in the original topology.

A. Topological Properties

We compare a number of topological properties between sampled topologies and the real Internet topologies with the same size. We include 6 sampling techniques: HBR0, HBR1, HBR2, HBR3, DDRV and DDRE. For each of the considered techniques, we sample a real Internet topology instance obtained on Dec 1, 2006 from Oregon Routeviews. We vary the sampling rate p from 0.1 to 1 to get smaller scale topologies with different sizes.

Overview of results: We find that **HBR0 and HBR1 are the best sampling methods**. HBR2 performs adequately, but consistently worse than HBR0. Finally, HBR3, DDRV, and DDRE perform significantly worse in most of the metrics.

We now provide the comparison of these methods in detail. Due to space limitations, we cannot show all the metrics that were used or provide intuitive explanations for the results in each case.

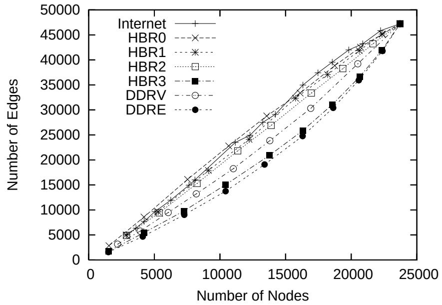

Fig. 1. Number of Edges.

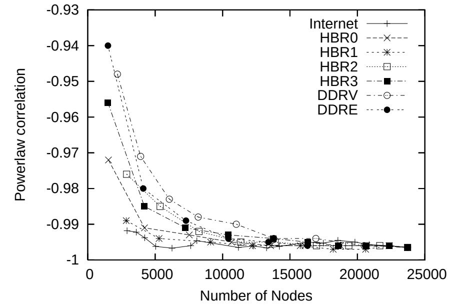

Fig. 2. Power-law Correlation.

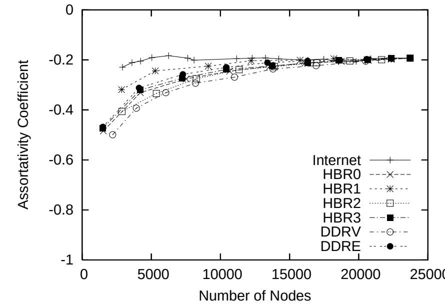

Fig. 3. Assortativity Coefficient.

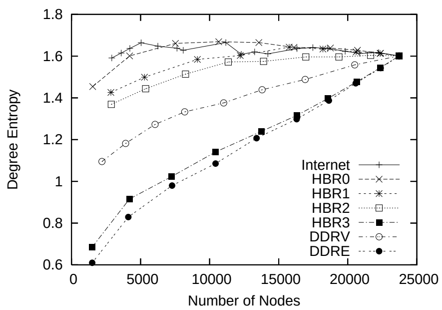

Fig. 4. Degree Entropy.

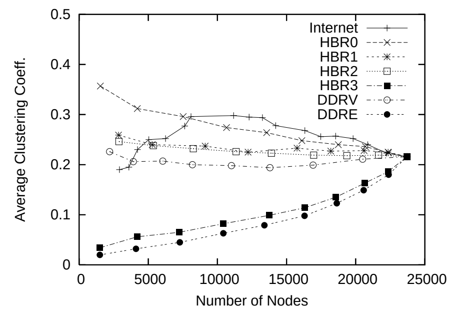

Fig. 5. Average Cluster Coefficient.

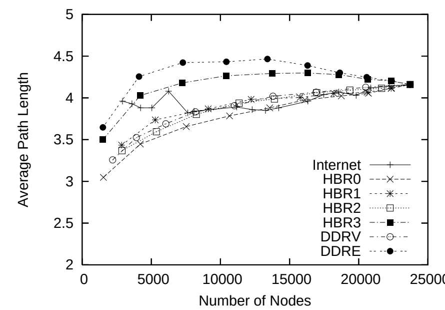

Fig. 6. Average Path Length.

Number of edges. The number of edges in a graph of a given size represents the density of a graph. In Fig. 1, we plot the number of edges against the number of nodes in the topologies. One can see that, throughout the years from 1997 to 2006, the number of edges grows almost linearly with the number of nodes in the historical Internet instances. Reduction methods HBR0, HBR1 and HBR2 follow the Internet data nicely while HBR3, DDRV and DDRE deviate from the evolution of Internet data.

Degree distribution. The degree distribution of the AS-level Internet topology is known to follow a power-law with a correlation coefficient larger than 99% [12], especially if we focus on customer-provider edges [19]. We calculate the power-law correlation coefficient for the complementary cumulative distribution (CCDF) function on the node degrees of each topological instance. In Fig. 2, we see that all Internet instances from the Oregon Routeviews follow power-law degree distributions. Topologies sampled from HBR0 and HBR1 follow the Internet instances very well (≥99%) until the sizes drop to 1/8th of the original's.

Assortativity. The *assortativity coefficient* r of a topology is defined as the Pearson's correlation coefficient of node degrees between all pairs of connected nodes. Intuitively, r captures the tendency of the nodes to attach to nodes with similar (*assortative mixing*, $0 < r \leq 1$) or different degrees (*disassortative mixing*, $-1 \leq r < 0$). In Fig. 3, we plot the r values for all historical Internet instances as well as for the one sampled by our reduction methods. We find that r in the Internet instances is fairly stable at approximately -0.2. Among all the reduction methods, HBR1 works best: the r values from HBR1 graphs follow the Internet values until the size of topology is reduced to 1/8th of the initial size.

Degree entropy. We define the *degree entropy* $\mathcal{H}$ of a topology as $\mathcal{H} = -\sum_k P(k)lnP(k)$, where $P(k)$ is the probability that a randomly selected node has a degree k in this topology. The degree entropy is a measure of the degree randomness of graphs. In Fig. 4, we plot $\mathcal{H}$ values for all topologies. $\mathcal{H}$ for the Internet instances is fairly stable at about 1.6. $\mathcal{H}$ for topologies produced by HBR0 and HBR1 are very stable and close to that from the Internet. On the other hand, HBR3, DDRE and DDRV perform badly as the degree entropy drops sharply in the sampled topologies that they produce.

Average clustering coefficient. We examine the *clustering coefficient* which has been used to characterize and compare generated and real topologies [20]. Intuitively, the clustering coefficient captures how tightly connected is the one-hop neighborhood of a node. For a node v_i with $n_i > 1$ neighbors, the clustering coefficient of v_i is $\gamma_i = \frac{m}{m_{max}}$, where $m_{max} = \frac{n_i(n_i-1)}{2}$, and m is the number of edges between these neighbors. A clustering coefficient of exactly one means that the neighborhood is a clique. The *average clustering coefficient* $\overline{\gamma}$ is the average γ_i of all nodes in the topology. In Fig. 5, we plot the average clustering coefficient against the number of nodes. For Internet instances before 2001 (there were about 8,000 ASes at that time), $\overline{\gamma}$ grows as the size of the topology grows. However, after 2001, $\overline{\gamma}$ slowly decreases. Explaining this observed change in the trend is intriguing but outside the scope of this paper. We limit ourselves to observing that HBR0 follows the most recent (2001 to 2006) trend of Internet the best, although HBR1, HBR2, and DDRV are not far behind.

AS path length. The AS path length d_{AB} from AS A to AS

B in a topology with routing policies is defined as the shortest steady-state AS path length from A to B consistent with the routing policies. Many previous studies only consider the shortest distance without policies, and they may underestimate the AS path length. The average AS path length is the average of d_{AB} for all AS pairs. The number of AS pairs in an n-node topology is $n(n-1)$. Note that in an AS topology with routing policies, d_{AB} is not always the same as d_{BA}. In contrast, d_{AB} is always the same as d_{BA} if no policy is considered. In Fig. 6, we plot the average AS path length for each topology instance. We see that the average distance between ASes is slowly increasing from 2001 when the size of Internet was about 8,000 ASes. The topologies produced by HBR0, HBR1, HBR2 and DDRV follow this trend very well.

V. Conclusion

In this work, we develop a method for conducting feasible and meaningful inter-domain routing evaluations.

We develop the HBR approach, which can reduce a topology successfully to 20% of its initial size. We validate the realism of our sampled topologies using: (a) an extensive list of graph metrics, and (b) actual BGP performance evaluations (the results not shown here due to lack of space).

Our work opens new frontiers in our ability to model the inter-domain topology accurately and evaluate BGP routing effectively. Towards these goals, we intend to release: (a) our HBR tool, and (b) a series of sampled graphs in an effort to establish a badly-needed community-wide simulation benchmark.

On-going and future work. We are in the process of extending the current work along several dimensions [18].

a. What are the effects of policies on BGP routing simulations? We claim that research studies *should* conduct policy-aware simulations. In fact, our initial results suggest that when policies are not considered, BGP simulation results can be significantly different not only in their numerical values but in the fundamental trends.

b. Can we guarantee any properties of the topologies generated by HBR? From a more theoretical standpoint, we want to identify the necessary and sufficient condition for a topology to be BGP-connected. Based on these conditions, we can prove that HBR will produce a BGP-connected topology, if the initial topology is BGP-connected. In addition, we would like to be able to prove that HBR can guarantee the preservation of certain graph metrics. For example, if the initial topology follows a power-law degree distribution, will the sampled topology do too?

c. How much does simulation time decrease when using a sampled topology? A back of the envelope calculation would be easy, if we assume that the simulation running time grows as a cubic or quadratic function of the network size. However, it would be interesting to quantify the benefits in practice and take into consideration other practical parameters such as the complexity of the scenario we study. Our initial results show a decrease in time by a few orders of magnitude.

References

[1] Genesis project, www.genesis-sim.org/genesis.

[2] http://www.ece.gatech.edu/research/labs/MANIACS/BGP++/.

[3] http://www.ssfnet.org.

[4] Oregon routeviews project, http://www.routeviews.org.

[5] G.D. Battista, M. Patrignani, and M. Pizzonia. Computing the types of the relationships between Autonomous Systems, 2003.

[6] A. Bremler-Barr, Y. Afek, and S. Schewarz. Improved BGP convergence via ghost flushing. In *IEEE Infocom*, 2003.

[7] T. Bu and D. Towsley. On distinguishing between Internet power law topology generators. In *IEEE Infocom*, 2002.

[8] J. Chandrashekar, Z. Duan, Z. Zhang, and J. Krasky. Limiting path exploration in BGP. In *IEEE Infocom*, 2005.

[9] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley. AS relationships: inference and validation. *ACM SIGCOMM Computer Communication Review (CCR)*, January 2007.

[10] X. Dimitropoulos and G. Riley. Efficient large-scale BGP simulations. *Computer Networks*, (12):2013–2027, 2006.

[11] X. Dimitropoulos and G. Riley. Modeling Autonomous-System relationships. In *Workshop on Principles of Advanced and Distributed Simulation*, 2006.

[12] M. Faloutsos, P. Faloutsos, C. Faloutsos, and G. Siganos. Power-laws of the Internet topology. *IEEE/ACM Trans. on Networking*, 1(4):514–524, 2003.

[13] N. Feamster, H. Balakrishnan, and J. Rexford. Some foundational problems in interdomain routing. In *Hotnet*, 2004.

[14] L. Gao. On inferring autonomous system relationships in the Internet. In *IEEE Global Internet*, November 2000.

[15] L. Gao, T. Griffin, and J. Rexford. Inherently safe backup routing with BGP. In *IEEE INFOCOM*, 2001.

[16] L. Gao and J. Rexford. Stable Internet routing without global coordination. In *ACM Sigmetrics*, 2000.

[17] T. Griffin and B. Premore. An Experimental analysis of BGP convergence time. In *IEEE ICNP*, 2001.

[18] Y. He, M. Faloutsos, S. V. Krishnamurthy, and M. Chrobak. Generating policy-aware topologies for AS-level Internet simulations. Technical report, Univ. of California, Dept. of CSE, 2008.

[19] Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy. A systematic framework for unearthing the missing links: measurements and impact. In *USENIX NSDI*, 2007.

[20] S. Jaiswal, A. Rosenberg, and D. Towsley. Comparing the structure of power law graphs and the Internet AS graph. In *ICNP*, 2004.

[21] C. Jin, Q. Chen, and S. Jamin. Inet: Internet topology generator. Technical report, Univ. of Michigan, 2000.

[22] V. Krishnamurthy, M. Faloutsos, M. Chrobak, L. Lao, J-H. Cui, and A.G. Percus. Reducing large Internet topologies for faster simulations. In *Networking*, 2005.

[23] C. Labovitz, A. Ahuja, R. Wattenhofer, and S. Venkatachary. The impact of Internet policy and topology on delayed routing convergence. In *IEEE Infocom*, 2001.

[24] P. Mahadevan, D. Krioukov, K. Fall, and A. Vahdat. Systematic topology analysis and generation using degree correlations. In *ACM Sigcom*, 2006.

[25] Z. Mao, R. Govindan, G. Varghese, and R. Katz. Route flap damping exacerbates Internet routing convergence. In *ACM Sigcomm*, 2002.

[26] A. Medina, A. Lakhina, I. Matta, and J. Byers. BRITE: An approach to universal topology generation. *MASCOTS*, 2001.

[27] W. Muhlbauer, A. Feldmann, O. Maennel, M Roughan, and Steve Uhlig. Building an AS-topology model that captures route diversity. In *ACM Sigcomm*, 2006.

[28] D. Pei, M. Azuma, D. Massey, and L. Zhang. BGP-RCN: Improving BGP convergence through root cause notification. *Computer Networks*, 48(2):175–194, 2005.

[29] D. Pei, X. Zhao, L. Wang, D. Massey, A. Mankin, F. Wu, and L. Zhang. Improving BGP convergence through assertions approach. In *IEEE Infocom*, 2002.

[30] BJ Premore. http://www.ssfnet.org/Exchange/gallery/asgraph/index.html.

[31] J. Qiu. http://www.bgpvista.com/simbgp.php.

[32] B. Quoitin and S. Uhlig. Modeling the routing of an Autonomous System with C-BGP. *IEEE Networks*, 19(6), 2005.

[33] L. Subramanian, S. Agarwal, J. Rexford, and R. Katz. Characterizing the internet hierarchy from multiple vantage points. In *IEEE Infocom*, 2002.

[34] W. Sun, Z. Mao, and K. Shin. Differentiated BGP update processing for improved routing convergence. In *IEEE ICNP*, 2006.

[35] F. Wang and L. Gao. Inferring and characterizing Internet routing policies. In *ACM IMW*, 2003.

[36] J. Xia and L. Gao. On the evaluation of as relationship inferences. In *IEEE Globecom*, November 2004.