

On the Scalability of BGP: the Role of Topology Growth

Ahmed Elmokashfi, *Member, IEEE*, Amund Kvalbein, *Member, IEEE*, and Constantine Dovrolis, *Member, IEEE*

Abstract—The scalability of BGP routing is a major concern for the Internet community. Scalability is an issue in two different aspects: increasing routing table size, and increasing rate of BGP updates. In this paper, we focus on the latter. Our objective is to characterize the churn increase experienced by ASes in different levels of the Internet hierarchy as the network grows. We look at several “what-if” growth scenarios that are either plausible directions in the evolution of the Internet or educational corner cases, and investigate their scalability implications and interaction with different failure types. Our findings explain the dramatically different impact of multihoming and peering on BGP scalability, highlight negative and positive effects of multihoming on churn and reachability, and identify which topological growth scenarios will lead to faster churn increase for different failure types.

Index Terms—Internetworking, Routing, Topology, BGP

I. INTRODUCTION

Recently, there is a significant concern among both Internet operators and researchers about the scalability of interdomain routing with BGP. A workshop organized by the Internet Architecture Board concluded that “*routing scalability is the most important problem facing the Internet today*” [24]. The concern is that we are soon approaching the point where the global routing system, and the core routers in particular, will no longer be able to keep up with routing dynamics. BGP scalability is an issue in two different aspects: *increasing routing table size, and increasing rate of BGP updates (churn)*. Note that, in general, an increase in the routing table size (number of routable prefixes) also increases churn, since the number of networks that can fail or trigger a route change increases. In this paper, *we focus on the issue of increasing churn*.

The goal of this study is to improve our understanding of the underlying reasons for the experienced growth in churn. Churn is a result of a complex interplay of 1) the *routing protocol*, including policy annotations and various BGP mechanisms like update rate limiting, route flap dampening etc. 2) *events* like prefix announcements, link failures, session resets, traffic engineering operations that generate routing updates, and 3) the characteristics of the Internet *topology*. The last factor,

in particular, is the primary focus of this paper. We aim to understand how topological characteristics of the AS-level graph influence the scalability of BGP churn.

Describing the AS-level Internet topology and how it evolves has been the subject of much research (and heated debate) in the last decade. We refer the reader to the following representative references [7], [20], [22], [29], [33]. In this paper, we do not use an existing topology generation model because we want to explore a wide range of “what-if” possibilities that none of the existing models captures in a parsimonious and intuitive manner. For the same reasons, we do not base our investigations on inferred historical internet topologies. Instead, we first identify four basic but fundamental characteristics of the Internet graph that have persisted over the last decade. Then, we design a simple and controllable topology generator that satisfies the previous properties, and at the same time allows us to easily navigate the topological space. The “knobs” of this generator are parameters with operational relevance in practice, such as the multihoming degree (MHD) of stubs versus transit providers, instead of abstract measures such as betweenness or assortativity.

Using our topology generator, we establish the factors that determine churn at different locations in the Internet hierarchy, and investigate the importance of each factor in a growth model that resembles the evolution of the Internet over the last decade. We then examine several deviations from this growth model, and investigate how the number of routing updates generated by different routing events grows with the size of the topology in each case. We ask questions such as: “What if the MHD of stub ASes increases with the network size instead of staying constant?” “What if the Internet becomes denser mostly due to peering links?” “What if tier-1 providers dominate the transit market, reducing the number of tier-2 providers?” We examine thoroughly two different routing event types that take place at the edge of the network.

The rest of the paper is organized as follows. In the next section, we explain our overall approach and describe the model that we base our investigation on. In Sec. III, we describe our topology generator and present our Baseline growth model. In Sec. IV, we present a model for the churn experienced at different locations in the Internet, and discuss the importance of different factors as the network grows. In Sec. V and Sec. VI, we examine several topology growth scenarios and investigate how they affect BGP churn generated by different event types that take place at the edge of the network. We review related work in Sec. VII, and draw our conclusions in Sec. VIII.

Manuscript received November 8, 2009; revised May 1, 2010.

This paper was presented in part at the ACM International Conference on emerging Networking EXperiments and Technologies (CoNEXT), Madrid, Spain, 2008.

A. Elmokashfi is with Simula Research Laboratory, 1364 Fornebu, Norway e-mail: (ahmed@simula.no).

A. Kvalbein is with Simula Research Laboratory, 1364 Fornebu, Norway e-mail: (amundk@simula.no).

C. Dovrolis is with the College of Computing, Georgia Institute of Technology, Atlanta, GA 30332 USA, e-mail: (dovrolis@cc.gatech.edu).




Fig. 1: Model for a node representing an AS.

II. APPROACH AND MODEL

We can only study the problems described above using simulations. Since our goal is to look at scalability under different hypothetical topology growth models, our investigation cannot be performed by doing measurements in the current Internet. Also, the complexity of BGP and large Internet-like topologies make it difficult to create a tractable and useful mathematical model. Such modeling has been attempted before, but only for regular topologies, and without taking the effects of MRAI into account [32]. Simulations of any system of the size and complexity of interdomain routing require to make several simplifying assumptions in our model. In this section we describe the choices and assumptions we make, and argue why the resulting model captures the effects we want to investigate in our study.

We study different growth models of the AS-level topology of the Internet, using our topology generator described in Sec. III. In order to do this in a scalable way, *we model each AS as a single node*, and *connections between two neighboring ASes as a single logical link*. This implies that we do not capture routing effects within an AS, introduced by iBGP or interactions with IGP routing protocols (e.g., hot-potato routing). However, while such effects do have an impact on how routing updates are generated, they are orthogonal to the effects we want to study.

We look at two different events that generate BGP updates. First, we focus on events where individual destination prefixes are withdrawn and then re-announced by the owner. This is the most basic routing event that can take place in the Internet, and at the same time the most radical; these changes must be communicated all over the network. Second, we study events where a single link connecting a stub AS to one of its providers fails and is restored. For these two event types, *we measure the number of routing updates received by nodes at different locations in the network*.

Figure 1 shows the structure of a node in our simulation model. A node exchanges routing messages with its neighbors. Incoming messages are placed in a FIFO queue and processed sequentially by a single processor. The time it takes to process an update message is uniformly distributed between 0 and 100 ms. Each node maintains a table with the routes learned from each neighbor. Upon receiving an update from a neighbor, a node will update this table, and re-run its decision process to select a new best route. The new preferred route is then installed in the forwarding table and announced to its neighbors. For each neighbor, we maintain an export filter that blocks the propagation of some updates according to the policies installed in the network. Outgoing messages are stored in an

output queue until the MRAI timer for that queue expires. If a queued update becomes invalid by a new update, the former is removed from the output queue.

For our study, we need a simulator that is capable of capturing the exchange of routing updates described above, and that scales to network sizes of thousands of nodes. Existing interdomain routing simulators fall into two broad categories. Either they only calculate steady state routes, and do not capture routing dynamics [26], or they include a detailed model of each eBGP session, and hence do not scale to network sizes in the order of today’s AS-level Internet topology of about 27000 nodes [2], [10]. Because of this, we have chosen to develop a new simulation model that suits our requirements¹. Using this simulator, we are able to efficiently simulate networks up to about 10000 nodes.

We consider policy-based routing, with the use of MRAI timers to limit the frequency with which a node sends updates to a neighbor. By “policies”, we refer to a configuration where relationships between neighboring ASes are either peer-to-peer or customer-provider. We use normal “no-valley” and “prefer-customer” policies. Routes learned from customers are announced to all neighbors, while routes learned from peers or providers are only announced to customers. A node prefers a route learned from a customer over a route learned from a peer, over a route learned from a provider. Ties among routes with the same local preference are broken by selecting the route with the shortest AS path, then based on a hashed value of the node IDs.

By “MRAI” or “rate-limiting”, we refer to a configuration where two route announcements from an AS to the same neighbor must be separated in time by at least one MRAI timer interval. We use a default MRAI timer value of 30 seconds. To avoid synchronization, we jitter the timer as specified in the BGP-4 standard. According to the BGP-4 standard [27], the MRAI timer should be implemented on a per-prefix basis. However, for efficiency reasons, router vendors typically implement it on a per-interface basis. We adopt this approach in our model. We follow the MRAI implementation recommended in the most recent RFC (RFC4271) [27], which specifies that both announcements and explicit withdrawals should be rate-limited.

III. CONTROLLABLE TOPOLOGIES

In this section, we first describe some key properties that characterize the AS-level Internet topology. We believe that these properties will remain valid in the future. We then describe a model that allows us to construct topologies with different configurable properties while still capturing these key properties.

Most existing topology generators are not capable of producing topologies annotated with business relations, which are essential in our study. Those who are [9], [15], do not have the flexibility we need for controlling different topological characteristics. It is possible to infer historical Internet topologies from routing update traces [12], but it is well known that such

¹The simulator code and the scripts used to generate the results in this paper is available at <http://simula.no/research/networks/software>

inference tends to underestimate the number of peering links, and it is difficult to infer topologies of a tractable size that are representative of today's Internet. We therefore implement our own topology generator.

The input parameters to our generator have “operational” semantics. Instead of specifying abstract graph properties like the clustering coefficient, the betweenness or the assortativity of the topology, we define our topology in a more hands-on, real-world related manner by specifying parameters like how many providers an AS has, how likely it is to peer with other types of ASes etc.

A. Stable topological properties

The AS-level Internet topology is far from a random graph. Over the past decade it has experienced tremendous growth, but the following key characteristics have remained constant:

Hierarchical structure. On a large scale, the nodes in the Internet graph form a hierarchical structure. By hierarchical we mean that customer-provider relationships are formed so that there are normally no provider loops, where A is the provider of B who is the provider of C who again is the provider of A.

Power-law degree distribution. The degree distribution in the Internet topology has been shown to follow a truncated power-law, with few very well-connected nodes, while the majority of nodes have only few connections [11]. The well connected nodes typically reside at the top of the hierarchy.

Strong clustering. The nodes in the Internet are grouped together in clusters, with nodes in the same cluster more likely to be connected to each other. One reason for this clustering is that networks operate in different geographical areas.

Constant average path length. Recent measurements show that in spite of a tremendous growth in the number of nodes, the AS-level path length has stayed virtually constant at about 4 hops for the last 10 years [8].

B. Topology generator

Next, we describe a flexible model for generating topologies that captures the above properties about the AS-level graph. Several design choices and parameters in our topology generator were guided by a recent measurement study [8].

We use four types of nodes in our model. At the top of the hierarchy are the tier-1 (T) nodes. T nodes do not have providers, and all T nodes are connected in a clique using peering links. Below the T nodes, we have the mid-level (M) nodes. All M nodes have one or more providers, which can be either T nodes or other M nodes. In addition, M nodes can have peering links with other M nodes. At the bottom of the hierarchy, we have two different types of stub nodes. We distinguish between customer networks (C) and content providers (CP). In this context, CP nodes would include content provider networks, but also networks providing Internet access or hosting services to non-BGP speaking customers. In our model, the difference between C and CP nodes is that only CP nodes can enter peering agreements with M nodes or CP nodes, while C nodes do not have peering links.

To capture clustering in our model, we introduce the notion of *regions*. The purpose of regions is to model geographical

	Meaning	Baseline value
n	Total number of nodes	1000 – 10000
n_T	Number of T nodes	4 – 6
n_M	Number of M nodes	$0.15n$
n_{CP}	Number of CP nodes	$0.05n$
n_C	Number of C nodes	$0.80n$
d_M	Avg M node MHD	$2 + 2.5n/10000$
d_{CP}	Avg CP node MHD	$2 + 1.5n/10000$
d_C	Avg C node MHD	$1 + 5n/100000$
p_M	Avg M-M peering degree	$1 + 2n/10000$
p_{CP-M}	Avg CP-M peering degree	$0.2 + 2n/10000$
p_{CP-CP}	Avg CP-CP peering degree	$0.05 + 5n/100000$
t_M	Prob. that M's provider is T	0.375
t_{CP}	Prob. that CP's provider is T	0.375
t_C	Prob. that C's provider is T	0.125

TABLE I: Topology parameters

constraints; networks that are only present in one region are not allowed to connect with networks that are not present in the same region. In our model T nodes are present in all regions. 20% of M nodes and 5% of CP nodes are present in two regions, the rest are present in only one region. C nodes are only present in one region.

We generate topologies top-down in two steps. First we add nodes and transit links, then we add peering links. The input parameters n_T , n_M , n_{CP} and n_C decide how many of the n nodes belong to each node type, respectively. First, we create the clique of T nodes. Next, we add M nodes one at a time. Each M node connects to an average of d_M providers, uniformly distributed between one and twice the specified average. M nodes can have providers among both T and M nodes, and we use a parameter t_M to decide the fraction of providers that are T node. M nodes can only select providers that are present in the same region.

M nodes select their providers using preferential attachment, which gives a power-law degree distribution [4].

We then add the CP and C nodes, which have an average number of providers d_{CP} or d_C , respectively. CP and C nodes can select T nodes as providers with a probability t_{CP} and t_C , respectively. Just like the M nodes, C and CP nodes select their providers using preferential attachment.

When all nodes have been added to the topology, we add peering links. We start by adding p_M peering links to each M node. As for the provider links, p_M is uniformly distributed between zero and twice the specified average. M nodes select their peers using preferential attachment, considering only the peering degree of each potential peer. Each CP node adds p_{CP-M} peering links terminating at M nodes, and p_{CP-CP} peering links terminating at other CP nodes. CP nodes select their peers among nodes in the same region with uniform probability. Importantly, we enforce the invariant that a node not peer with another node in its customer tree. Such peering would prey on the revenue the node gets from its customer traffic, and hence such peering agreements are not likely in practice.

C. Baseline growth scenario

Next, we define a Baseline growth model that will later be used as a reference scenario for looking at how different topological factors influence BGP churn. Our aim is to look

at the scalability of different hypothetical growth models, and it is not our goal that the Baseline model should be an exact copy of the historical Internet. Still, the parameters used are inspired by recent measurements of the evolution of the Internet topology over the last decade [8]. The Baseline growth model is characterized by a slow increase in the MHD of stub nodes, and a faster growth in the MHD of middle nodes and the number of peering links. In the Baseline topology we use 5 regions, containing one fifth of all nodes each. Table I gives the parameter values for the Baseline growth model.

Before looking at the churn characteristics of the Baseline model, we validate that the generated topologies capture the four stable properties of the Internet topology discussed above. We compare some properties of the Baseline model to inferred Internet topologies. We look at two Baseline topologies of sizes 5000 and 10000 nodes respectively, and compare against two inferred AS-level topologies of sizes 3247 and 17446 nodes. The smaller topology is provided by Dhamdhare and Dovrolis [8] and it is based on RouteViews [1] and RIPE [28] BGP routing tables from January to March 1998. The second inferred topology is provided by Mahadevan et al. [23] and based on RouteViews BGP routing tables from March 2004. Note that the inferred topologies miss a large fraction of peering links, which distorts their characteristics quantitatively [8]. Therefore, our aim is that the Baseline model matches the major topological properties of the Internet qualitatively rather than quantitatively.

Hierarchical structure. This is trivially fulfilled through the way we construct the topologies.

Power-law degree distribution. The top plot in Fig. 2 shows the CCDF of the node degree on a log-log scale. We observe that our Baseline growth model captures the power-law scaling of the node degrees reasonably well, and is comparable to that of the inferred Internet topologies. The use of preferential attachment when selecting which nodes to connect to gives the observed power-law degree distribution [3].

Strong clustering. We measure the local clustering (or clustering coefficient) of each node in a topology. The local clustering of a node is defined as the ratio of the number of links between that node’s neighbors to the maximum possible number of such links (i.e. a full clique). Hence, the local clustering measures how well connected a nodes neighborhood is. The middle plot in Fig. 2 reports the average local clustering, across all nodes of the same degree, as a function of node degree. To keep the figure readable, we plot results for only two topologies (the other pair of topologies show similar results). Our Baseline growth model matches qualitatively the trends seen in the inferred topologies: first, local clustering decreases with the node’s degree, and second, the clustering versus degree relation follows a power-law. It should be noted however that the Baseline model produces lower clustering than the inferred Internet topologies.

Constant average path length. The average path length in our Baseline topologies is constant at around four hops as the network grows from 1000 nodes to 10000 nodes. This matches closely the average path length in the inferred Internet topology at least since 1998 [8].

In addition to confirming that the Baseline topologies cap-

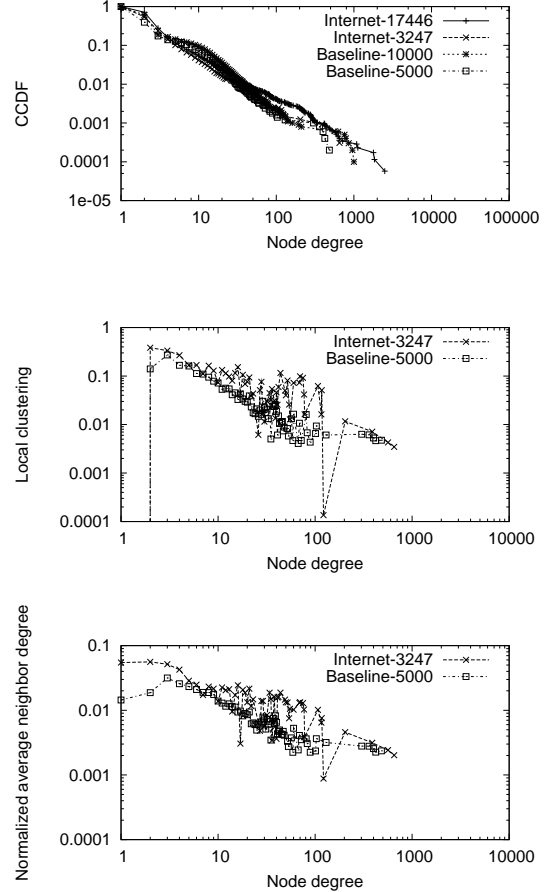


Fig. 2: Node degree distribution (top), Local clustering (middle), Normalized average neighbor connectivity (bottom)

ture the previous four stable properties, we also investigate the *average neighbor connectivity* [23], which has been difficult to capture by existing topology generators [14]. The average neighbor connectivity of a node is simply the average degree of its neighbors. This metric relates to the assortativity of a graph. It measures whether a node of a certain degree prefers to connect with higher or lower degree nodes. The bottom plot in Fig. 2 shows the average neighbor connectivity as a function of the node degree. We normalize the average neighbor connectivity by the maximum possible value which is (the total number of nodes in the graph - 1), in order to compare topologies of different sizes. The Baseline growth model gives an average neighbor connectivity that matches well the inferred Internet topologies, with smaller degree nodes having a higher average local connectivity than the higher degree nodes (referred to as negative assortativity).

IV. EXPLAINING CHURN IN A GROWING NETWORK

In this section, we first present our analytical model describing the number of updates received at a node. Then we use the Baseline growth model to show how this model can be simplified for the different node types, and to determine the most important factors driving the churn growth.




Fig. 3: Illustration of network based on our topology model.

Our main metric in this section is the number of updates received at a node after withdrawing a prefix from a C-type node, letting the network converge, and then re-announcing the prefix again. The experiment is repeated for 100 different C nodes (increasing this number does not change the results), and the number of received updates is measured at every node in the network. We then average over all nodes of a given type, and report this average. In the following, we refer to this procedure as a ‘‘C-event’’. Note that due to the heavy-tailed node degree distribution, we expect a significant variation in the churn experienced across nodes of the same type. We return to the other event type in Sec. VI.

A. A framework for update analysis

We give a formulation for the number of updates received at a node, and discuss how the different churn increase factors depend on the use of policies, the topological properties of the network, and the convergence properties of the routing protocol used.

Figure 3 shows a generic network of the type described in Sec. III. Transit links are represented as solid lines, while peer-to-peer links are dotted. For each node, we have indicated the preferred path to the *event originator* Z , which is the node announcing the active prefix. The routing updates that establish these paths flow in the opposite direction. We observe that due to the use of policies, updates (and the resulting paths) will follow a particular pattern: a node N will only announce a route to its providers and peers after an event at node Z if N has Z in its customer tree. On the other hand, N will always send an update to its customers, unless its preferred path to Z goes through the customer itself.

Let $U(X)$ denote the number of updates a node of type X receives after a C-event. X can be either of the four node types in our model; T, M, CP or C. We distinguish between the number of updates received from customers $U_c(X)$, peers $U_p(X)$ and providers $U_d(X)$ respectively. The total number of updates will be the sum of these: $U(X) = U_c(X) + U_p(X) + U_d(X)$. Each of these values will depend on three factors - the number $\mathbf{m}_{y,X}$ of direct neighbors of a given business relation y , the fraction $\mathbf{q}_{y,X}$ of these neighbors that sends updates during convergence, and the number of updates $\mathbf{e}_{y,X}$ each of these neighbors contribute. The expected number of updates from a certain class of neighbors will be the product of these three factors, and we can write

$$U(X) = \mathbf{m}_{c,X} \mathbf{q}_{c,X} \mathbf{e}_{c,X} + \mathbf{m}_{p,X} \mathbf{q}_{p,X} \mathbf{e}_{p,X} + \mathbf{m}_{d,X} \mathbf{q}_{d,X} \mathbf{e}_{d,X} \quad (1)$$

Note that for some node types, some of these terms will be 0, e.g., T nodes have no providers, and stub nodes have no customers. In the sequel, we will discuss how each of these factors depend on various topological characteristics and their interactions with properties of the routing protocol.




Fig. 4: Number of updates received at T, M, CP and C nodes.





Fig. 5: Number of updates received from peers and customers for T nodes (top), and from providers, peers and customers for M nodes (bottom).

B. Churn at different node types

We focus our discussion on the churn experienced by transit providers (T and M nodes), and content providers (CP nodes). These are the AS types that are most likely to be affected by increasing churn rates, since they must maintain larger routing tables with few or no default routes. Also, as seen in Fig. 4, these are the nodes that experience the stronger growth in the number of updates received after a C-event.

We have calculated 95% confidence intervals for the values shown in Fig. 4, and they are too narrow to be shown in the graph. This tells us that increasing the number of event originators beyond the 100 used in this experiment will not reduce the observed variance. This variance is a result of the often significant differences between topology instances of different size, caused by the heavy-tailed node degree distribution.

T nodes have no providers, so we have $U(T) = U_p(T) + U_c(T) = \mathbf{m}_{p,T} \mathbf{q}_{p,T} \mathbf{e}_{p,T} + \mathbf{m}_{c,T} \mathbf{q}_{c,T} \mathbf{e}_{c,T}$. The top panel in Fig. 5 shows $U_c(T)$ and $U_p(T)$, for topologies of increasing size created with our Baseline topology model. We observe that both $U_c(T)$ and $U_p(T)$ increase with network size, and that both these factors contribute significantly to the total number of updates. As the network grows, the increased multihoming increases the number of routes that a T node learns from both its customers and peers. $U_p(T)$ is the larger factor for small network sizes, and it grows approximately




Fig. 6: Relative increase in $U_c(T)$, $U_p(T)$ and $U_d(M)$.

linearly with network size, with a coefficient of determination $R^2 = 0.93$. The strongest growth is seen in $U_c(T)$, which dominates for larger network sizes. Regression analysis shows that the growth of $U_c(T)$ is *quadratic*, with a coefficient of determination $R^2 = 0.95$.

While routes are only exported to peers and providers if they are received from a customer, routes are always exported to customers. As we can see in the bottom panel of Fig. 5, M nodes receive the large majority of their updates from their providers. Hence, a good estimate for the number of updates at M nodes is $U(M) = U_d(M) = \mathbf{m}_{d,M} \mathbf{q}_{d,M} \mathbf{e}_{d,M}$. The intuition behind this is that M nodes reach the “main part of the Internet” through their providers, and hence also receive the majority of routing updates from them. This is a major simplification, that makes our analysis much simpler. The same is true for CP nodes, so we limit our discussion to M nodes in the following.

Figure 6 shows the increase ratio in $U_c(T)$, $U_p(T)$ and $U_d(M)$. Each term is normalized so that the number of updates is 1 for $n = 1000$. To explain the observed trends for these terms, we look at the different factors described in Eq. 1 to find out how much of the growth is caused by each of them.

First, we look at the increase in the number of neighbors of different types. Figure 7 (top) shows the relative increase in the $\mathbf{m}_{c,T}$, $\mathbf{m}_{p,T}$ and $\mathbf{m}_{d,M}$ factors as the network grows. $\mathbf{m}_{c,T}$ grows much faster than the other factors. With our Baseline topology growth model, $\mathbf{m}_{c,T}$ grows approximately linearly with n in the range of network sizes we consider. The number of peers $\mathbf{m}_{p,T}$ is given directly by $n_T - 1$, which grows very slowly with n . Similarly, $\mathbf{m}_{d,M}$ is determined by the MHD of M nodes $d_M = 2 + 2.5n/10000$, which also grows linearly with n .

The middle panel in Fig. 7 shows the relative increase in $\mathbf{e}_{c,T}$, $\mathbf{e}_{p,T}$ and $\mathbf{e}_{d,M}$, representing the average number of updates received from each neighbor of a given type that exports a route.

The increase in the e factors we see here is caused by path exploration. The increase is stronger for $\mathbf{e}_{p,T}$ and $\mathbf{e}_{d,M}$, since they represent links that are further away from the event originator, giving more chances for more paths to be explored during convergence.

We also see how the increase in the number of received updates is stronger from neighbors that have a larger number of policy-compliant paths that can be explored. The number of valid paths from a T node to an event originator increases superlinearly, which also causes a superlinear growth in $\mathbf{e}_{p,T}$, while the slower growth in the number of paths exported by

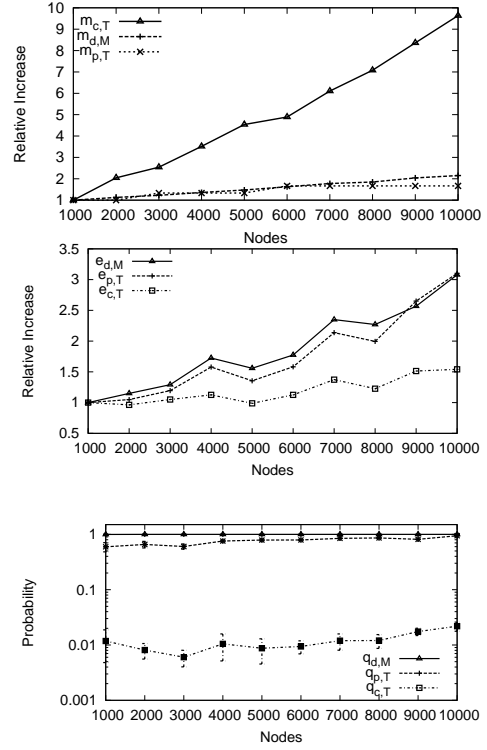


Fig. 7: Relative increase in the factors that determine the update growth rate.

customers gives a slower growth in $\mathbf{e}_{c,T}$. This is compliant with discussion of path exploration in [17].

The bottom panel in Fig. 7 shows the fraction of neighbors of a given type that announces a route after a C-event, represented by the $\mathbf{q}_{c,T}$, $\mathbf{q}_{p,T}$ and $\mathbf{q}_{d,M}$ factors. A provider will always announce a route to its customer, unless it prefers the path through the customer itself. Hence $\mathbf{q}_{d,M}$ is almost constant, and always larger than 0.99. $\mathbf{q}_{c,T}$ and $\mathbf{q}_{p,T}$ are both generally increasing with network size. This illustrates how *increased multihoming makes it increasingly likely that the event originator is in the customer tree of a given customer or peer of a T node*. This probability is much higher for peers than for customers of T nodes, since the peers, which are T nodes themselves, have a much larger number of nodes in their customer tree.

To sum up our discussion, we have shown that the churn at M nodes is dominated by the updates received from providers. The number of updates $U_d(M)$ grows with network size, since both the number of providers $\mathbf{m}_{d,M}$ and the average number of updates $\mathbf{e}_{d,M}$ received from each active provider grows, while the probability that a provider will announce a path is constant. The growth in $U_d(M)$ (a factor 6.7 in our range of $n = 1000$ to $n = 10000$, as seen in Fig. 6) is dominated by the growth in $\mathbf{e}_{d,M}$ (factor 3.1) and the (linear) growth in the MHD (a factor 2.2), which makes the total growth seem slightly superlinear.

For T nodes, both the updates $U_c(T)$ received from customers and the updates $U_p(T)$ received from peers are important, and both grow with network size. The strongest growth is contributed by $U_c(T)$, with a factor 27. Much of this growth

can be attributed to a strong linear growth in the number of customers (a factor 9.5). Combined with the generally increasing trend for $q_{c,T}$ (a factor 1.85) and an increase in $e_{c,T}$ (a factor 1.5), this gives a clearly superlinear growth in $U_c(T)$.

The number of updates $U_p(T)$ received from peers also grows, but at a slower rate (a factor 8.2). This is mainly because of the much slower growth in the number of peers - while the number of customers $m_{c,T}$ increases with a factor of 9.5 over our range of topology sizes, the number of peers $m_{p,T}$ grows only by a factor 1.7. Furthermore, factor $q_{p,T}$ also contributes to the growth in $U_p(T)$ by a factor of 1.6. However, most of the growth can be attributed to the increase in $e_{p,T}$ (a factor of 3.1).

This section has shown how the T nodes experience the highest growth in churn as the network grows with our Baseline growth model. This increase is driven mainly by an increased number of updates from customers. M and CP nodes also see increased churn, driven mainly by their increased MHD. In the next section, we will see how changes in the topology growth model affect the various churn factors.

V. TOPOLOGY GROWTH SCENARIOS

In this section, we look at several single-dimensional deviations from the Baseline model presented above. By looking at how BGP churn increases at various hypothetical growth models, we are able to answer different “what-if” questions about Internet growth. For example, what if multihoming to several providers becomes much more common than today for stub networks? Or what if buying transit services from tier-1 nodes becomes so cheap that they drive regional providers out of business? Our goal is not always to create realistic growth scenarios, but also to highlight the effect of altering different topological properties. Hence, we sometimes look at the effect of large changes to a single property at a time.

As seen in Sec. IV, T nodes experience both the strongest churn in absolute terms, and the strongest increase as the network grows. Hence, we focus mainly on the number of updates received at T nodes.

A. The effect of the AS population mix

First, we look at how the mix of different node types affects churn, by considering four different deviations from the Baseline model with respect to the mix of T, M, CP and C nodes. These deviations illustrate how economic factors can create a very different fauna of networks than what we see today. To implement these scenarios in our model, we change the parameters n_T , n_M , n_{CP} and n_C , while keeping all other parameters fixed.

NO-MIDDLE In the first deviation, we look at a network without M nodes, by setting $n_M = 0$. This illustrates a scenario where the price for transit services from the globally present tier-1 nodes is so low that they have driven regional transit providers out of business.

RICH-MIDDLE In the second deviation, we focus on the opposite scenario, where the ISP market is booming and there is room for a plethora of M nodes. We implement this by




Fig. 8: The effect of the AS population mix on T nodes.

multiplying n_M by 3 ($n_M = 0.45n$), and reducing n_{CP} and n_C accordingly (while keeping their ratio constant).

STATIC-MIDDLE In the third deviation, we look at a situation where all network growth happens at the edges of the network. The number of transit providers (T and M nodes) is kept fixed, and the network grows only by adding CP and C nodes. This could be a plausible scenario for the future, if the ISP population becomes stable.

TRANSIT-CLIQUE In the fourth and final deviation, we let all transit nodes be part of the top-level clique. This scenario may seem far-fetched, but it is important because it shows what would happen if the transit provider hierarchy collapses to a clique of “equals” connected by peering links. We implement this by setting $n_T = 0.15n$ and $n_M = 0$.

Figure 8 shows the average number of updates seen after a C-event at a T node for each deviation as the network grows.

A first observation from the graphs is that the node mix has a substantial influence on churn. In particular, the comparison of RICH-MIDDLE, Baseline, and STATIC-MIDDLE shows that the number of M nodes is crucial. There are two ways in which M nodes increase churn at T nodes. First, an increasing number of M nodes increases the customers $m_{c,T}$ of T nodes. For instance, in the RICH-MIDDLE deviation $m_{c,T}$ increases by a factor of 10.2 when n increases from 1000 to 10000. On the other hand, $m_{c,T}$ increases only by a factor of 5.3 in the STATIC-MIDDLE deviation. Second, an increasing number of M nodes, when they are multihomed to other providers (M or T nodes), tends to also increase the factor $q_{c,T}$. The reason is that M nodes create additional valid paths from the source of a C-event (at stub networks) to T nodes, and so it becomes more likely that a T node will receive updates from its peers and customers after a C-event. Regression analysis shows that the growth of $U(T)$ in the RICH-MIDDLE, Baseline and STATIC-MIDDLE deviations can be modeled as *quadratic*, with different scaling factors.

We also observe that the number of T nodes in the network does not have any impact on the number of updates by itself. The only difference between deviations NO-MIDDLE and TRANSIT-CLIQUE is in the number of T nodes, and we see that the number of updates is the same in these two scenarios. In the absence of M nodes, T nodes will receive one update for each provider the event originator has - either directly from the event originator, or from a peer. This number only increases as a function of the multihoming degree of the event originator, and is not influenced by the network size per se.

An important conclusion from the above observations is that *the increased number of updates does not primarily come from an increased number of transit nodes, but from the hierarchical*

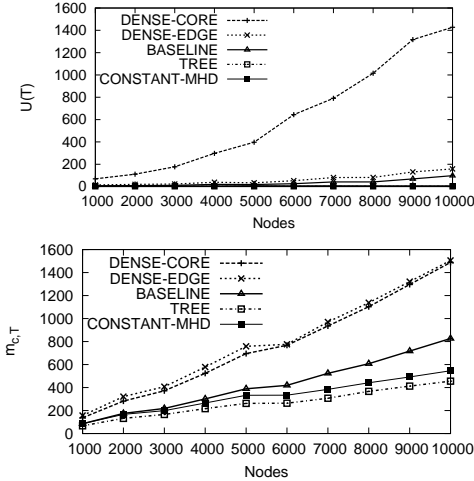


Fig. 9: The effect of the multihoming degree at T nodes.

structure in which they are organized. An Internet with several tiers of providers buying transit services from other providers gives a much higher update rate than a more flat topology where most stub networks connect directly to tier-1 providers. Whether the Internet will move towards a more hierarchical or flat topology in the future is hard to tell. We do know however, that the average path length, measured in AS-level hops, has remained roughly constant, at around 4 hops, during the last ten years [8]. This implies that the Internet retains some hierarchical structure, and that the depth of that structure does not seem to vary with the size of the network.

B. The effect of the multihoming degree

Next, we look at the effect of varying the number of transit links each node brings to the network. Both stub and mid-tier nodes have an incentive to connect to several providers to increase their reliability and load balancing capability. We implement these scenarios by varying the d_M , d_{CP} and d_C parameters, while keeping all other parameters fixed.

DENSE-CORE We look at the effect of much stronger multihoming in the core of the network (M nodes). We implement this deviation by multiplying d_M by 3.

DENSE-EDGE We look at the effect of densification at the edges of the network. In this deviation, stub nodes increase their multihoming degree. We implement this by multiplying d_C and d_{CP} by 3.

TREE We look at a tree-like graph, where all nodes have only a single provider. Here, d_M , d_{CP} and d_C are all set to 1. This is clearly not a realistic scenario, but helps us explore the extreme version of a trend.

CONSTANT-MHD Finally, we look at a scenario where the multihoming degree of all nodes stays constant. We implement this by removing the component of d_M , d_{CP} and d_C that depends on n .

Figure 9 shows the number of received updates (top) and the number of customers m_c (bottom) for T nodes in the different scenarios. First, note that there is a clear connection between the MHD and the number of updates seen at a T node - for the same network size, a higher MHD causes larger churn. Second, even though the number of customers $m_{c,T}$ is about




Fig. 10: The effect of peering relations.

the same in DENSE-CORE and DENSE-EDGE, the churn is significantly higher in the former. This fact illustrates how the meshed connectivity of multihomed M nodes increases the likelihood that a T node will receive updates from a peer or customer. In other words, increased multihoming at the core of the network causes a larger growth in the factor $q_{c,T}$ than increased multihoming at the edges of the network. Specifically, we measured that $q_{c,T}$ increased by a factor of 1.6 in DENSE-CORE, while it increased by a factor of 1.3 in DENSE-EDGE.

When the MHD degree stays constant (in TREE and CONSTANT-MHD), the churn at T nodes is much less. In the extreme case of the TREE model, the churn at T nodes remains constant at two updates per C-event, because the T node learns about the event from exactly one peer or customer (once for the DOWN event and once for the UP event). In the CONSTANT-MHD model, the number of updates is also roughly constant because the increase in the number of customers $m_{c,T}$ as the network grows is offset by a corresponding decrease in the probability $q_{c,T}$ that any given customer of the T node will have the source of that C-event in its customer tree.

According to a recent measurement study [8], the average MHD of both stub nodes and providers has been increasing during the last decade (from 1.4 to 1.8 for stub nodes and from 1.8 to 3.2 for providers). The fact that the MHD has been increasing more rapidly in the core of the network implies that the Internet is closer to the DENSE-CORE model than to the DENSE-EDGE or the CONSTANT-MHD deviations. This can be viewed as bad news, at least in terms of BGP churn.

C. The effect of peering relations

In this subsection, we look at the impact of varying the peering degree between different types of nodes. The fraction of peering links in the Internet has increased over the last decade [8]. However, various difficulties in detecting such links do not allow us to know which peering model is most realistic.

NO-PEERING There are no peering interconnections, except in the clique of T nodes. This is clearly not realistic, but it serves as a reference point.

STRONG-CORE-PEERING We look at densification through more peering links in the core of the network. We model this deviation by doubling p_M .

STRONG-EDGE-PEERING Another possibility is densification through more peering links at the network edges. We model this deviation by multiplying p_{P-M} and p_{P-C} by 3.

Since the peering degree is only changed at M and CP nodes, we show the number of updates received at M nodes




Fig. 11: The effect of provider preference

rather than T nodes. Figure 10 shows the number of updates received at M nodes as a function of network size for the Baseline and each deviation. The main conclusion is that the peering degree does *not* cause a significant change in the generated churn. Adding or removing a significant number of peering links at the edge or at the core of the network does not give major differences in the number of updates. This conclusion also holds for other node types. To explain this observation recall that updates are propagated over peering links only for customer routes. Hence, the fraction of peering links that are active during a C-event is low. Moreover, such updates have limited export-scope (only to customers), compared to routes received from customers.

D. The effect of provider preference

Next, we look at the effect of provider preferences, i.e., the probability that a node chooses to buy transit services from a T or an M node. This choice has implications for how the network will grow; a higher preference for T nodes gives a more “flat” structure, while a higher preference for M nodes diverts more paths through several layers of hierarchy. We define two deviations of the Baseline model:

PREFER-MIDDLE In the first deviation, nodes prefer to buy transit services from M nodes rather than T nodes. We implement this by setting $t_P = t_C = 0$, and limiting the number of T providers for M nodes to one at most.

PREFER-TOP In this deviation, nodes prefer to buy transit services directly from T nodes. We implement this by limiting the number of M providers for M, CP and C nodes to be at most one.




Fig. 12: Updates received at T, M, CP and C nodes.

The top panel in figure 11 shows that a scenario where most nodes buy transit from M nodes results in a higher churn at T nodes, while more direct connections to T nodes decreases churn. *If we are moving towards an Internet in which customers and content providers at the edges prefer to connect to mid-tier ISPs, the number of BGP updates at T nodes will be much higher than if they prefer to connect to tier-1 ISPs.* Looking at the different factors that determine $U(T)$, we observe that the PREFER-TOP deviation gives a much higher $m_{c,T}$ than PREFER-MIDDLE, but that this is more than offset by a strong decrease in $q_{c,T}$, as shown in the middle, and bottom panels in figure 11. An M node is more likely to notify its provider about a C-event than a stub node, because an M node has several potential event sources in its customer tree.

A recent study [8] observed that content providers and regional transit providers tend to buy transit service from either tier-1 or tier-2 providers with almost equal probability. The equivalent of C nodes (“Enterprise Customers”) however, show a preference for tier-2 providers during the last 3-4 years, justifying the selection of the corresponding probabilities in the Baseline model.

VI. EDGE LINK FAILURE

In this section, we focus on a different type of event than the C-event considered so far, namely the failure and subsequent restoration of a link connecting a stub node to one of its providers. We will refer to this event as an “L-event”. We argue that this is a relevant event type to study, even if a single link in our topology model can sometimes represent several physical connections between two ASes. At the edge of the network, these links will often not be replicated, and hence such events are not unlikely to occur in practice.

Unlike a C-event, an L-event does not have to be communicated to all nodes in the network. Stub nodes may have more than one provider, and some nodes will in this case prefer a path that is not affected by the failure. Hence, we expect that an L-event will result in a lower number of updates being propagated in the network compared to a C-event.

We employ the framework introduced in Sec.IV to examine the impact of an L-event in our Baseline growth model. The experiment is repeated for 100 different links (increasing this number does not change the results), and the average number of updates received by each type of node over all experiments is reported. Figure 12 shows the number of updates received at different node types after an L-event. The general trends remain similar to those presented for C-events in Fig. 4.




Fig. 13: q factors.

T nodes experience the highest number of updates, and the strongest increase as the network grows. However, the number of received updates $U(T)$ is significantly lower than after a C-event, and the increase rate is lower, with a growth factor of 4.4 in our range of $n = 1000$ to $n = 10000$ compared to a factor of 14.4 after a C-event.

We break these numbers down according to Eq. 1. Since the m factors depend on the topology and hence are the same for C-events and L-events, the lower number of updates is caused by the reduced growth and absolute numbers for the e and q factors. We find that both the number of updates received from peers $U_p(T)$ and the number received from customers $U_c(T)$ increases with topology size. $U_c(T)$ shows the stronger growth with a factor of 6.6 for our range of topology sizes, compared to a factor 27 for C-events. This is caused by a strong linear growth in $m_{c,T}$ with a factor 9.5, a slow increase in $e_{c,T}$ by a factor 1.3, combined with an overall decrease in $q_{c,T}$. As seen in Fig. 13, $q_{c,T}$ shows significant variations across different topology instances. $U_p(T)$ increases with a factor 2.0 for L-events, compared to a growth of 8.2 for C-events. This can be attributed to a growth factor of 1.6 in $m_{p,T}$, a growth of 1.2 in $q_{p,T}$, and a constant $e_{p,T}$.

The q factors reflect how the number of neighbors that send updates to a T node is reduced for an L-event compared to a C-event. In other words, many nodes prefer other paths to the stub node, and are not affected by the link failure. The e factors reflect how the affected nodes explore much fewer paths before the network converges after an L-event.

Next, we investigate the interaction between a set of topology growth scenario and L-events.

A. Topology growth scenarios

The upper panel in Fig. 14 shows how the number of updates received at a T node increases as the network grows, for different scenarios with respect to the multihoming degree.

First, note that the number of updates received is significantly lower after an L-event than after a C-event in all growth scenarios, as expected. Further, we observe that the DENSE-CORE scenario gives a significantly higher churn growth rate than the other scenarios, while CONSTANT-MHD gives the lowest churn rate. The strong growth in the DENSE-CORE case is caused by the same effect as in the C-event discussed in Sec. V-B; the provider connected to the failed link has 3x more providers, and will hence send an update to many more nodes. Correspondingly, the low number of providers in the CONSTANT-MHD scenario gives a very limited churn.

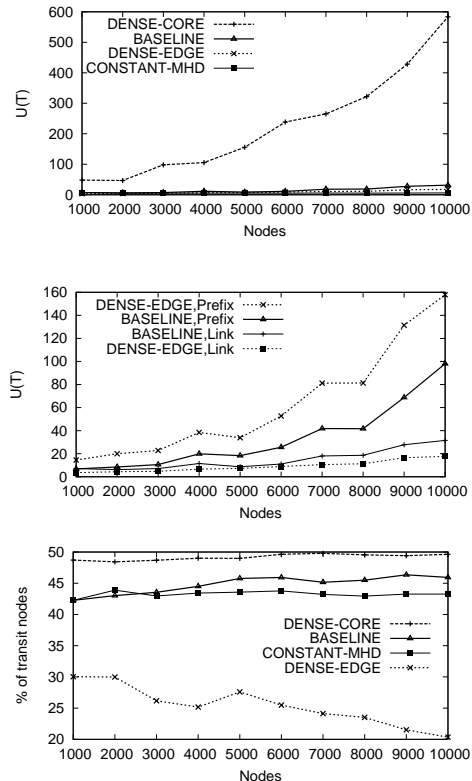


Fig. 14: The effect of multihoming degree on churn after an L-event

The DENSE-EDGE scenario, interestingly, has a lower growth rate than the Baseline scenario. This is different from the situation in a C-event, as shown in the middle panel of Fig. 14. This reduction in churn can be explained by observing that the provider connected to the failed link will send updates to the same number of (transit) neighbors in the DENSE-EDGE and Baseline scenarios, since the multihoming degree in the core is the same in the two scenarios. However, the probability that nodes receiving these updates will change their preferred path is lower in the DENSE-EDGE case, since the stub node can also be reached through several other providers.

The differences between the different growth scenarios is also visible in the bottom panel in Fig. 14, which shows the fraction of transit nodes in the network that select a new preferred path after the L-event. While this fraction is constant or slightly increasing with network size in the CONSTANT-MHD, Baseline and DENSE-CORE scenarios, it decreases from an already lower level in the DENSE-EDGE scenario.

These results show that the effect of densification through increased multihoming degree on the experienced churn level is different depending on the type of event that triggers the re-convergence. *While densification at the edge increases churn after a prefix failure, it gives reduced churn after a link failure.*

We have also investigated churn after an L-event in the other topology growth scenarios discussed in Sec. V. Generally, we find that the number of updates received by a T node is lower for L-events than C-events. However, the general trends observed in churn after an L-event remain similar to what is observed after a C-event.

A recent study [25] showed that the number of routing events that ends with completely withdrawing a prefix, is lower than the number of events that ends with a new preferred path. Further, the extent of path exploration for the later type of events is lower than the former one. These observations fit well with our results.

VII. RELATED WORK

Interdomain routing dynamics and scalability has been a topic in the literature for the last decade or so, after it first was shown that BGP suffers from excessive churn caused by pathological protocol behavior [19]. The phenomenon of *path exploration* was discussed in [17], and upper and lower bounds for the number of updates exchanged during convergence were given. In a follow-up work, it was shown that the duration of path exploration depends on the length of the longest possible backup path to the affected destination [18]. The impact of MRAI timers to limit path exploration was discussed in [13]. In a more recent measurement study, it was shown that path exploration is less severe in the core of the network than on the edges [25]. Another study [5] showed that a small fraction of ASes is responsible for most of the churn seen in the Internet. A recent measurement study concluded that the state of BGP routing is now “healthier” than it was a decade ago, with less update traffic caused by configuration mistakes or protocol pathologies [21]. While all previous papers deal with measuring and explaining BGP routing dynamics, our work differs in that it focuses on the relation of topology growth and BGP dynamics. Zhao et al. [32] investigated the impact of link failures location on BGP dynamics, but only for regular topologies. The work in [31] studied the Internet AS-level topology resilience to certain types of failure. However, it considered mainly reachability issues and had no focus on routing dynamics and topology growth. Increased churn is also a main motivation for completely new routing architectures, like [30]. Other work has focused on BGP scalability in the context of increasing routing table size. One study finds that address fragmentation is the largest contributor to routing table growth [6]. As a response to the increased routing table sizes, a radically different routing strategy called compact routing has been proposed [16]. This approach can give routing table sizes that scale logarithmically with the number of routable addresses, but performs poorly under dynamic conditions.

VIII. CONCLUSIONS

We have examined the role of topology growth on the scalability of BGP. We started by looking at the number of updates received at nodes at different locations in the AS hierarchy after a C-event. For different node types, we have identified the most significant sources of churn, and described how different factors contribute to increased churn as the network grows. We have shown that nodes at the top of the AS hierarchy experience both the highest churn in absolute terms, and the strongest increase as the network grows. We further looked into the impact of L-events on routing scalability, and demonstrated that certain topology growth scenarios scale differently depending on failure types.

Using our flexible topology model, we have explored scalability in several plausible and educational “what-if” scenarios for the growth of the AS-level topology. We have shown that *the most important topological factor deciding the number of updates generated is connectivity in the core of the network*. In particular, the number mid-tier transit providers and the multihoming degree of these nodes plays a crucial role, since *transit nodes in the mid-level of the Internet hierarchy have a special role in multiplying update messages*.

Another important finding from this study is that peering links play a very different role than transit links with respect to scalability. *The peering degree in the Internet does not influence churn*. We have also shown that the depth of the hierarchical structure in the Internet plays a significant role. *If we are moving towards an Internet in which customers and content providers at the edges prefer to connect to mid-tier ISPs, the number of BGP updates at T nodes will be much higher than if they prefer to connect to tier-1 ISPs*. Finally, we have demonstrated that densification through increased multihoming degree will have a different impact on routing scalability depending on its location and the failure type. *While densification at the edge increases churn after a prefix failure, it gives reduced churn after a link failure. On the other hand, a denser core increases churn because after both edge prefix and edge link failures*.

In our future work we plan to look at more complex mix of events, more elaborate topologies that include intradomain protocols and iBGP configurations, and other BGP mechanisms and configurations, such as Route Flap Dampening and BGP multipath extensions.

REFERENCES

- [1] Routeviews project page. <http://www.routeviews.org>.
- [2] SSFNet website. <http://www.ssfnet.org/>.
- [3] R. Albert and A. L. Barabasi. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47, 2002.
- [4] A. L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, October 1999.
- [5] A. Broido, E. Nemeth, and kc. claffy. Internet expansion, refinement, and churn. *European Transactions on Telecommunications*, January 2002.
- [6] T. Bu, L. Gao, and D. Towsley. On characterizing BGP routing table growth. *Computer Networks*, 45(1), May 2004.
- [7] H. Chang, S. Jamin, and W. Willinger. Internet Connectivity at the AS-level: An Optimization-Driven Modeling Approach. In *ACM SIGCOMM Workshop on MoMeTools*, 2003.
- [8] A. Dhamdhere and C. Dovrolis. Ten years in the evolution of the Internet ecosystem. In *IMC 2008*, 2008.
- [9] X. Dimitropoulos, D. Krioukov, A. Vahdat, and G. Riley. Graph annotations in modeling complex network topologies. *arXiv:0708.3879*, 2008.
- [10] X. Dimitropoulos and G. Riley. Efficient large-scale BGP simulations. *Elsevier Computer Networks, Special Issue on Network Modeling and Simulation*, 50(12):2013–2027, 2006.
- [11] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *ACM SIGCOMM*, pages 251–262, 1999.
- [12] Lixin Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking (TON)*, 9(6), December 2001.
- [13] T. Griffin and B. Premore. An experimental analysis of BGP convergence time. In *ICNP*, 2001.
- [14] H. Haddadi, D. Fay, A. Jamakovic, O. Maennel, A. W. Moore, R. Mortier, and S. Uhlig. On the importance of local connectivity for Internet topology models. In *the 21st International Teletraffic Congress*, pages 1–8, September 2009.
- [15] Y. He, S. V. Krishnamurthy, M. Faloutsos, and M. Chrobak. Policy-aware topologies for efficient inter-domain routing evaluations. In *IEEE INFOCOM 2008 Mini-Conference*, Phoenix, AZ, USA, April 2008.

- [16] D. Krioukov, kc. claffy, K. Fall, and A. Brady. On compact routing for the Internet. *Computer Communications Review*, 37(3), July 2007.
- [17] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed Internet routing convergence. In *ACM SIGCOMM*, pages 175–187, August 2000.
- [18] C. Labovitz, A. Ahuja, R. Wattenhofer, and S. Venkatachary. The impact of Internet policy and topology on delayed routing convergence. In *INFOCOM*, Anchorage, AK, USA, April 2001.
- [19] C. Labovitz, G. R. Malan, and F. Jahanian. Internet routing instability. *SIGCOMM Comput. Commun. Rev.*, 27(4):115–126, 1997.
- [20] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graph Evolution: Densification and Shrinking Diameters. *ACM Transactions on Knowledge Discovery from Data (ACM TKDD)*, 2007.
- [21] J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkrantz. BGP routing dynamics revisited. *Computer Communications Review*, April 2007.
- [22] L. Li, D. Alderson, R. Tanaka, J. C. Doyle, and W. Willinger. Towards a theory of scale-free graphs: Definition, properties, and implications (extended version). *Internet Mathematics*, 2005.
- [23] P. Mahadevan, D. Krioukov, M. Fomenkov, X. Dimitropoulos, kc. claffy, and A. Vahdat. The Internet AS-level topology: three data sources and one definitive metric. *SIGCOMM Comput. Commun. Rev.*, 36(1):17–26, 2006.
- [24] D. Meyer, L. Zhang, and K. Fall. Report from the IAB workshop on routing and addressing. <http://tools.ietf.org/id/draft-iab-raws-report-02.txt>, April 2007.
- [25] R. Oliveira, B. Zhang, D. Pei, R. Izhak-Ratzin, and L. Zhang. Quantifying path exploration in the Internet. In *IMC*, Rio de Janeiro, Brazil, October 2006.
- [26] B. Quoitin and S. Uhlig. Modeling the routing of an autonomous system with C-BGP. *IEEE Network*, 19(6), November 2005.
- [27] Y. Rekhter, T. Li, and S. Hares. A border gateway protocol 4 (BGP-4). RFC4271, January 2006.
- [28] RIPE's Routing Information Service. <http://www.ripe.net/ris/>.
- [29] G. Siganos, M. Faloutsos, and C. Faloutsos. The Evolution of the Internet: Topology and Routing. *University of California, Riverside technical report*, 2002.
- [30] L. Subramanian, M. Caesar, C. Tien Ee, M. Handley, Z. Morley Mao, S. Shenker, and I. Stoica. HLP: a next generation inter-domain routing protocol. In *Proceedings SIGCOMM*, pages 13–24, Philadelphia, USA, August 2005. ACM.
- [31] J. Wu, Y. Zhang, Z. Mao, and K. G. Shin. Internet routing resilience to failures: analysis and implications. In *CoNEXT '07*, pages 1–12, New York, NY, USA, 2007. ACM.
- [32] X. Zhao, B. Zhang, A. Terzis, D. Massey, and L. Zhang. The impact of link failure location on routing dynamics: A formal analysis. In *ACM SIGCOMM Asia Workshop*, April 2005.
- [33] S. Zhou and R. Mondragon. Accurately Modeling the Internet Topology. *Physical Review E*, vol. 70, 2004.



Dr. Constantine Dovrolis is an Associate Professor at the College of Computing of the Georgia Institute of Technology. He received the Computer Engineering degree from the Technical University of Crete in 1995, the M.S. degree from the University of Rochester in 1996, and the Ph.D. degree from the University of Wisconsin-Madison in 2000. He joined Georgia Tech in August 2002, after serving at the faculty of the University of Delaware for about two years. He has held visiting positions at Thomson Research in Paris, Simula Research in Oslo, and FORTH in Crete. His current research focuses on the evolution of the Internet, Internet economics, and on applications of network measurement. He is also interested in network science and in applications of that emerging discipline in the understanding of complex systems. Dr. Dovrolis has been an editor for the IEEE/ACM Transactions on Networking, the ACM Communications Review (CCR), and he served as the Program co-Chair for PAM'05, IMC'07, and as the General Chair for HotNets'07. He received the National Science Foundation CAREER Award in 2003.



Ahmed Elmokashfi is a Ph.D student at Simula Research Laboratory and University of Oslo. He received his B.Sc. degree in Electrical and Electronics Engineering from the University of Khartoum in 2003, and his M.Sc degree in Telecommunication from Blekinge Institute of technology in 2007. Elmokashfi research interests include IP routing, routing scalability, measurements, and network science.



Amund Kvalbein received his M.Sc. degree from the University of Oslo in 2003, and his Ph.D. degree from the same university in 2007. The main focus of his Ph.D. thesis is fast recovery mechanisms. He is now a Post Doc at Simula Research Laboratory, where he is leading a project on network resilience. His main research interests are network layer protocols, in particular issues related to fault tolerance, scalability and robustness under unexpected operational environments.