

Observing the Evolution of Internet AS Topology *

Ricardo Oliveira
University of California
Los Angeles, CA
rveloso@cs.ucla.edu

Beichuan Zhang
University of Arizona
Tucson, AZ
bzhang@cs.arizona.edu

Lixia Zhang
University of California
Los Angeles, CA
lixia@cs.ucla.edu

ABSTRACT

Characterizing the evolution of Internet topology is important to our understanding of the Internet architecture and its interplay with technical, economic and social forces. A major challenge in obtaining empirical data on topology evolution is to identify *real* topology changes from the *observed* topology changes, since the latter can be due to either topology changes or transient routing dynamics. In this paper, we formulate the topology *liveness problem* and propose a solution based on the analysis of BGP data. We find that the impact of transient routing dynamics on topology observation decreases exponentially over time, and that the real topology dynamics consist of a constant-rate birth process and a constant-rate death process. Our model enables us to infer real topology changes from observation data with a given confidence level. We demonstrate the usefulness of the model by applying it to three applications: providing more accurate views of the topology, evaluating theoretical evolution models, and empirically characterizing the trends of topology evolution. We find that customer networks and provider networks have distinct evolution trends, which can provide an important input to the design of future Internet routing architecture.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Network topology

General Terms

Measurement, Theory

Keywords

Internet topology, topology evolution

*This work is partially supported by the National Science Foundation under Contract No. CNS-0520318. Opinions and findings expressed in this paper are those of the authors and do not necessarily reflect the views of NSF.

Ricardo Oliveira was supported by a scholarship from FCT/Portugal, ref. SFRH/BD/17268/2004.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'07, August 27–31, 2007, Kyoto, Japan.
Copyright 2007 ACM 978-1-59593-713-1/07/0008 ...\$5.00.

1. INTRODUCTION

The Internet has been evolving rapidly over time like a living organism, so has its topology. Characterizing the dynamics and evolution trends of the Internet topology is an important research topic for several reasons. It provides an essential input to the understanding of limitations of existing routing protocols, the evaluations of new designs, as well as the projection of future needs; and it will help advance our understanding of the interplay between networking technology, the resulting topology, and the economic forces behind them.

In recent years there have been a plethora of efforts in understanding the Internet AS topology and its evolution. On the empirical side, most work has focused on examining graph properties of topology snapshots (*e.g.*, node degree distribution [17], likelihood [22], and degree correlation [24]). On the theoretical side, a number of topology evolution models (*e.g.*, [7, 12, 13, 31]) have been proposed for the addition and removal processes of nodes and links in the Internet AS topology. However there is a missing link between these two categories of work: an empirical description of the topology evolution, that is when and where AS nodes and inter-AS links are added or removed over time *in reality*. Without this knowledge, it is difficult to judge the accuracy of topology snapshots taken at different times in the graph property computation, or to validate theoretical evolution models. For example, up to now, theoretical evolution models have been evaluated by comparing certain graph properties of the topology derived from the evolution model with that of the AS topology snapshots taken over time. As shown in [22], however, significantly different topologies can share the same graph properties, such as power-law distribution of node degrees. A more accurate evaluation approach would be to compare the modeling results against empirical topology evolution data directly.

Obtaining empirical topology evolution data is a challenging task. Besides the well known issue that the observed AS topology snapshots only capture a subset of the real Internet topology [10, 33, 25, 28, 32], a new problem arises when we try to measure topology *changes* over time: *the changes in the observed topology* do not necessarily reflect *the changes in the real topology* and vice versa. Because the observed topology is normally inferred from routing or data paths, its changes can be due to either real topology changes or transient routing dynamics (*e.g.*, caused by link failures or router crashes). Therefore the challenge is, given all the changes in the observed topology over time, how to differentiate those caused by real topology changes from those caused by transient routing dynamics, which we call the *liveness problem*. Only after solving this problem can we provide empirical topology evolution data such as when and where an AS or an inter-AS link is added or removed from the Internet.

In this paper we developed a solution to the liveness problem based on the analysis of available data. Our analysis shows that the effect of transient routing dynamics on the observed topology decreases exponentially over time, and the real topology changes can be modeled as the combination of a constant-rate birth process and a constant-rate death process. Fitting the model to data, we obtain the values of several key parameters that enable us to identify real topology changes from observed changes with a given confidence level, and to estimate the birth and death rates of links and nodes in the Internet AS topology for the first time.

We show the usefulness of our empirical topology evolution model through three applications. First, it allows us to obtain a more accurate topology by combining snapshots taken at different times and discounting dead links and nodes. Second, it can be used to evaluate theoretical topology evolution models. Third, it reveals evolution trends of the Internet. More specifically, we find that the edge of the Internet and the core have distinct growth patterns. The Internet’s growth in size is mainly attributed to customer networks at the edge of the AS topology, where the majority of node births and deaths occur. Provider networks, which occupy the core in the topology, grow slowly in number but adjust their inter-connectivity frequently, and most adjustments involve peer-to-peer AS links.

We would like to be the first to point out an intrinsic limitation of our results. The AS topology we collected and studied is inevitably incomplete, and the numerical results based on this partial topology could be potentially off from that of the real Internet topology. However this inherent limitation is shared by all empirical studies on the Internet topology and it should not stop us from solving the liveness problem. We also make our best efforts to mitigate the problem. We use the data collected from hundreds of publicly available BGP monitors, and we observed diminishing returns by additional monitors once the number reaches a certain level. We develop our model using different numbers and combinations of the monitors to assure that the model has no bias towards any particular set of monitors. Furthermore, we compare the results based on BGP data with that based on data from traceroute and Internet registries. Section 3.5 provides further discussions on the model’s limitations.

The rest of the paper is organized as follows. Section 2 defines the liveness problem in studying Internet topology. Section 3 develops and verifies our topology evolution model using three different data sources: BGP logs, traceroute, and Internet registries. It also provides an evaluation of the AS topology data collected from these sources regarding their suitability for solving the liveness problem. Section 4 presents three applications of our model. Section 5 discusses the related work, and Section 6 concludes the paper.

2. TOPOLOGY LIVENESS AND COMPLETENESS PROBLEMS

One can abstract the Internet topology at the inter-domain level as an AS connectivity graph. In this graph, a *node* represents an Autonomous System (AS), and a *link* the existence of one or multiple BGP routing sessions between two neighbor ASes. An inter-AS link represents a contractual relation between two ASes to forward data for each other over the link. Because individual ASes apply private routing policies to BGP updates, generally speaking one cannot observe the complete AS topology. We denote the complete real Internet AS topology graph by G_{real} , and the topology graph that one infers from measurement data by G_{obsv} . The observed portion of the AS topology is a subset of the real topology, *i.e.*, $G_{obsv} \subset G_{real}$. Knowing how much these two topologies differ is what we term the *Completeness Problem*.




Figure 1: Observing Topology Over Time

G_{obsv} can be constructed in multiple ways. One way is to have data collectors establish BGP sessions with a set of operational routers, which we call *monitors*, to obtain their BGP routing tables and updates. Another way is to have a set of vantage points send traceroute probes and then to convert the obtained router paths to AS paths¹. For example, in Fig. 1, at time t_0 , we measure the topology from monitor A by either examining A ’s routing table or probing other two nodes B and C . The resulting G_{obsv} misses one link, $B-C$, from G_{real} . To study graph properties of the AS topology it is important to minimize the number of missing links. Existing efforts in this area include deploying additional monitors and incorporating data from other sources (*e.g.*, routing registry [32]). For example, if B is also a monitor, then one can observe the existence of link $B-C$.

As a direct consequence of our inability to observe the complete topology, another problem, which we call the *Liveness Problem*, arises when we study topology evolution over time. That is, an observed change in G_{obsv} does not necessarily reflect a change in G_{real} . For example, in Fig. 1, at time t_1 , link $A-C$ goes down due to a physical failure, but this failure does not change the contractual relationship between A and C , *i.e.*, link $A-C$ still exists in G_{real} . However, the routing protocol will adapt to the failure and link $A-C$ disappears from the observation. As a result, comparing $G_{obsv}(t_0)$ with $G_{obsv}(t_1)$, we will see one link removal ($A-C$) and one link addition ($B-C$). In another example, consider the changes from time t_2 to time t_3 in Fig. 1. D changes its service provider by switching from C to B . This is a real topology change and results in one link removal ($D-C$) and one link addition ($D-B$) in both G_{real} and G_{obsv} . In both cases, what we observe are changes in G_{obsv} , and the question is how to tell which ones are real topology changes happened in G_{real} .

We use *appearance* and *disappearance* to name the addition and removal of elements (*i.e.*, links and nodes) in G_{obsv} respectively, and *birth* and *death* to name the addition and removal of elements in G_{real} respectively. The liveness problem concerns how to infer the real births and deaths from observed appearances and disappearances. More specifically, when a link or node disappears from G_{obsv} , is it still alive in G_{real} ? When a link or node appears for the very first time, has it been alive in G_{real} before? Answering these questions is critical to studying topology evolution, as we need to know when and where births and deaths occur in G_{real} .

¹Different from BGP monitors, traceroute vantage points are usually end hosts. However in this paper we term both as monitors.

The liveness problem and completeness problem are related in that solving one will help solve the other. If the liveness of links and nodes is known, we can combine observations made at different times to form a more complete topology estimate. For example, in Fig. 1, combining $G_{obsv}(t_0)$ and $G_{obsv}(t_1)$ will give a more complete topology at time t_1 , provided that we know link $A-C$ is still alive at time t_1 . Similarly, if the complete topology is known, we will be able to differentiate real topology changes from transient routing changes. For example, if we know the complete topology in Fig. 1, we will not take the appearance of link $B-C$ at time t_1 as a birth.

However the liveness problem and completeness problem are also fundamentally different. On the one hand, even if we know the liveness of all the observed links and nodes over time and are able to combine observations made through a long time period, we still do not know whether the combined topology is complete, or how incomplete it may be. For example, in Fig. 1, from time t_2 to time t_3 , knowing the liveness of links and nodes does not help tell whether link $B-C$ exists. On the other hand, even if monitors are placed at every node to capture all the links (except those having failures at the moment), when link $A-C$ disappears from the observation at time t_1 , we still cannot tell instantly whether it is due to an operational failure or the termination of the inter-AS contract, although observations over time can provide a good estimate as described later.

Both the liveness problem and completeness problem are important to a full understanding of the Internet topology and its evolution. An ideal solution would be having all the ISPs register their inter-AS connectivity at a central registry and keep their entries up-to-date, which, unfortunately, does not seem feasible in the current Internet. A near ideal solution would be placing a monitor in each AS, which is also infeasible in reality. A number of research efforts have been devoted to making G_{obsv} more complete, without knowing exactly *how* close the obtained G_{obsv} is to G_{real} . However, to our knowledge, no one has addressed the liveness problem, which has been a major hurdle to empirical studies of topology evolution. In this paper, we focus on the liveness problem and propose a solution based on the analysis of available topology data.

Intuitively, real topology changes generally occur over relatively long time intervals (*e.g.*, months or even years), while transient routing changes happen within much shorter periods (*e.g.*, minutes or hours). Thus if we keep observing the topology over time, we should be able to differentiate topology changes from transient routing changes. For example, if a link disappears and re-appears after a short period of time, it is most likely that the disappearance is not a death. If a link disappears and never re-appears again over a long time period, it is most likely that the link no longer exists. The research question is how long one should wait before declaring a birth or death with a given level of confidence. We develop an empirical model that captures the effects of long-term topology changes and short-term routing changes on observed topologies.

Internet topology can be abstracted at different granularity, *e.g.*, router-level topology, AS-level topology, and ISP-level topology (a number of ISPs have multiple ASes). Although this paper focuses on the AS-level topology, the liveness problem is a general problem that exists independently from whether the nodes in Fig. 1 are routers, ASes, or ISPs. Thus we believe that solving the problem at the AS-level could lead a way to liveness solutions at other granularity. For example, if we can identify real topology changes for each AS, by combining the behavior of ASes that belong to the same ISP, we will get the topology changes for ISP-level topology. One of our future work is to apply the methodology developed in this paper to other types of topologies.

3. AN EMPIRICAL MODEL OF OBSERVED TOPOLOGY DYNAMICS

In this section, we develop an empirical model to capture the observed topology dynamics over time, and based on the model we propose a solution to the liveness problem. We develop the model using BGP log data, verify its consistency with information extracted from Internet registries, and evaluate the suitability of existing traceroute data sets in solving the liveness problem.

3.1 Data Sets

We use data from three different types of sources: BGP, traceroute, and Internet registries. The BGP data consists of both routing tables and updates collected by RouteViews [4] and RIPE [3] from a few hundreds of monitors between January 1, 2004 and December 1, 2006, a period of almost three years². From BGP routing tables and updates, we extract topology information (*i.e.*, AS nodes and links) and record the timestamps of appearances and disappearances of links and nodes. There are totally 27,972 nodes and 123,182 links in the entire data set. To evaluate the effects of different monitors, we group BGP data into three sets.

- *ATT*: data from a single monitor residing in AS7018.
- *Set-54*: data from a set of 54 monitors residing in 35 ASes; these monitors are present throughout the entire measurement period.
- *ALL*: data from all monitors.

The traceroute data is collected and kindly provided to us by three research projects: Skitter [5], DIMES [29], and iPlane [23]. They all have monitors around the globe to periodically traceroute thousands of destination IP addresses, and convert router paths to AS paths. They differ in the number of monitors, locations of monitors, probing frequency, and the list of destinations to probe. Both Skitter and DIMES have data from January 1, 2004 to December 1, 2006, but iPlane's data collection only started from late June, 2006. Each data set comes with an AS adjacency list describing the AS topology it observes.

We also extract AS number allocation data from Regional Internet Registries (RIR) [2], and AS connectivity information from Internet Routing Registries (IRR) [1].

3.2 An Empirical Model

We first use BGP data to develop an empirical model for observed topology changes. Before starting the model development, we would like to note an important difference between links and nodes in terms of their observability. Due to the relatively small number of existing monitors and the rich connectivity among ASes, many links are not seen on the first day of observation; some of them get revealed through routing dynamics over time. However, because most ASes (over 99%) originate one or more prefixes, they appear in the global routing table on the first day of observation; the small number of remaining transit ASes behave in the same way as links in terms of their observability. As a result, the same model applies to both links and nodes. We will focus on developing the model for links, and only show the results of applying the model to nodes.

3.2.1 The Appearance of Links and Nodes

Observations: Fig. 2 shows the *cumulative* number of unique links captured by different monitor sets over time. Taking the *ATT* curve

²The main reason for starting from 2004 instead of earlier is to have an adequate number of monitors for the entire measurement period.




Figure 2: Number of links captured by different sets of monitors




Figure 3: Number of monitors in RouteViews and RIPE combined




Figure 4: Number of links, AT&T monitor with different starting times




Figure 5: Visible links seen by all monitors




Figure 6: Link disappearance period




Figure 7: Link disappearance period, by all monitors

for instance: on the first day, the observed links are those in the monitor's routing table on January 1, 2004; a point (200, 40000) on the curve means that during the first 200 days, this monitor has seen 40000 unique links in total from its BGP routing tables and updates.

As shown in Fig. 2, all the three curves share a common pattern: they start with a relatively high growing rate, but slows down over time and settle on a more or less constant growth rate. For the *ALL* curve, despite that the number of monitors has been changing over time (see Fig. 3), its overall shape is the same as the other two's, except slight bumpiness at the beginning. The same pattern also holds across different starting times of the observation, as shown in Fig. 4. Therefore, this pattern hints at something fundamental to topology observation.

Intuitively, we can interpret the linear portion of the curve as due to real topology changes (*i.e.*, link births) and the initial fast growth as caused by originally hidden links being revealed by transient routing dynamics. The curves show that, within the first 100 to 200 days, most links that could be revealed have shown up. After that point the effect of the *revelation process* becomes minimal, and the curves would have flattened out eventually had there been no link birth. The sustained linear increase of the curves gives a strong indication of topology changes by link births. We derive an empirical model to quantify this intuition as follows.

Modeling: According to their observability, we sort all links into three types: *Visible* (links that have been observed), *Invisible* (links that cannot be observed by the given set of monitors³) and *Hidden*

³Invisible links exist because of routing policies, *e.g.*, a peer-to-peer link between two ASes will not be advertised to their providers, thus it is impossible to be observed by the providers.

(links that are possible to be observed but have not yet). Fig. 2 and Fig. 4 show the cumulative number of unique visible links over time. We make the following two simple assumptions:

- *Constant Birth Rate:* Let b_v be the birth rate of visible links, b_h the birth rate of hidden links, then the total birth rate of visible and hidden links $b = b_v + b_h$.
- *Uniform Revelation Probability:* The probability for each hidden link to be revealed during a small time interval Δt is $\lambda \Delta t$.

At a given time t , let $v(t)$ be the cumulative number of visible links observed from time 0 to time t , and $h(t)$ be the number of hidden links at time t . Consider a small time interval from t to $t + \Delta t$. During this period, $\lambda \cdot h(t) \Delta t$ hidden links are revealed; at the same time, $b_h \cdot \Delta t$ new hidden links are born. Therefore,

$$\Delta h = h(t + \Delta t) - h(t) = -\lambda h(t) \Delta t + b_h \Delta t = (b_h - \lambda h(t)) \Delta t$$

$$\frac{\Delta h}{b_h - \lambda h(t)} = \Delta t$$

Integrating both sides from time 0 to time t , we have:

$$h(t) = h_0 e^{-\lambda t} + \frac{b_h}{\lambda} (1 - e^{-\lambda t})$$

where h_0 is the number of hidden links at time 0. Since $h(t \rightarrow \infty) = \frac{b_h}{\lambda}$, we can re-write the above equation as

$$h(t) = h_0 e^{-\lambda t} + h_\infty (1 - e^{-\lambda t})$$

Now consider the number of observed links $v(t)$, between time t and $t + \Delta t$,

$$\Delta v = \lambda h(t) \Delta t + b_v \Delta t = \lambda (h_0 - h_\infty) e^{-\lambda t} \Delta t + b \Delta t$$

Integrating both sides from time 0 to time t , we get

$$v(t) = v_0 + bt + (h_0 - h_\infty)(1 - e^{-\lambda t}) \quad (1)$$

where v_0 is the number of links observed on the first day, bt reflects the linear birth process, h_0 is the initial number of hidden links, h_∞ is the number of hidden links as observation time $t \rightarrow \infty$, and the impact of revelation process decreases exponentially over time.

Results: We perform non-linear regressions on the data based on Eq. 1, and the fit is very good for all three sets of monitors, including the *ALL* curve (Fig. 5), which has a changing set of monitors. All regression results presented in this paper (e.g., Fig. 5, 7, 13 and 14) have high coefficient of determination, $R^2 > 99.5\%$. The good fitting means that the simple model approximates real data satisfactorily. As explained earlier, the same model should apply to nodes as well, since hidden nodes are revealed in the same way as hidden links by routing dynamics. The fit to node data is also very good. Values of model parameters are obtained from regressions and listed in Fig. 9.

3.2.2 The Disappearance of Links and Nodes

Observations: A link disappeared from $G_{obsv}(t)$ can be a real death, or it can be still alive in G_{real} , not observed by any monitor at the moment but may re-appear sometime in the future. Assuming the observation period ends on day n , we define that a link has a *disappearance period* of $(n - m)$ days if the link disappeared on day m and has not re-appeared until the end of the observation. Note that even though a link may appear and disappear many times in the entire observation period, only the *last* disappearance counts in calculating the link's disappearance period.

Fig. 6 shows the cumulative number of links over the disappearance period. For instance, a data point at (200, 21000) on the *ATT* curve means that, at the end of the observation, 21000 links have a disappearance period less than or equal to 200 days as seen by the *ATT* monitor. Interestingly, the curve also exhibits the pattern of an initial exponential component plus a stable linear component over time, and this same pattern holds across different monitor sets and different observation ending times.

Modeling: We divide visible links into two subtypes: *in-sight* (links that are in the currently observed topology $G_{obsv}(t)$), and *out-of-sight* (links that have been seen previously, but are not in $G_{obsv}(t)$, and may come back to G_{obsv} sometime later). We make two simple assumptions:

- **Constant Death Rate:** In-sight links disappear from the monitors' view at a rate of $d + f_0$, where d is the number of link deaths, f_0 the number of links that become out-of-sight, per unit time.
- **Uniform Revelation Probability:** For each out-of-sight link, the probability of being revealed (i.e., become in-sight again) during a small time interval Δt is $\mu\Delta t$. This revelation process is essentially the same as the one described earlier for appearance. We use different notations, λ and μ , since the former is computed from the *first appearance* of links, and the latter is computed from the *re-appearance* of links.

Suppose the observation ends at time t_{end} . Consider the f_0 links that become out-of-sight at time t_d , and $s = t_{end} - t_d$. Let $f(x)$ be the number of these links that have not re-appeared since time t_d through time $t_d + x$. After a short time period Δx , some of these links may be revealed by routing dynamics. Therefore,

$$\Delta f(x) = f(x + \Delta x) - f(x) = -\mu f(x)\Delta x$$

Integrating from $x = 0$ to $x = s$,

$$f(s) = f_0 e^{-\mu s}$$

At the end of observation, the number of links with disappearance period of s is equal to $d + f(s)$. Let $z(s)$ be the cumulative number of links whose disappearance period is less than or equal to s , then

$$z(s) = \int_{y=0}^{y=s} (d + f(y)) dy = \frac{f_0}{\mu} (1 - e^{-\mu s}) + sd \quad (2)$$

where the death process is captured by a linear term, and the revelation process (of disappeared links) is captured by an exponential term.

Results: Eq. 2 fits data well and the results are shown in Fig. 7. The same model can also be applied to nodes. Model parameters, for both appearance and disappearance of links and nodes, are listed in Fig. 9. Even though λ is estimated from first appearance and μ is estimated from re-appearance, they have similar numerical values, which is consistent with our model that both parameters characterize the same revelation process. Note that in deriving the model for link appearance, we did not take into account the death process of visible or hidden links for clarity. The death of visible links does not affect Eq. 1 because Eq. 1 is about *cumulative* number of observed links. Assuming the death rate for hidden links is d_h , the only change it makes in Eq. 1 is to replace b_h by $(b_h - d_h)$, e.g., $b = b_v + b_h - d_h$ instead of $b_v + b_h$.

3.2.3 Distinguishing Topology Changes from Transient Routing Changes

Based on our empirical model, the effects of transient routing dynamics on observed topology decrease exponentially over time, while the real birth and death occur at constant rates. If one observes the topology long enough, the observed changes will be dominated by real topology changes. Assume the observation starts at time t_{start} , ideally one can wait for a sufficiently long time B so that every *new* link appearance at time $t > t_{start} + B$ can be considered a birth with high confidence. Similarly, after a link disappears, one can wait for a sufficiently long time D and if the link does not re-appear during this time, it can be considered a death with high confidence. Now we are ready to quantify B and D with certain confidence.

According to our model, on each day the newly discovered links come from two sources: b_v from birth and $\lambda h(t)$ from revelation. If we count all the newly discovered links on day t as birth, the chance of being correct is

$$confidence(t) = \frac{b_v}{b_v + \lambda h(t)} = \frac{b_v}{b + \lambda(h_0 - h_\infty)e^{-\lambda t}} \quad (3)$$

From regression on data, we can obtain the values of b , λ , and $(h_0 - h_\infty)$. To estimate b_v , we assume $b_v \simeq b$. This is based on the following observation. Since $b = b_v + b_h - d_h$, our assumption is equivalent to $b_h \simeq d_h$. If b_h and d_h differ significantly, the number of hidden links at the beginning of observation would vary significantly with different starting dates, i.e., h_0 as well as $(h_0 - h_\infty)$ would change significantly over different starting dates. However, our examination of data shows that this is not the case, i.e., $(h_0 - h_\infty)$ remains relatively stable over different starting dates, which validates the assumption of $b_h \simeq d_h$ and $b_v \simeq b$.

Knowing the values of the parameters, we can now calculate B for a given confidence level using Eq. 3. Similarly, D can be calculated for a given confidence level based on our model of disappearance. Fig. 8 shows the values of B and D for different confidence values. For instance, after 205 days from January 1, 2004, we can count all newly discovered links as real births with 90% or higher




Figure 8: Observation period as a function of confidence level for links

Parameters	Links	Nodes
Birth rate b (day^{-1})	67.3	10.3
Revelation λ (day^{-1})	0.0151	0.0223
$(h_0 - h_\infty)$	11013	240
Death rate d (day^{-1})	45.7	2.87
Revelation μ (day^{-1})	0.0196	0.0172
f_0/μ	10545	797

Figure 9: Model Parameters




Figure 10: Node birth from RIR




Figure 11: Link birth from IRR




Figure 12: Link death from IRR




Figure 13: Visible links in Skitter, $\lambda = 0.00598$, $b = 39.86$.

probability of being correct. If a link disappears and does not show up after 189 days, with 90% chance this is a real death.

So far we have considered links and nodes separately. For a given observation period (B or D), the confidence level of nodes usually is higher than that of links. Thus considering both together sometimes can improve the confidence on links. For instance, a node birth is always accompanied by some link births. Therefore we can decide these link births with higher confidence than what we would have by only considering links.

3.3 Comparison with Internet Registry Data

Our model is built upon the assumption that a linear term captures real topology changes in G_{real} . Here we use Internet registry data to check the soundness of this assumption. The registry data is particularly useful because it is not affected by routing dynamics.

Regional Internet Registries (RIRs) [2] maintain complete history of AS number allocations. Fig. 10 shows the cumulative number of allocations since January 1, 2004. It can be approximated by a straight line with slope of 10.4 nodes/day, very close to 10.3 nodes/day obtained from our model⁴. Since there is a variable delay between an AS number's allocation and its announcement in the global routing system, we cannot use the allocation date to verify the birth date of nodes. However, the fact that the AS allocation exhibits a growth rate extremely close to our model provides a supporting evidence for our assumption of linear node birth. We are not able to check node death rate with RIR data since deallocation of AS numbers is not mandated and usually is not done in practice.

⁴Some AS numbers are allocated but never used for global routing. The cumulative number of such ASes grows over time in our data, which may explain the slight difference between node birth rates obtained from RIR and BGP data.

Internet Routing Registries (IRR) [1] are databases for registering inter-AS connections and routing policies. Registration with IRR is done voluntarily. It is known that information in IRR is incomplete, and out-of-date for some entries. Historic IRR files are not publicly available, but we have been downloading a daily copy since April 5, 2005. Fig. 11 shows link appearances and Fig. 12 shows link disappearances based on IRR data. Again, one cannot use the IRR data to verify the birth/death dates of individual links, since registering a link and bringing a link up usually happen on different days. Both figures exhibit linear behavior over time, which is consistent with our assumption of linear link birth and death. The rates obtained from IRR data are lower than that from BGP data, which is likely due to the incompleteness of IRR data. More specifically, the birth rate estimated in Fig. 11 is about 86% of that from BGP data, whereas the death rate estimated in Figure 12 is only 40% of that from BGP data. This indicates that even some operators are willing to register their connectivity, they still tend to overlook the removal of stale information.

3.4 Evaluation of Traceroute Data

Besides BGP routing tables and updates, traceroute data is another major source for AS topology information. In this subsection, we analyze existing traceroute data with regard to their effectiveness in solving the liveness problem.

3.4.1 Measuring AS Topology by Traceroute

BGP and traceroute measurements work differently. A BGP data collector passively listens to routing updates regarding all the globally announced IP prefixes from individual monitors and logs the updates as well as routing tables for each monitor. A traceroute monitor actively sends UDP or ICMP probes to a list of IP ad-




Figure 14: Link disappearance period, by Skitter, $\mu = 0.0385$, $d = 57.61$.




Figure 15: Comparison of appearance timestamps between Skitter and BGP.




Figure 16: Comparison of disappearance timestamps between Skitter and BGP.

dresses and records the router-level paths, which is then converted to AS-level paths.

BGP and traceroute data can be complementary in topology measurement. Since usually their monitors are placed in very different locations, they may be able to see different parts of the Internet topology. Also, BGP data records the routing paths, while traceroute records the data paths, which can be different in some cases. For example, if a provider AS P aggregates a customer AS C 's prefix in P 's routing announcement, a BGP monitor may not be able to see the link $P-C$, but traceroute can reveal its existence. However, as pointed out in [26, 21, 11], accurately converting router paths to AS paths remains an open issue, and there can be many pitfalls in this process. One commonly used method of converting router IP addresses to AS numbers is to look up BGP routing table and RIR address allocation database, which, as shown in previous work, may introduce false AS links. For example, assuming three ASes are connected as $A-B-C$, if B 's border router uses one of A 's IP addresses, then the simple conversion will give a false AS link $A-C$.

As a sanity check on traceroute data, we compared the links observed in BGP data with that reported by Skitter during January 2007, and manually verified the difference between the two data sets by contacting the operators of 20 ASes. For these 20 ASes, Skitter reported 447 links that were not in BGP data. However, only 16 out of these 447 links (4%) were confirmed by the operators. Unfortunately, all three traceroute data sets used BGP tables and WHOIS lookups for the IP address to AS number conversion, thus they may potentially suffer the same conversion errors⁵. Due to the potentially false links in the data sets, we did not include traceroute data in our model development. However, since traceroute data potentially can be a very valuable source for AS topology information, we evaluate the three data sets and discuss the impacts of two important measurement factors: probing frequency and destination list.

3.4.2 Skitter

During our measurement period, Skitter has about 20 monitors around the globe. Every day each monitor probes every IP address in a fixed list of around 970,000 addresses [6]. We apply our model to Skitter data. Fig. 13 shows the cumulative number of unique links observed, and Fig. 14 shows the disappearance period of links. The curves have the same shape as BGP data does and our model fits the data well, which means that the model can be applied

⁵iPlane improves the conversion by first doing alias resolution of IP addresses and then mapping each IP address to the AS that majority of the addresses in the alias map to.

to the topology dynamics observed by traceroute too. However, the observed parameters may not reflect the real rates of birth and death due to the data set's limitations.

First, Skitter's revelation parameters are different from that obtained from BGP data. Its λ is less than the half of BGP's, and its μ is about twice as BGP's. These differences can be explained as follows. When a routing change happens, a BGP monitor will be notified by triggered routing updates. But for a traceroute monitor to see the change, it must probe the path when the change is still in effect. For instance, if a hidden backup link is exposed for 2 hours and the monitor probes once every 24 hours, the chance to discover this link is only $2/24 = 8\%$. Since in the Internet, primary paths (links) are being used in the majority of time and Skitter only probes once every day, the chance of observing backup paths (links) is small. Therefore, Skitter is slow in discovering backup links (*i.e.*, small λ), but quick in picking up recovered primary links (*i.e.*, large μ). This is further verified by examining the links observed by both BGP and Skitter. Fig. 15 shows that, for the same link, how much the difference between BGP's and Skitter's timestamps is when they discover the link for the first time. BGP discovers 60% of links earlier than Skitter, Skitter discovers 10% of links earlier than BGP, and they discover 30% of links on the same day. Fig. 16 shows the difference between BGP's and Skitter's timestamps when they see a particular link for the last time. BGP observes 50% of links for longer time, Skitter observes 10% of links for longer time, and they see 40% of links on the same day for the last time. Clearly, to be more effective in observing topology dynamics, traceroute data collection needs to probe destinations frequently⁶.

Second, Skitter has higher death rate (57.61/day) than birth rate (39.86/day), which implies the observed topology shrinks every day, and we have verified it by counting the number of unique nodes and links observed by Skitter everyday. This is a result of using a *fixed* list of destination IP addresses over several years. Using a fixed destination list underestimates the birth rate because new ASes (which announce new prefixes) and their links will not be probed. It also overestimates the death rate because over time some IP addresses on the fixed list become unreachable due to various reasons. For example, a noticeable percentage of prefixes stop being announced over time [27], or an ISP may decide to block ICMP traffic. Fig. 17 clearly shows the decline of reachable addresses in Skitter's destination list. Many links that Skitter no longer observes are due to its shrinking probing scope, not because the links are dead. To be effective in observing topology dynamics, tracer-

⁶High probing frequency may cause security concerns from networks being probed.




Figure 17: Number of reachable addresses in Skitter destination list.




Figure 18: Number of collected links in DIMES.

oute data collections must update the destination list constantly to include all ASes.

3.4.3 DIMES and iPlane

DIMES [29] is a distributed measurement infrastructure consisting of a large number of agents installed by end users on their computers. These agents periodically traceroute or ping a given set of destinations. The number of DIMES agents has been growing rapidly, from a few hundreds in early 2005 to almost 12,000 in late 2006. Each DIMES agent probes a different destination list, which is computed and updated periodically [29]. The probing frequency varies for different agents, and can be as low as once per week. The purpose of using partial destination lists and low probing frequency for each agent is to distribute the measurement load among agents while maintaining a good coverage of the topology. Fig. 18 shows the cumulative number of links observed by DIMES over time. Although its overall trend is similar to BGP's and Skitter's, the curve is much more irregular due to the fast growth of agents, partial and changing destination lists, and different probing frequencies.

iPlane [23] employs around 200 monitors installed on PlanetLab nodes [15]. All iPlane monitors use a common destination list extracted from BGP routing tables daily, and probe one address inside every /24 address block. The destinations are being probed once every day by every monitor. iPlane's measurement scheme seems very promising in terms of observing topology dynamics. However, its data is available only from late June, 2006 and onward, not long enough for our study.

In summary, traceroute data is an important source of topology information. We can potentially use it to enhance our understanding of the liveness problem. However, existing data sets that we have examined do not seem suitable for studying topology dynamics due to a few limitations. To be effective in capturing topology dynamics, we must first find an effective means to accurately con-

vert router paths to corresponding AS paths. Furthermore, traceroute monitors collectively should have an observation view that is *broad* (i.e., destinations represent all ASes), *fresh* (i.e., destination list is updated) and *frequently probed*.

3.5 Discussion

The model developed in this section captures the main characteristics of observed AS topology changes by three dynamic processes: birth, death, and revelation. This model can help us obtain key parameters of different dynamic processes and separate real topology changes from transient routing changes with a given confidence. At the same time, we must also understand the limitations of this model.

The model is mainly descriptive and not derived from first principles. It matches the data well and is useful in studying topology dynamics. However, it is possible to have other models that also fit the data well, e.g., node birth may be modeled by an exponential function with a small exponent [20]. Our model does not provide an explanation for why birth and death rates are constant, or why the revelation probability are uniform. Answering these questions is likely to require looking deep into economic, technological, and operational factors behind the Internet evolution, and our model can serve as an important input to such a study.

We may also articulate the reasons that made the model work well. One sound reason relates to the model's macroscopic granularity and the large scale of the Internet. Individual factors influencing the AS topology evolution are probably not constant or uniformly distributed, however, given the large scale of the Internet, and the large number of (perhaps independent) factors in action, fluctuations caused by individual factors may even out when we measure macroscopic properties using data aggregated from many different views. Had we restricted our study to a small geographic area, or to a subgraph of the Internet (e.g., an academic network), we might have obtained very different results. The model also may not apply to very different time scales, e.g., looking at link/node appearances at hourly basis (or over decades) might show a different pattern. We develop the model based on data from recent three years, and we are yet to see whether the model will hold in the future. Our model also makes some simplified assumptions. For instance, in reality different links may have different revelation probabilities (λ and μ), depending on their connectivity and routing policies.

The apparent success of the model should not be taken without caution; rather, it is a call for additional verifications. Besides BGP data, we use registry and traceroute data to show evidence that the observed data is indeed explained by the model. However, it remains a research challenge regarding how to perform a thorough validation, given the lack of public information about ISP connectivity changes. One specific validation step we plan to carry out in the near future is to contact a number of ISPs to verify the accuracy of the link birth/death timestamps inferred by the model.

4. APPLICATIONS

This section uses three applications to illustrate the importance of solving the liveness problem. First, knowing the liveness of links and nodes helps obtain a more complete topology and its properties more accurately. Second, the birth and death dates inferred from our empirical model can be used to evaluate theoretical models of topology evolution. Third, the same data of birth and death dates can be used to empirically characterize topology growth trends.

We would like to note that, although these applications demonstrate the usefulness of our model, the accuracy of the numerical results may be affected by the raw data available to us. Results




Figure 19: Trade-off between liveness and completeness for topology snapshot.




Figure 20: Fraction of multi-homed customers.




Figure 21: Attachment probability distribution for a target node degree.

in this section are based on BGP data, the *ALL* data set. Recently He *et al.* [32] point out that topology derived from BGP data may miss a significant number of peer-to-peer links⁷. We plan to incorporate quality data from other sources when they become available.

4.1 More Accurate View of the Topology

The AS topology on a particular day is often referred to as a topology *snapshot*. Understanding the graph properties of a topology snapshot (static view) and the changes of the properties over time (dynamic view) is an active research area. Knowing the liveness of links and nodes can help capture more accurate views of the topology.

4.1.1 Static View

Some previous work (*e.g.*, [17], [14], [25]) obtain a topology snapshot by extracting AS nodes and links from BGP routing tables of a *single day*. This approach yields an incomplete static view of the topology. Besides the invisible links, which the monitors are unable to capture, there are many hidden links that can be captured, but are missing in the routing tables on the particular sampling day. One way to obtain a more complete topology snapshot on day t is to include live links and nodes appeared in routing tables and routing updates of recent past, *i.e.*, since day $t - L$. The value of L depends on how confident we want to be that the links added are still alive on day t . For instance, $L = 0$ means that we have 100% confidence that all the links are alive on day t , however the topology will be rather incomplete. As L increases, the topology becomes more complete, however the confidence on link liveness decreases as the topology may contain links that are already dead. By adjusting the value of L we can make trade-offs between the liveness and the completeness of the resulting topology snapshot. Fig. 19 shows the number of links in a topology snapshot of November 30, 2006, as a function of the liveness confidence. For instance, a point (0.6, 75000) represents a snapshot with 75000 links in which all links have more than 60% chance of being alive on November 30, 2006. The liveness confidence is obtained by $(1 - \text{death confidence})$, where death confidence is calculated using the equivalent of Eq. 3 for disappearance⁸. Depending on the liveness confidence we want to put in the snapshot, the number of

⁷The topology in [32] is collected from BGP routing tables, IRR, and traceroute data. Our topology is collected from both BGP routing tables and *routing updates*. A comparison of the two topology data sets on May 12, 2005 shows that our set has 9154 additional links but misses 7056 links.

⁸The gap in the curve from $x = 0.85$ to $x = 1$ is due to using the parameters extracted from BGP data, Eq. 3 does not have a solution for death confidence lower than 15% with time resolution of a day.

links in the topology graph can vary from about 64000 to more than 88000.

4.1.2 Dynamic View

In order to study how graph properties of the topology change over time, the effects of revelation process on the observed topology must be taken into consideration. As an example, assume we want to measure the percentage of multi-homed stub ASes over time. A stub AS is the one that *always* appears as the last AS in an AS path; it corresponds to a customer network at the bottom tier of the Internet routing hierarchy. A stub AS may connect to multiple network service providers, but it does not forward data traffic between its providers. Fig. 20 shows the percentage of multi-homed stub ASes without considering the effects of revelation process: starting with an initial snapshot on the first day, the topology is updated every day by adding links as they first appear and removing links according to their last-seen time, and the percentage of stub ASes is calculated and plotted for each day's topology. The curve shows a fast increase at the beginning of the study period and a fast decrease at the end, which might look puzzling at first, but can be easily explained by the revelation process. At the beginning, there are many hidden links that take some time to appear, and as they appear, we discover that existing single-homed stub ASes are in fact multi-homed. Near the end, many in-sight links become out-of-sight, and are prematurely discounted from the topology graph, resulting in false classification of multi-homed stub ASes into single-homed. To take into account the effects of revelation process, in Fig. 20 we draw two vertical lines corresponding to the 95% confidence margins calculated from Eq. 3. Only the part of the curve between these two vertical lines reflects the real percentage of multi-homed stub ASes with a high confidence level. This example illustrates the importance and usefulness of our revelation model in topology measurement and other similar types of topological studies.

4.2 Evaluating Theoretical Models

A number of theoretical models have been proposed for network topology evolution. They generally model the decision process of where to add new links/nodes into the topology and which old links/nodes to be removed from the topology over time. Due to the lack of empirical data on link/node births and deaths, these models have been evaluated based on the outcome of the evolution, instead of the decision process itself. For example, a common evaluation method is to let the network grow to a certain size by simulating the link/node addition and removal processes, then compare graph properties of the resulting network topology with that of an observed topology of a similar size. This approach to evaluation,




Figure 22: Model evaluation

(per day)	Transit	Stub	Total
Node Birth	2.4	8.3	10.7
Node Death	0.8	2.5	3.3
Net Growth	1.6	5.8	7.4
Link Birth	37.7	29.2	66.9
Link Death	29.0	16.7	45.7
Net Growth	8.7	12.5	21.2

Figure 23: Comparison Between Stub and Transit changes.




Figure 24: Node net growth

however, is limited in its effectiveness because different evolution processes can generate topologies that share certain graph properties. A much better approach is to compare the decisions of the theoretical models directly with the observed link/node births and deaths, which is made possible once the liveness problem is solved. In this section, we demonstrate the value of this new approach by applying it to evaluate two theoretical models.

Barabasi *et al.* [8] proposed an evolution model (the BA model) that explains the emergence of the power-law distribution of node degrees in complex networks. One of the key elements in this model is the *preferential attachment* or *rich-get-richer* paradigm. The basic idea is that new nodes tend to connect to existing nodes that already have high degrees. More precisely, according to the BA model, a new node attaches to an existing node i at time t with the probability $p_i(t) = d_i(t) / \sum_j d_j(t)$, where $d_i(t)$ is the degree of node i at time t , and the summation is over all the existing nodes in the topology. Thus, the probability that a new node attaches to a node with degree d at time t is

$$P(d, t) = N(d, t) \cdot \frac{d}{\sum_j d_j(t)}$$

where $N(d, t)$ is the number of existing nodes with degree d at time t . Fig. 21 shows the distribution of $P(d, t)$ on a particular day from BGP data, and the vertical line at $d_{med} = 49$ represents the median of the distribution. Since the distribution of node degrees is heavy-tailed, *i.e.*, there exists a small number of nodes with high degrees and a large number of nodes with small degrees, on average a new node has a higher chance to attach to a small degree node than attaching to a high degree node. To evaluate the BA model, we extract the birth and death events of links and nodes from BGP data using 90% confidence margins provided by our model. In other words, we record a topological change only when we are at least 90% sure that it is a real topological change. Appearances and disappearances near the beginning and the end of the study period are discarded to eliminate the effect of revelation process. One main problem in evaluating a probabilistic theoretical model is that, on the one hand, we need a large number of node births to make a meaningful sample set; on the other hand, the degree distribution changes over time as nodes join the topology, therefore each sample of node degrees is only good for a specific instance in time. To overcome this problem, we use the following evaluation scheme inspired by the Monte Carlo method: (1) Compute the distribution $P(d, t)$ for day t and its median $d_{med}(t)$. (2) For each node born on day $t + 1$, check if the node it connects to has a degree higher or lower than $d_{med}(t)$; if it is higher, increment counter h , otherwise increment counter l . (3) Repeat for every day. By the law of large numbers, if the node-join process follows the BA model, the ratio $h/(h + l)$ should converge to 50%.

Fig. 22 plots the value of $h/(h + l)$ over time and shows that the BA model actually converges to 58% in the long run. This indicates that, during the evolution of Internet AS topology, high degree nodes attracted more new nodes than what the BA model describes. This result is also consistent with the conclusion in [14], which used an earlier data set and different evaluation techniques.

Bu *et al.* [9] proposed the Generalized Linear Preference (GLP) model, in which the probability that a new node attaches to a node with degree d is given by:

$$P(d, t) = N(d, t) \cdot \frac{d - \beta}{\sum_j (d_j(t) - \beta)}$$

where $\beta \approx 0.8$. Fig. 22 plots the result of applying our technique to the GLP model. Note that Fig. 22 only checks for the median of the $P(d, t)$ distribution, and it shows that GLP model matches the median of the empirical data for node attachment. To further evaluate the model, we need to compare other percentiles of the distribution as well.

4.3 Characterizing Evolution Trends

We use the same data of link/node births and deaths to empirically characterize the trends of topology evolution. Generally speaking, ASes can be classified into two categories: *stub* and *transit*. A stub AS *only* appears as the last AS in an AS path, while a transit AS will appear in the middle of some AS paths. A stub AS corresponds to a customer network, which does not forward traffic between its neighbors. A transit AS corresponds to a network service provider, which provides data delivery service for other networks. In the context of the AS topology graph, we refer to them as *stub nodes* and *transit nodes*, respectively. Links between transit nodes are called *transit links*, and links between transit and stub nodes are called *stub links*⁹.

Provider networks and customer networks are fundamentally different business entities in the Internet, and our data shows that their growth trends in the Internet topology are also very different. Fig. 23 shows the breakdown of birth, death, and net growth rates for transit and stub nodes of the topology, using our model with 90% confidence margin. It is clear that most node dynamics, including birth, death, and net growth, happen mainly to the stub nodes. In particular, one may note that the stub nodes' net growth rate is 3.6 times of the transit nodes'. Fig. 24 shows the node net growth on a daily basis, with the curves already compensated by the number of hidden and out-of-sight nodes (which is minimal in the case of nodes). These results indicate that the size of the Inter-

⁹There also exist links between stub nodes, usually these links are not observable in BGP data as they are not announced to other ASes.




Figure 25: Link net growth




Figure 26: Net growth of node wirings.




Figure 27: Frequency of link changes

net, as measured by the number of AS nodes, is rapidly increasing mainly due to new customers joining the Internet.

Fig. 23 also shows the breakdown of link growth rates, and Fig. 25 plots the daily link net growth, with the curves already adjusted to compensate the impact of the hidden and out-of-sight links. Although the transit nodes only make about 28% of the total nodes and their percentage is decreasing, their link birth rate is 29% higher than that of stub links, and the death rate 74% higher. Note that the link birth and death counts lump together all the link changes that can be either due to node birth/death or due to link adjustment between *existing* nodes. To quantify the latter, we define *wiring* as the addition of a new link between two existing nodes, and *unwiring* as the removal of a link between two live nodes (*i.e.*, the two incident nodes are still alive in the topology after the link is removed). Stub wirings and unwirings reflect customer’s actions of changing providers, while transit wirings and unwirings reflect provider’s actions of adjusting their inter-ISP connectivity.

Fig. 26 shows the net growth of wirings between July 24, 2004 and May 25, 2006, a period that falls within our 90% confidence margins. Here we consider only the nodes that are present throughout this measurement period and find 10349 such stub nodes and 6319 transit nodes. We then count the wiring and unwiring events among these nodes. Fig. 26 shows that provider networks get more densely connected over time, perhaps as a result of serving ever increasing customer demands¹⁰. Keeping in mind that the number of transit nodes is much lower than the number of stub nodes, and that a transit wiring event means adding a link between two transit nodes only, we see a high net growth rate in transit links, despite the high death rate shown in Fig. 23. This confirms a general observation that over recent years, provider networks have been actively adjusting their inter-connectivity.

What types of inter-ISP connectivity that providers are actively adjusting? Based on the type of inter-AS business relationships, links are generally classified into three classes: customer-to-provider (c2p), peer-to-peer (p2p), and sibling-to-sibling (*e.g.*, ASes that belong to the same company). Since the sibling relationship is relatively rare, here we focus on c2p and p2p links only. In c2p relationship, the customer (or lower-tier ISP) pays its provider (or upper-tier ISP) to gain the global reachability. In p2p relationship, data traffic is exchanged free of charge between the two peers, but only traffic originated from a peer AS (or its customers) is allowed on this link. We applied the PTE algorithm [18] to infer AS link relationships, and the algorithm was able to infer the relationships for 75% of all the links involved in wirings and unwirings. We clas-

sify wiring and unwiring events according to their link types. For each type, we calculate the time interval between two consecutive events and plot the distribution of the intervals in Fig. 27. We can see that, among c2p links, stub links are more stable compared to c2p transit links, and that the p2p links have much shorter intervals between their connectivity adjustments than all c2p links. According to [30], a common ISP operational practice is to set up a p2p link and re-evaluate it periodically (*e.g.*, once every few months). Based on whether the p2p link helps reduce the overall cost, it may be either kept or terminated. Equipped with the empirical data of link births and deaths, we are able to identify and quantify this phenomenon in the global routing system.

5. RELATED WORK

One of the early work in AS topology measurement studies is [19], in which topology snapshots are constructed by using BGP routing tables and updates within 21 days. Since then many topology characterization results have been obtained through using topology snapshots from either BGP routing tables or traceroute data (*e.g.*, [17, 24]). However, as pointed by [10] and [28], the existing snapshot approach may miss a significant number of links, especially the peer-to-peer links between ASes at lower tiers. To collect a more complete AS topology, we can either gather topology information from additional sources (*e.g.*, IRR, looking glasses, exchange points) [33, 25, 32, 10], or place additional monitors at different places [29]. Different from the above mentioned efforts which focus on the completeness problem, our work reported in this paper focuses on the topology liveness problem. However our results can be used to help capture more complete topologies with a given confidence on link/node liveness.

There are a number of topology evolution models. For example, the work reported in [7, 9, 34] models the addition of nodes and links so that the resulting topology can have certain desired graph properties, such as degree distribution, clustering coefficient, rich-club connectivity and betweenness centrality. Chang *et al.* [12] extend the “heuristically optimized trade-off” (HOT) model [16] to allow each AS to have multiple geographical locations; more recently in [13] the authors also identify a set of criteria used in the decision process when new peering relationships are established. Wang *et al.* [31] propose an evolution model based on random walks and wealth evolution.

Our empirical model differs from the previous topology evolution models in that we model how fast the topology changes as measured by the birth and death rates of nodes and links, while previous work models where the changes happen. We consider the effects of transient routing dynamics, which has not been addressed by any previous effort. Using our empirical model, we can obtain

¹⁰Other relevant factors, such as increases in link capacity and the number of BGP sessions between neighboring ASes, are not observable in BGP data.

the dates of individual node or link changes, which can then be used to evaluate theoretical models. Our work is complementary to previous efforts and can help advance realistic modeling of the Internet topology and its evolution.

Chen *et al.* [14] evaluate the BA model using BGP data. They ignore the first ten days and the last ten days of data to reduce the impact of the topology revelation, instead of addressing the liveness problem directly. The choice of ten days is mainly based on intuitive judgement. Our model provides a means to choose the threshold with a target confidence level. Huston [20] studies the consumption of AS numbers and shows that AS allocation over time can be fitted by either a linear function or an exponential function with a small exponent. This is similar to our linear modeling of node birth. Our model is more complete in that, in addition to birth, it also includes the death process and revelation process, and can be applied to links' birth, death, and revelation as well.

6. CONCLUSION

In this paper we formulated the topology liveness problem and contrasted it with the completeness problem that has been studied intensively by the previous efforts. We developed an empirical model that characterizes the changes in the observed AS topology by three processes: birth, death, and revelation. Based on the model, we are able to distinguish real topology changes from transient routing changes with a given confidence level. The success of the model calls for further investigations on why the evolution of the Internet, a vast distributed and autonomous system, can be approximated well by such a seemingly over-simplified model.

Besides providing better views of the topology and helping validate topology evolution models, our results of the topology evolution trends can be an important input to the design of future Internet routing architecture. We show that the main contributor to the growth of the Internet size is the fast increasing number of customer networks, which puts a heavy burden on the global routing system's scalability. Our results also show the highly active connectivity adjustment among transit networks and the rapid growth in density among them. A future global routing system may take into account the distinct growth patterns of these two types of networks in order to achieve long-term scalability.

Acknowledgements

We would like to thank RouteViews (David Meyer, Mike Witt), RIPE (Vasco Austuriano), DIMES team (Yuval Shavitt, Eran Shir, Ela Malinowski), CAIDA (Bradley Huffaker), iPlane team (Harsha Madhyastha) and UC Riverside (Yihua He, Georgos Siganos) for making their data available to us. We are also grateful to Dan Massey, Rafit Izhak-Ratzin, Mohit Lad and Eric Osterweil for their input to this work. Finally, we thank our shepherd Sharad Agarwal and anonymous reviewers for their valuable comments.

7. REFERENCES

- [1] Internet Routing Registry. <http://www.irr.net/>.
- [2] Regional Internet Registry data. <ftp://www.ripe.net/pub/stats>.
- [3] RIPE routing information service project. <http://www.ripe.net/>.
- [4] RouteViews routing table archive. <http://www.routeviews.org/>.
- [5] Skitter AS adjacency list. http://www.caida.org/tools/measurement/skitter/as_adjacencies.xml.
- [6] Skitter destination list. <http://www.caida.org/analysis/topology/macroscope/list.xml>.
- [7] R. Albert and A.-L. Barabasi. Topology of evolving networks: local events and universality. *Physical Review Letters*, 85(24):5234–5237, 2000.
- [8] A.-L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [9] T. Bu and D. Towsley. On distinguishing between Internet power law topology generators. In *Proc. of IEEE INFOCOM*, 2002.
- [10] H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger. Towards capturing representative AS-level Internet topologies. *Elsevier Computer Networks Journal*, 44(6):737–755, 2004.
- [11] H. Chang, S. Jamin, and W. Willinger. Inferring AS-level Internet topology from router-level path traces. In *Proc. of SPIE ITCOM*, 2001.
- [12] H. Chang, S. Jamin, and W. Willinger. Internet connectivity at the AS-level: an optimization-driven modeling approach. In *Proc ACM SIGCOMM MoMeTools workshop*, 2003.
- [13] H. Chang, S. Jamin, and W. Willinger. To peer or not to peer: modeling the evolution of the Internet's AS-level topology. In *Proc. of IEEE INFOCOM*, 2006.
- [14] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origin of power-laws in Internet topologies revisited. In *Proc. of IEEE INFOCOM*, 2002.
- [15] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman. Planetlab: an overlay testbed for broad-coverage services. *ACM SIGCOMM Computer Comm. Review (CCR)*, 33(3):3–12, 2003.
- [16] A. Fabrikant, E. Koutsoupias, and C. Papadimitriou. Heuristically optimized tradeoffs: a new paradigm for power laws in the Internet. In *Proc. of ICALP*, 2002.
- [17] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *Proc. of ACM SIGCOMM*, 1999.
- [18] L. Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, 2001.
- [19] R. Govindan and A. Reddy. An analysis of Internet inter-domain topology and route stability. In *Proc. of IEEE INFOCOM*, 1997.
- [20] G. Huston. AS Number Consumption. *The ISP Column*, September 2005.
- [21] Y. Hyun, A. Broido, and kc claffy. On third-party addresses in traceroute paths. In *Proc. of Passive and Active Measurement Workshop (PAM)*, 2003.
- [22] L. Li, D. Alderson, W. Willinger, and J. Doyle. A first-principles approach to understanding the Internet's router-level topology. In *Proc. of ACM SIGCOMM*, 2004.
- [23] H. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: an information plane for distributed services. In *Proc. of OSDI*, 2006.
- [24] P. Mahadevan, D. Krioukov, K. Fall, and A. Vahdat. Systematic topology analysis and generation using degree correlations. In *Proc. of ACM SIGCOMM*, 2006.
- [25] 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. *ACM SIGCOMM Computer Comm. Review (CCR)*, 36(1):17–26, 2006.
- [26] Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz. Towards an accurate AS-level traceroute tool. In *Proc. of ACM SIGCOMM*, 2003.
- [27] X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, and L. Zhang. IPv4 address allocation and the BGP routing table evolution. *SIGCOMM Computer Comm. Review (CCR)*, 35(1):71–80, 2005.
- [28] D. Raz and R. Cohen. The Internet dark matter: on the missing links in the AS connectivity map. In *Proc. of IEEE INFOCOM*, 2006.
- [29] Y. Shavitt and E. Shir. DIMES: Let the Internet measure itself. *ACM SIGCOMM Computer Comm. Review (CCR)*, 35(5):71–74, 2005.
- [30] W. B. Norton. The art of peering: the peering playbook, 2002.
- [31] X. Wang and D. Loguinov. Wealth-based evolution model for the Internet AS-level topology. In *Proc. of IEEE INFOCOM*, 2006.
- [32] Y. He, G. Siganos, M. Faloutsos, S. V. Krishnamurthy. A systematic framework for unearthing the missing links: measurements and impact. In *Proc. of NSDI*, 2007.
- [33] B. Zhang, R. Liu, D. Massey, and L. Zhang. Collecting the Internet AS-level topology. *ACM SIGCOMM Computer Comm. Review (CCR)*, 35(1):53–61, 2005.
- [34] S. Zhou and R. J. Mondragón. Accurately modeling the Internet topology. *Physical Review E*, 70(6):066108, 2004.