Modeling BGP Table Fluctuations

Ashley Flavel, Matthew Roughan, Nigel Bean and Olaf Maennel

School of Mathematical Sciences
University of Adelaide

Abstract. In this paper we develop a mathematical model to capture BGP table fluctuations. This provides the necessary foundations to study short- and long-term routing table growth. We reason that this growth is operationally critical for network administrators who need to gauge the amount of memory to install in routers as well as being a potential deciding factor in determining when the Internet community will run out of IPv4 address space.
We demonstrate that a simple model using a simple arrival process with heavy tailed service times is sufficient to reproduce BGP dynamics including the "spiky" characteristics of the original trace data. We derive our model using a classification technique that separates newly added or removed prefixes, short-term spikes and long-term stable prefixes. We develop a model of non-stable prefixes and show it has similar properties in their magnitude and duration to those observed in recorded BGP traces.

1 Introduction

The Border Gateway Protocol (BGP) [1] automatically discovers paths within the Internet, allowing end-hosts to communicate, whilst respecting policy requirements of Autonomous Systems (ASs). However, events such as link failures, newly added networks and policy changes can alter the path towards a particular destination in routers throughout the Internet, which in turn changes traffic flow and performance [2]. There exists a need for network operators to understand which events may lead to performance disruptions and traffic shifts [3–5] or whether a change in routing configuration may lead to an unforeseen interaction between policies [6,7]. Despite this, the properties of routing updates [8–10] and the extent to which routers actually scale [11,12] are still poorly understood. Several questions operators are still struggling to answer [13] include:

– What is the maximum BGP table a router can handle?
– How much memory is needed to store the Forwarding Information Base (FIB) on the line-cards?
– What are the future hardware requirements for routers?

One of the reasons answers to these questions are missing, is that it is hard to create field conditions for realistic tests [14]. Efforts within the IETF, for example from the Benchmarking Methodology Working Group (BMWG) [15], try to overcome the discrepancy between the field and testing conditions by recommending metrics and test setups for test-beds. However, we argue that a good model of BGP events is necessary to understand BGP dynamics, which in turn will lead to the development of superior test tools as well as improved estimation of future Internet trends such as IPv4 address space usage.

In contrast to existing tools [16] we create a model characterizing BGP table fluctuations that is extendible to estimate the size of a BGP table at any particular point in time along with confidence interval estimates which are needed to understand the likelihood of extreme fluctuations. A typical example often cited in literature is the incident in April 1997 [17] where $AS7007$ accidentally announced almost all prefixes in the Internet (belonging to all other ASes) for approximately two hours. Although events of such a magnitude are rare, our data analysis shows that short-term fluctuations in the order of up to several thousand prefixes are not abnormal. It is unclear whether hardware limitations [12, 14], protocol interactions [18], BGP implementations [19] or other factors [20] are to blame for this behavior. Consequently, measuring magnitudes of previous events without understanding their underlying nature will not predict the likelihood of future events. A mathematical model, however, is capable of being trained using current behavior, to provide insight into possible future behavior, e.g. the likelihood of large routing events.

Large routing events can have a serious effect on routers and potentially even cause service interruptions. Nowadays, routers in the Internet have a specially designed data structure to store forwarding information. This Forwarding Information Base (FIB) is often stored in separate memory across the various line-cards to improve packet lookup times and thus forwarding performance (see for example the design of the Cisco GSR [21]). The FIB itself needs to be constructed from routing information comprising manual router configuration (including static routes), Interior Gateway Protocols (IGPs) and the BGP table or Routing Information Base (RIB). Typically the RIB contributes the largest proportion of the FIB table. It is also an "unknown factor", as a network administrator has limited control over what is learned from the outside. However, if the memory limit on the line-cards is reached, the router cannot perform its designed tasks and service outages occur.

In Section 4 we present a new classification technique to separate prefix behavior. By studying recorded RIBs, and applying our classification technique, we derive statistical properties of routing tables and in Section 5 we introduce a model capable of capturing the RIBs short-term dynamics. We concentrate our efforts on the previously un-examined short-term fluctuations, however, our model provides a mechanism to fully characterize all changes and predict the future components of the BGP table based on its current state.

2 Background and Related Work

Routing in the Internet is accomplished on a per-prefix basis. Routing protocols, such as OSPF and IS-IS, are used to find the shortest path internally within an AS, while BGP [1] is used to exchange reachability information between ASs. As BGP is a policy-routing protocol, it gives operators some freedom to express their company requirements and policies. To accomplish this, BGP allows attachment of several attributes for each route, and is based on a path-vector protocol. Upon startup, a router establishes sessions with all configured neighbors and all appropriate table-entries are exchanged. Hence, each neighbor sends its RIB to adjacent routers, which in turn store each table in memory (sometimes referred to as the Adj-RIB-In). Next, the router may modify or filter attributes in the

Adj-RIB-In before selecting a single "best path" used to create its own RIB. The RIB, together with static and IGP routes, are combined to form the Forwarding Information Base (FIB). The FIB is typically a high-speed lookup data structure which consists of prefix-next-hop pairs enabling forwarding of packets to the appropriate next-hop: Special memory is used for storage of the FIB directly on the line-cards of the routers. The selected best routes, if not subjected to out-bound filtering, are then propagated to other neighbors.

BGP's flexibility, coupled with the fact that network administrators (mis-)use BGP in numerous ways means it is often difficult to determine the underlying cause of routing behavior [9]. As BGP propagates changes to the best path, a single router may send multiple updates based on one triggering event [8]. Further, the propagation of one update may cause induced updates at other locations [22]. It is even possible for policy conflicts to occur that can potentially disrupt the entire Internet [7] which has led to a considerable body of research (see [10] for further details).

Pioneering work related to the size of the Internet RIB was undertaken by Fuller *et al.* [23] who measured the number of routes in the RIB on a monthly granularity over the period 1988-1992. Additionally, Huston [24] has used information obtained from the University of Oregon's RouteViews project [25] to display the long-term growth of numerous Autonomous Systems' RIBs since 1994 [26]. In contrast, our work examines the fluctuations within the Internet RIB on a fine timescale.

Several *prefix*-clustering techniques based on time correlations between updates have been previously described [27,28] while other methods which cluster *updates* based on their likelihood of being caused by a single event have also been considered [22,29]. Clustering updates can provide insight into underlying features, however, determining boundaries of clusters can be a difficult task. In contrast, we describe a simple classification technique in Section 4 which has clearly defined boundaries separating prefixes with different behaviors.

3 Data Sets

RouteViews [1] uses software [31] capable of conducting BGP sessions with routers throughout the Internet to collect BGP data. RouteViews archive all routing information and make their data publicly available to benefit the entire Internet community. Fig. 1 provides an overview of the data used in the ensuing analysis and unless otherwise stated, we use the **Verio-Trace** as an example dataset throughout this paper.

As we cannot obtain the actual RIB of any router in the Internet, we need to approximate a *potential* RIB from the available raw data. Each monitored router (or peering router) provides a snapshot of their perspective every 2 hours (hence the granularity of [24]). In addition to the snapshots, updates with timestamps to the nearest second are recorded. A snapshot, with updates applied in chronological order, provides a finer granularity than simply using 2 hourly snapshots. The *potential* RIB is *not* an exact representation of the RIB on a remote

[1] RIPE [30] collects similar data to that in RouteViews with a European focus. For this investigation we have only used RouteViews data.

Name	Verio-Trace	Verio-Trace-Prediction
Start Time (UTC)	1 June 2004 01:23:03	1 August 2004 01:48:06
Finish Time (UTC)	1 August 2004 01:48:06	1 October 2004 00:20:56
Start RIB size	137820 entries	140484 entries
Finish RIB size	140484 entries	129191 entries

Fig. 1. Detailed Data Sources Information: All traces obtained from RouteViews [25] for a Verio router (IP: 129.250.0.85, AS: 2914)

router, it is however a good approximation. Route monitors typically do not collect internal prefixes of remote peers and the Minimum Route Advertisement Interval (MRAI) [1] reduces the frequency of update messages recorded. We assume that a majority of routing table fluctuations occur externally and on time scales longer than 30 seconds (a typical setting for MRAI). Hence, we argue our approximation provides a good representation to an actual RIB stored on a remote router. Further, session resets between the route monitor and peering router can cause discrepancies between the constructed and recorded tables. We use intermediate table dumps to infer where missing withdrawals are likely to have occurred, and the impact of these session resets for the data analyzed in this paper is minimal.

4 Classification

Fig. 2 depicts the change in size of the RIB on a per second basis over a two-month interval. Huston [24] has previously shown that the RIB experiences a growth trend when monitored at hourly intervals. However, visual examinination of the time series for the number of entries in the RIB reveals several other key features such as "spikes" which are not evident at a coarser granularity. These spikes indicate short-term availability of prefixes and occur in highly varying magnitudes and durations. In Fig. 2 (b) a large upward spike in the early morning hours (UTC) is clearly visible. The short-term event consisted of approximately $2,000$ prefixes which appeared in the routing table for less than 10 minutes.

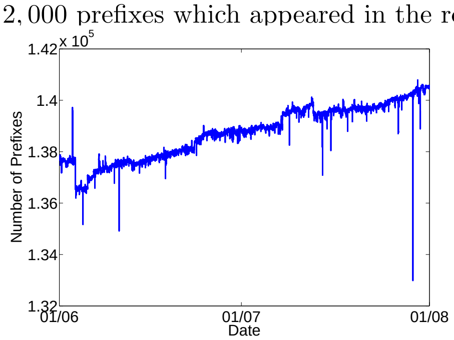

(a)

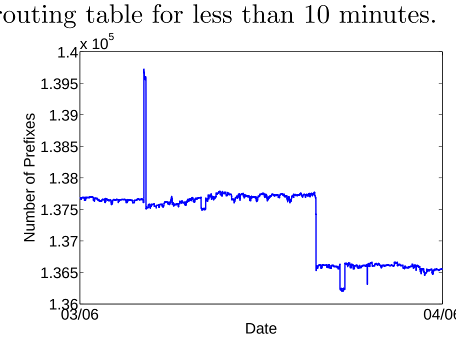

(b)

Fig. 2. Short-term RIB fluctuations from **Verio Trace** for (a) 1 June 2004 - 1 August 2004 and (b) a closer examination of 3 June 2004

Recall that our end-goal is to predict the size of the RIB at any given point in time in the future, as well as estimate intermediate table fluctuations. However, the observed time series are a complex combination of various components (trend, upward spikes and downward spikes). To achieve a clean simple model, we need to extract the various components so that we can identify their key characteristics and build models. Thus, in this section we develop a simple, straightforward classification technique of the raw BGP data that, at the same time, eases the construction of our model.

Given a number of observations at times $t_1, ..., t_n$, we first require our model to predict the number of table entries at some future time t_{n+i}. Then we are able to estimate the probability a router reaches a predefined memory limit between t_n and t_{n+i}. To achieve this, we need to know:

1. How many new prefixes are added?
2. What happens to existing prefixes?
3. What short-term changes prefixes exhibit?

Consequently, we need to understand the behavior of the RIB over the entire interval $[t_1, t_{n+i}]$. This behavior leads us to the following classifications:

Definition 1 (Stable Prefix). *Let RIB_t be the set of prefixes for which we have an explicit route at time t. If a prefix $p \in RIB_{t_1}$ and $p \in RIB_{t_2}$ then $p \in S^{[t_1,t_2]}_{stable}$ over time interval $[t_1, t_2]$.*

Within the RIB, there is a large proportion of prefixes which are permanently or almost permanently routable. Definition 1 is designed to separate the majority of prefixes which exhibit this behavior from the entire set of prefixes. A stable prefix is present within the RIB at the start and end of the time interval under consideration. Consequently, the number of stable prefixes can never exceed the initial (and final) number of prefixes. Stable prefixes can, however, leave the RIB during the time interval. As a result, the number of stable prefixes within the RIB captures the downward spikes within the time series (see Fig. 3 (a)). Not surprisingly, a large proportion of prefixes within the RIB are classified as stable.

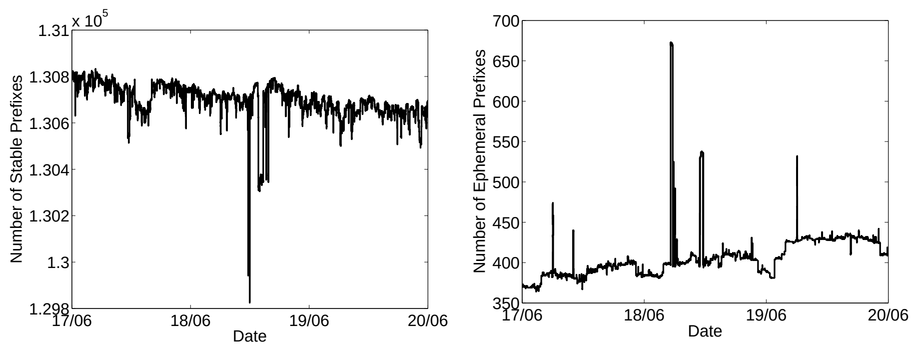

(a) (b)

Fig. 3. (a) Stable prefixes and (b) Ephemeral prefixes from 17 - 20 June 2004 classified on interval 1 June - 1 August 2004 from **Verio-Trace**

Definition 2 (Transient Prefix). *If a prefix $p \in RIB_{t_1}$ and $p \notin RIB_{t_2}$ or $p \notin RIB_{t_1}$ and $p \in RIB_{t_2}$ then $p \in S^{[t_1,t_2]}_{transient}$ over time interval $[t_1, t_2]$.*

New prefixes are announced and others permanently withdrawn when new networks are created and old networks aggregated or dismantled. The transient prefixes as defined in Definition 2 captures these newly announced and withdrawn prefixes. As previously shown [24], the RIB grows over time and thus the transient prefixes capture this long term growth (plot not shown as we focus on the short term process in this paper).

Definition 3 (Ephemeral Prefix). *If a prefix $p \notin RIB_{t_1}$ and $p \notin RIB_{t_2}$ and $p \in RIB_t$ for some $t \in (t_1, t_2)$ then $p \in S^{[t_1,t_2]}_{ephemeral}$ over time interval $[t_1, t_2]$.*

A majority of routing dynamics are caused by a minority of prefixes [32–34]. We aim to separate these prefixes with our definition of ephemeral prefixes. These prefixes as shown in Fig. 3 (b) tend to have short lifetimes although we show later that their lifetimes are best modeled as a heavy-tailed distribution. Consequently, ephemeral prefixes form the upward spikes within the RIB timeseries and may cause router line-cards to exceed their memory limit (before any expected long-term trend would exceed the bound).

All classifications described above are merely approximations to what operators may intuitively label as stable, transient or ephemeral. The times t_1 and t_2 are arbitrary times, hence the interval $[t_1, t_2]$ can be of any length. If we let the interval be infinite, every prefix would be classified as ephemeral. Conversely, if $t_1 = t_2$, every prefix would be classified as stable. In Section 5, we demonstrate that the approximate nature of the classifications, and the arbitrary time interval is not vital to our model's success.

5 Model

In this section we develop our model. We focus on the ephemeral prefixes, as they are most relevant in terms of short-term memory consumption. As described above, the classification of ephemeral prefixes is conceptually aimed at separating those prefixes which experience short-term presence in the RIB.

We define a *spike of prefixes* as a group of prefixes which arrive and depart at the same time. Visual detection of these spikes is somewhat trivial, however determining a simple and effective technique to automatically identify such spikes is not. The difficulty arises as updates relating to multiple events may overlap.

Definition 4 (A Spike of Prefixes). *The spike, $S_{a,w}$, is the set of prefixes which enter the table at time a and leave at time w where $a \leq w$. The spike duration is $r = w - a$ and the spike size is the number of prefixes in the set $S_{a,w}$.*

Updates are witnessed to occur in bursts which may last a number of seconds, after which the MRAI timer prevents any announcements from being sent, but updates caused by a single event may not be advertised or withdrawn at exactly the same time. As Cisco uses a jittered 30 seconds as their default MRAI, we use bins of 30s to capture the start and conclusion of spikes, sacrificing the ability to detect spikes at a finer granularity than 30s, however, large spikes which take several seconds to announce or withdraw are identified as a single set which is especially critical in identifying the probability of large spikes.

Our model assumes that the three components of spikes, (1) spike arrival times; (2) spike sizes; and (3) spike durations, are independent. Superimposing independent spikes defined by the three components above forms the basis for our model to predict the future statistical properties of the RIB.

5.1 Spike Arrival Times

Fig 4 (a) depicts the Complementary Cumulative Distribution Function (CCDF) for spike arrival times. It can be seen that the arrival times are uniformly distributed across the interval. Also, the mean inter-arrival time between spikes can be shown to be approximately 19.6 seconds. In this section, the uniform distribution of spike arrival times is satisfactory for our purposes. In Section 6, however, we require a spike arrival process. Given the difficulty in accurately determining the actual arrival time of each announcement, let alone spike, the precise form of the spike arrival process is impossible to identify. For simplicity, we have chosen to use the classical telecommunications arrival process, namely the Poisson process, to model the spike arrival process.

5.2 Spike Sizes

Fig. 4 (b) shows the CCDF for the spike sizes on log-log axes. We can see that the distribution of the size of spikes is heavy-tailed and consequently, we model the size of spikes as a Pareto distribution. We estimate parameters using logarithmic transformed data (dashed line). Note that, although we use the Pareto distribution, we do not claim this is the ideal distribution as there remains some disparity between the actual and fitted curves. Note a single outlier is responsible for the large discrepancy from the model in bottom right of the plot. We do, however, assert the distribution is heavy-tailed, and the Pareto distribution is a simple, parsimonious distribution with this feature.

5.3 Spike Durations

Recall from Definition 3, ephemeral prefixes are not present at the start and end of the time period. They may experience multiple lifetimes within the time period $[t_1, t_2]$ through multiple announcements and withdrawals. Consequently, some prefixes with multiple short lifetimes will provide more data points than other prefixes that have long lifetimes and contribute one or few data points. For our model we assume independence between events. The CCDF, plotted on log-log axes, for the lifetimes of ephemerals is shown in Fig. 4 (c).

The lifetimes of ephemeral prefixes are artificially limited by the size of the time period we use for classification. As a result, the CCDF representing the lifetimes of ephemeral prefixes is dependent on the arbitrary choice of time period $[t_1, t_2]$. Thus, we require a model to separate the censorship or truncation effect caused by the arbitrary choice of time interval and the underlying process defining the lifetimes of ephemeral prefixes.

We consider, without loss of generality, an arbitrary time interval $[0, T]$ and assume the start times are distributed according to the uniform distribution on $[0, T]$. Hence the probability an individual arrival occurs before time u obeys the probability distribution function

$$Pr\{t \leq u\} = u/T. \tag{1}$$

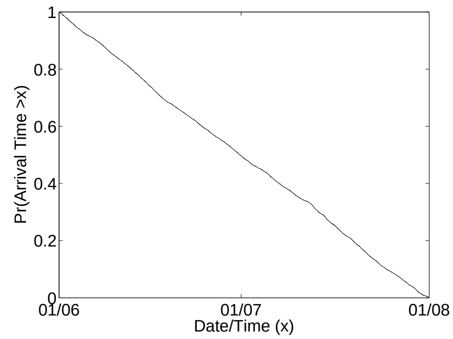

(a)

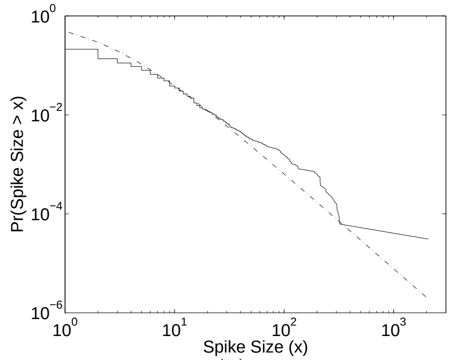

(b)

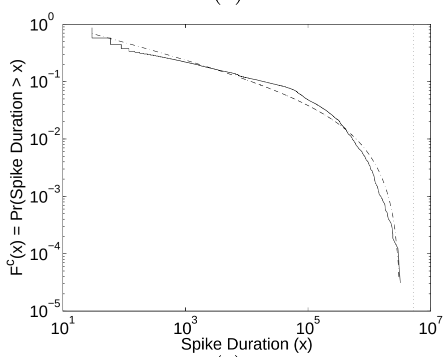

(c)

Fig. 4. Spike components from **Verio-Trace**. (a) Spike Arrival time CCDF (b) Empirical (solid) and Pareto Model (dashed) of Ephemeral Spike Size CCDF ($Mode = 2.26, Exponent = 1.93$). (c) Empirical (solid) and Truncated Pareto Model (dashed) of Ephemeral Spike Durations ($Mode = 8.24, Exponent = 0.28$).

We also assume independent durations of spikes, r and let a, w denote the start and stop times, respectively. If we further assume the duration of spikes are Pareto distributed, then the probability density function is given by

$$f(x) = \tfrac{cb^c}{x^{c+1}}, \quad x \geq b, \tag{2}$$

where b is the mode and c is the exponent. However, a prefix will only be classified as ephemeral if it is absent at the end of the classification interval $[0, T]$, i.e., $w \leq T$. So the distribution we observe is the conditional distribution

$$F^c(x) = Pr\{r \leq x | w \leq T\} = \frac{Pr\{r \leq x \cap w \leq T\}}{Pr\{w \leq T\}}. \tag{3}$$

First, consider the numerator, so

$$\begin{aligned} Pr\{r \leq x \cap w \leq T\} &= Pr\{r \leq x \cap a + r \leq T\} \\ &= \int_b^x Pr\{r = u\} Pr\{a \leq T - u\} du. \end{aligned} \tag{4}$$

Substituting (1) and (2) into (4) yields

$$\begin{aligned} Pr\{r \leq x \cap w \leq T\} &= \int_b^x \left(\frac{cb^c}{u^{c+1}} \frac{T-u}{T} \right) du \\ &= \frac{T-b}{T} - \frac{T-x}{T} \left(\frac{b}{x} \right)^c - \frac{b^c}{T} \left(\frac{x^{c+1} - b^{-c+1}}{1-c} \right). \end{aligned} \tag{5}$$

The normalization factor $Pr\{w \leq T\}$ is the particular case of (4) where $x = T$. Thus, if we set

$$Q_{b,c}(x) = \frac{T-b}{T} - \frac{T-x}{T}\left(\frac{b}{x}\right)^{c} - \frac{b^c}{T}\left(\frac{x^{c+1} - b^{-c+1}}{1-c}\right) \tag{6}$$

then from (3)

$$F(x) = Pr\{r \leq x | w \leq T\} = \frac{Q_{b,c}(x)}{Q_{b,c}(T)}. \tag{7}$$

Using a nonlinear regression based on (7) on a log scale, we are able to estimate parameters for ephemeral prefixes to fit our model CCDF ($F^c(x) = 1 - F(x)$) to the empirical data. The example shown in Fig 4 (c) demonstrates that our model is successful in capturing the distribution for the lifetimes of ephemerals including the truncation caused by the classification. Furthermore, a Pareto distribution with parameters estimated using our *truncated* model of spike durations provides an intuitive description for the *non-truncated* duration short-lived prefixes spent in the table.

6 Results

Fig. 5 (a) shows the timeseries for the number of ephemeral prefixes in **Verio Trace**. Two model generated time series (Fig. 6) based on fitted parameters of **Verio Trace** found in Section 5, contain the same features as in the empirical data (Fig. 5 (a)). The empirical data contains a single large spike of magnitude greater than 2000 prefixes. The model generated time series in Fig. 6(a) also predicts (in one case) a large spike of magnitude greater than the limits of the plot. These large spikes are particularly important from a modeling perspective as they potentially cause memory capacity problems in router line cards. Our model provides the predictive ability to determine the probability an abnormally large spike will occur. Also, many short duration spikes and few spikes of several hundred prefixes that last from seconds to weeks are witnessed. Most clearly seen in Fig. 6(b) is performance of the model when capturing overlapping different duration spikes (i.e. the 15 day period in June). An obvious artifact caused by the classification technique is the 'bump' in the center of the timeseries. As seen in each realization (Fig. 6) our model is able to successfully account for this.

Recall our goal is to predict statistical characteristics of the future timeseries as shown in Fig. 5 (b). Other than a single spike lasting approximately 20 days towards the end of the timeseries, the Figs. 5 (a) and (b) are similar: - they have similar spike sizes, durations and truncation effect (the 'bump'). Note we are not aiming to predict exact locations of spikes, but rather their statistical characteristics. We believe *based only on parameters estimated from* **Verio-Trace**, our model is able to predict the statistical properties of future short-term fluctuations in the RIB.

The marginal distribution for our model is shown in Fig. 7. We used 500 model realizations to find a numerical approximation to the mean, maximum and minimum marginal distributions. The empirical data from **Verio-Trace** and **Verio-Trace-Predicion** are also plotted, both of which fall inside the ranges for the marginal distribution. It is thus arguable that our model is able to reproduce the highly varying nature in the number of ephemeral prefixes.

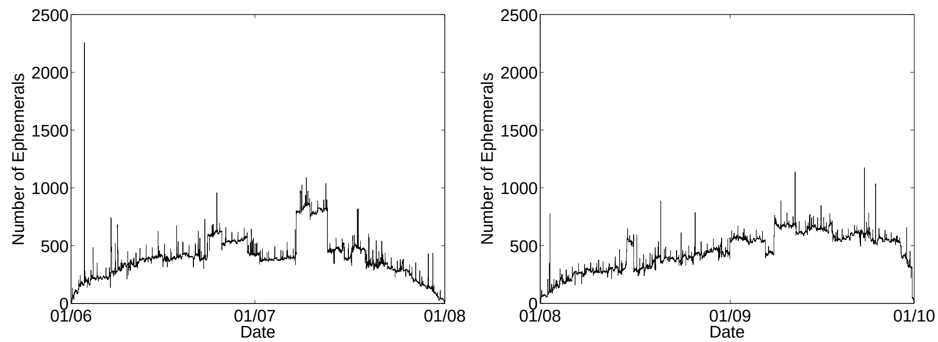

(a) **Verio-Trace** (b) **Verio-Trace-Prediction**

Fig. 5. Empirical number of ephemeral prefixes: Each plot demonstrates how the number of ephemeral prefixes within the RIB changes over two 2 month periods.

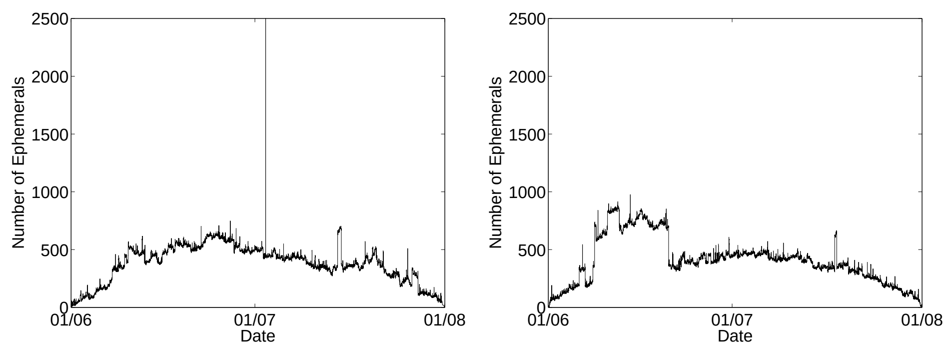

(a) (b)

Fig. 6. Model generated number of ephemeral prefixes: Each plot demonstrates a different realization of the model for the number of ephemeral prefixes within the RIB over a 2 month period. The parameters for the model are obtained from **Verio-Trace**.

More validation of the model is needed, but space limitations prevent us presenting more extensive results. These preliminary results indicate our model is capable of encapsulating the non-stable prefixes and extendable to the entire set of routable prefixes. We reiterate that although we use the Pareto distribution for lifetimes and sizes of spikes, we do not claim that this is the ideal distribution to use. However, we do assert that they exhibit heavy-tailed properties and the Pareto distribution is just one common distribution which has this property.

7 Conclusion and Future Work

In this paper we presented a classification technique to separate long-term growth trends from short-term state changes arising from newly added or removed BGP prefixes. We demonstrated the efficacy of our technique over a 2-month time interval using RouteViews BGP data. Our analysis confirmed the results of previous work such as [34] which supported the validity of our model. We further

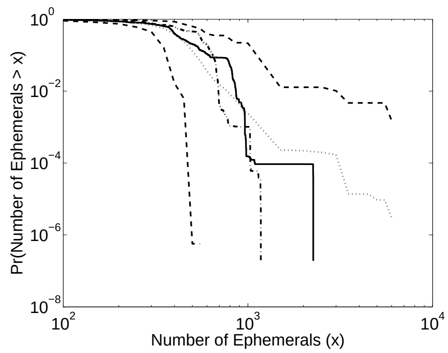

Fig. 7. Mean marginal distribution (dotted) for the number of ephemeral prefixes in RIB found from 500 realizations of our model. Also shown is the minimum and maximum marginal distributions (dashed) together with the marginal distribution for the empirical data from **Verio Trace** (solid) and **Verio-Trace-Prediction** (solid-dotted).

elicited the presence of heavy-tailed features in the number of short-term fluctuations in terms of size and duration.

The main contribution of this paper was to demonstrate that a simple arrival process with heavy-tailed service times is sufficient to capture BGP routing dynamics. To this end, we derived the parameters for our model from observable BGP data and successfully reproduced BGP table fluctuations including the "spiky" characteristics of the original traces.

In future work we will expand our model to estimate long-term table growth as well as the probability that short-term spikes do not exceed a fixed number of table entries. Also, we will investigate changing the classification interval start times and durations together with considering the evolution of prefixes as they enter the table. This issue is critically important in order to successfully predict when the Internet community will run out of IPv4 address space as well as the amount of memory needed for BGP tables on router line-cards.

References

1. Rekhter, Y., Li, T.: A Border Gateway Protocol 4 (1995) RFC 1771.
2. Li, J., Bush, R., Mao, Z.M., Griffin, T., Roughan, M., Stutzbach, D., Purpus, E.: Watching Data Streams Toward a Multi-Homed Sink Under Routing Changes Introduced by a BGP Beacon. PAM (2006)
3. Teixeira, R., Shaikh, A., Griffin, T.G., Rexford, J.: Dynamics of Hot-Potato Routing in IP Networks. In: Proc. ACM SIGMETRICS. (2004)
4. Teixeira, R., Shaikh, A., Griffin, T.G., Voelker, G.M.: Network sensitivity to hot-potato disruptions. In: Proc. ACM SIGCOMM. (2004)
5. Teixeira, R., Duffield, N.G., Rexford, J., Roughan, M.: Traffic matrix reloaded: Impact of routing changes. In: Proc. (PAM). (2005)
6. Griffin, T.G., Shepherd, F.B., Wilfong, G.: Policy Disputes in Path Vector Protocols. In: Proc. ICNP. (1999)
7. Griffin, T.G., Huston, G.: BGP Wedgies (2005) RFC 4264.
8. Mao, Z.M., Bush, R., Griffin, T.G., Roughan, M.: BGP Beacons. In: Proc. ACM IMC. (2003)
9. Griffin, T.G.: What is the Sound of One Route Flapping? (2002) IPAM.

10. (T. G. Griffin Interdomain routing links)
http://www.cl.cam.ac.uk/users/tgg22/interdomain/.
11. Chang, D.F., Govindan, R., Heidemann, J.: An empirical study of router response to large BGP routing table load. In: IMW'02: Proc. 2nd ACM SIGCOMM Workshop on Internet measurment. (2002)
12. Agarwal, S., Chuah, C.N., Bhattacharyya, S., Diot, C.: Impact of BGP dynamics on router CPU utilization. In: Proc. (PAM). (2004)
13. Jaeggli, J.: NANOG 39 BOF: Pushing the FIB limits, perspectives on pressures confronting modern routers. (http://www.nanog.org/mtg-0702/jaeggli.html)
14. Feldmann, A., Kong, H., Maennel, O., Tudor, A.: Measuring BGP Pass-Through Times. In: Proc. (PAM). (2004)
15. (IETF Benchmarking Methodology Working Group (bmwg))
http://www.ietf.org/html.charters/bmwg-charter.html.
16. Maennel, O., Feldmann, A.: Realistic BGP traffic for test labs. In: Proc. ACM SIGCOMM. (2002)
17. Misel, S.A.: (Wow, AS7007!) http://merit.edu/mail.archives/nanog/1997-04/msg00340.html.
18. Griffin, T.G., Wilfong, G.: An analysis of BGP convergence properties. In: Proc. ACM SIGCOMM. (1999)
19. Labovitz, C.: Scalability of the Internet backbone routing infrastructure. In: PhD Thesis, University of Michigan. (1999)
20. Wetherall, D., Mahajan, R., Anderson, T.: Understanding BGP misconfigurations. In: Proc. ACM SIGCOMM. (2002)
21. (Cisco 12000 Series Routers)
http://www.cisco.com/en/US/products/hw/routers/ps167/.
22. Feldmann, A., Maennel, O., Mao, M., Berger, A., Maggs, B.: Locating Internet Routing Instabilities. In: Proc. ACM SIGCOMM. (2004)
23. Fuller, V., Li, T., Yu, J., Varadhan, K.: Supernetting: an Address Assignment and Aggregation Strategy (1992) RFC 1338.
24. Huston, G.: Analyzing the Internet BGP Routing Table. The Internet Protocol Journal **4**(1) (2001)
25. (University of Oregon RouteViews project) http://www.routeviews.org/.
26. Huston, G.: (http://bgp.potaroo.net)
27. Zhang, J., Rexford, J., Feigenbaum, J.: Learning-Based Anomaly Detection in BGP Updates. in Proc. of SIGCOMM Workshops (2005)
28. Andersen, D., Feamster, N., Balakrishnan, H.: Topology Inference from BGP Routing Dynamics. In: 2nd ACM SIGCOMM Internet Measurement Workshop, Boston, MA (2002)
29. Caesar, M., Subramanian, L., Katz, R.H.: Root cause analysis of Internet routing dynamics. Technical report, UCB/CSD-04-1302 (2003)
30. (RIPE's Routing Information Service)
http://www.ripe.net/ris/.
31. Ishiguro, K.: (Zebra routing software) http://www.zebra.org/.
32. Rexford, J., Wang, J., Xiao, Z., Zhang, Y.: BGP routing stability of popular destinations. In: Proc. ACM IMW. (2002)
33. Wang, L., Zhao, X., Pei, D., Bush, R., Massey, D., Mankin, A., Wu, S.F., Zhang, L.: Observation and analysis of BGP behavior under stress. In: Proc. ACM IMW. (2002)
34. Huston, G.: The BGP Instability Report (2006) http://bgpupdates.potaroo.net/instability/bgpupd.html.