

IPv4 Address Allocation and the BGP Routing Table Evolution

Xiaoqiao Meng*, Zhiguo Xu*, Beichuan Zhang*, Geoff Huston‡, Songwu Lu*, Lixia Zhang*

*Computer Science Dept., UCLA
Los Angeles, CA 90095
{xqmeng,zhiguo,bzhang,slu,lixia}@cs.ucla.edu

‡APNIC
Brisbane, Australia
gih@apnic.net

ABSTRACT

The IP address consumption and the global routing table size are two of the vital parameters of the Internet growth. In this paper we quantitatively characterize the IPv4 address allocations made over the past six years and the global BGP routing table size changes during the same period of time. About 63,000 address blocks have been allocated since the beginning of the Internet, of which about 18,000 address blocks were allocated during our study period, from November 1997 to August 2004. Among these 18,000 allocations, 90% of them started being announced into the BGP routing table within 75 days after the allocation, while 8% of them has not been used up to now. Among all the address blocks that have ever been used, 45% of them were split into fragments smaller than the original allocated blocks; without these fragmentations, the current BGP table would have been about half of its current size. Furthermore, we found that the evolution of BGP routing table consists of both the appearance of new prefixes and the disappearance of old prefixes. While the change of the BGP routing table size only reflects the combined results of the two processes, the dynamics of either process is much higher than that of the BGP table size. Finally, we classify routing prefixes into *covering* and *covered* ones, and examine their evolution separately. For the covered prefixes, which account for almost half of the BGP table size, we infer their practical motives such as multihoming, load balancing, and traffic engineering, etc., via a classification method.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Operations—*Network management, Network monitoring* ; C.2.5 [Computer-Communication Networks]: Local and Wide-Area Networks

General Terms

Measurement

This work is partially supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. N66001-04-1-8926 and by National Science Foundation (NSF) under Contract No. ANI-0221453. Any opinions, conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the DARPA or NSF.

Keywords

BGP, IPv4 address allocation

1. INTRODUCTION

The global Internet has experienced tremendous growth over the last decade. One evidence of the Internet growth is its global routing table size, which is nearly tripled between January 1998 and August 2004. In this paper we collect data from IPv4 address allocation records and BGP routing table logs to measure the IPv4 address allocation and its impact on the global routing table size.

We obtained IPv4 address allocation records from the four regional Internet registries (RIRs) [1]: ARIN, RIPE, APNIC, and LACNIC. These four RIRs are responsible for allocating IPv4 address blocks to Internet service providers (ISPs); ISPs in turn assign IP addresses from their allocated blocks to small ISPs, organizations, and end users. Allocated address blocks, represented by address prefixes, are utilized when the prefixes are advertised into the global routing system, which uses BGP as the inter-domain routing protocol [10].

We obtained BGP routing table data from Oregon RouteViews project [2]. RouteViews operates a number of BGP data collection points which peer with BGP routers at various ISPs. RouteViews started archiving BGP routing tables on November 8, 1997. To assess the relation between IP address allocation and BGP routing table growth, we collect both data during the period of November 8, 1997 to August 31, 2004. In the rest of this paper we call all the address allocations made during this period *new allocations*.

BGP routing table data shows that, when a prefix is advertised into the routing system, it does not necessarily match the size of an allocated address block. A prefix may represent an address block as allocated, a fragment of an allocated address block, or an aggregation of multiple allocated address blocks. Moreover, the address space represented by an advertised prefix can be a sub-block of another existing prefix. In this case we call the former a *covered* prefix, and the latter a *covering* prefix. Covered prefixes are typically used to achieve certain specific load balancing or traffic engineering goals, which will be further discussed in Section 6.

Our study consists of three steps. We first characterize the usage of newly allocated IP address blocks, including the usage latency, the number of unused allocations, and the

percentage of prefixes advertised as identical, fragmented, and aggregated. We then analyze the impact of allocation on the growth of the routing table size. We finally examine the changes of the BGP routing contents over time. Our main findings can be summarized as follows:

1. About 63,000 IP address blocks have been allocated since the beginning of the Internet, and about 18,000 of them were allocated between November 8, 1997 and August 31, 2004. Over 90% of these 18,000 allocations showed up in BGP routing table within 75 days, while about 8% of them have not been used.
2. When newly allocated address blocks were advertised into the routing system, about 43% of them represented each allocated block as a single prefix, 45% of them were fragments of the allocated blocks, and the remaining 12% were aggregates with separately allocated address blocks. Overall, address fragmentation significantly increased the BGP table size. If none of the allocations made in the last 6 years were fragmented, the current BGP table size would have been half of its current size.
3. The BGP table size increase reflects the combined effect of advertisement of new prefixes and disappearance of old prefixes. 60%-70% of address blocks led to new prefix advertisement since they were allocated, and the percentage decreased to about 18% when the age of the allocated blocks reached 6 years. On the other side, for a certain fraction of allocated address blocks, some of their generated routing prefixes will disappear from the BGP table. Such a fraction gracefully increased with the age of the allocations. More specifically, when the age increased to 6 years, the percentage slowly grew to 15%.
4. A large number of covering prefixes were fragments of allocated address blocks, and the covered prefixes appear and disappear more dynamic than covering prefixes. We further classified the covered prefixes into 7 different categories based on their advertisement paths relative to that of their corresponding covering prefixes. The classification help understand the practical motives behind advertising the covered prefixes.

The rest of the paper is organized as follows. Section 2 provides brief background information about IPv4 address allocation. Section 3 presents several statistics for address allocation, the allocation policies and the BGP table growth. Section 4 analyzes the usage behavior of allocated address blocks by using multiple metrics. Section 5 analyzes the BGP table size growth from the allocation perspective. Section 6 analyzes the covering covered prefixes separately. Section 7 presents the related work. Section 8 concludes the paper.

2. BACKGROUND

In this section, we briefly give background on BGP routing table and IPv4 address allocation.

2.1 BGP routing table

BGP is the *de facto* standard for inter-domain routing in the global Internet. It uses a *prefix*, a 32-bit IP address with a mask length, to identify a block of IP addresses that reside in the same destination network. For example, 10.0.0.0/8 represents a block of contiguous IP addresses ranging from 10.0.0.0 to 10.255.255.255. A BGP routing table contains information about how to reach all the advertised prefixes (or destination networks) on the Internet. Each entry in the routing table contains, among other information, a prefix and an AS path attribute, which is a list of ASes through which the prefix can be reached.

2.2 Address allocation

The original IP design divided address space into three different address classes – Class A, B, and C. Each class had a fixed boundary between the network-prefix and the host-number within the 32-bit IP address. However, because there were a very limited number of Class A address blocks, and Class C blocks were too small for most organizations, Class B address blocks were allocated most of the time and approached exhaustion. Therefore the Internet Registry started allocating multiple Class C blocks to individual service providers, resulted in rapid global routing table growth in a short time period.

Classless Inter-Domain Routing (CIDR) was introduced around 1993-1994. CIDR enables a flexible boundary between the network-prefix and the host-number field in the IP address space by using an address mask parameter. For example, 64.4.176.0/20 represents a network prefix of 64.4.176.0 with 20 bit prefix length. This prefix contains 2^{12} unique IP addresses. CIDR's flexible address allocation enables more efficient allocation of the address space. The deployment of CIDR has effectively slowed down the growth of the global routing table.

A typical allocation record reads as follows:

```
“arin|US|ipv4|24.220.0.0|65536|19981115|allocated”
```

The above record indicates that ARIN allocated a block of IPv4 addresses to an ISP in US on November 15, 1998. The address block starts from 24.220.0.0 and contains 65536 unique addresses, implying that the mask length is 16 bits. This allocated address block can be represented by 24.220.0.0/16.

2.3 BGP Data Collection

RouteViews collects BGP tables from up to 65 peering BGP routers, although this set of BGP peers have changed over time. Since each BGP router only has a partial view of the global BGP reachability, to maximize the observation of all the advertised prefixes, we merged the BGP tables collected from all RouteViews' peers into one.

Another part of our study is the evolution of the covering and covered prefixes in BGP table. In this case, we use the BGP views provided by two BGP routers that are more persistently archived by the Route-Views project.

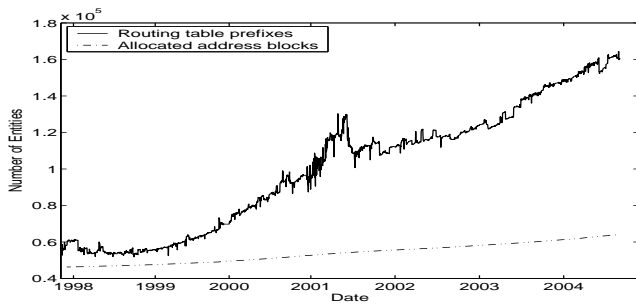


Figure 1: Number of allocated blocks vs. BGP routing table size.




Figure 2: Allocated address space vs. address space contained by BGP table.

3. ADDRESS ALLOCATION AND BGP ROUTING TABLE GROWTH

Figure 1 shows the number of allocated address blocks and the BGP table size over the last six years. As one can see, the BGP table size is much larger than the number of allocated address blocks, reaching around 160,000 by August 31, 2004. Furthermore, the gap between the number of routing table entries and the number of allocated blocks has been steadily increasing over time. In early 1998, the table size is 1.33 times of the number of the allocations; by end of August 2004, this ratio increased to 2.54.

Figure 2 compares the amount of IPv4 address space that has been allocated against that being announced in the BGP table over the last 6 years. The figure shows that about 75% to 77% of allocated address space is announced in the global routing system, and this ratio has been fairly stable over our study period.

Figure 3 shows the distribution of the address allocation sizes during four sample time periods: before 1993, in 1997, in 1999, and in 2004 (until August 31). This figure shows clearly the impact of the two major allocation policy changes in the past. The first policy change is the adoption of CIDR during 1993-1994. Comparing the allocation size distribution before 1993 and after 1993, we see that /16 and /24 were two dominant allocation sizes before 1993, but no longer so after 1993. The second policy change occurred when the minimum allocation size was reduced from /19 to /20 in year 2000. The most popular allocation size was /19 in 1999, and /20 in 2004. We refer to [11] for more detailed study on the impact of allocation policy changes on the BGP table.

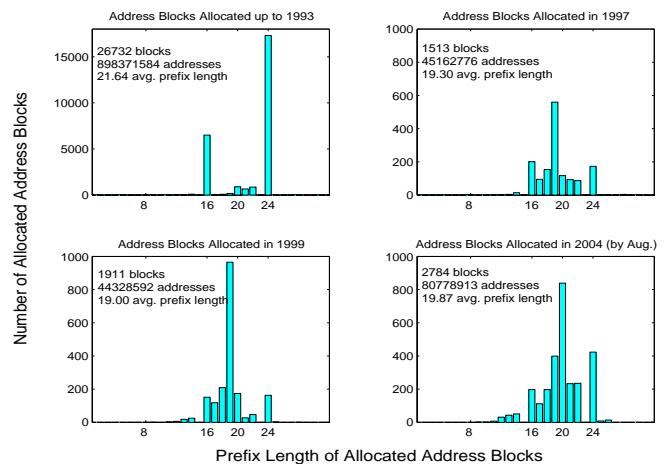


Figure 3: Distribution of allocation size




Figure 4: Histogram of usage latency

4. USAGE OF ADDRESS ALLOCATION

Allocated IP address blocks become globally reachable when they are advertised as routing prefixes in the BGP routing table. In this section, we measure the advertisement of address blocks allocated during our study period [11/08/1997, 08/31/2004] to answer the following three questions:

1. How long does it take a newly allocated block to be advertised in the routing table?
2. Are there unused allocated blocks?
3. When allocated blocks show up in the routing table, are they fragmented, aggregated with block of different allocations, or identical to the allocations?

4.1 Usage latency

We define *usage latency* of an allocated address block as the time interval between the allocation time and the first time a part of, or the entire, allocated block shows up in the BGP routing table. Among the total 18,000 (18K) address blocks allocated during [11/08/1997, 08/31/2004], 5% of them were seen in the routing table before their recorded allocation time. After checking with ARIN, we believe this phenomenon is most likely due to the incompleteness of




Figure 5: Allocation time and allocation size for unused address blocks

RIRs’ data recording. The address allocation files we obtained only record the (last) allocation time for each address block, and do not mark the *reallocation* of address blocks that were returned to RIRs after an earlier allocation. Although rare, some allocated IP address blocks do get returned to RIRs occasionally, and these returned address blocks can then be reallocated. Thus the observed usage of an address block can be due to its previous allocation even though RIRs only recorded the reallocation time. There are also 1494 allocated address blocks (8% of the total blocks) that have not been used, which we discuss in the next section.

We plot the histogram of the usage latency for the remaining 87% of the allocations in Figure 4. There are 11,622 allocated blocks advertised in the BGP table immediately after their allocation. The average usage latency is 52 days, 90% of allocated blocks are advertised within 75 days, and the longest usage latency is 2158 days.

4.2 Unused address allocations

As we mentioned earlier, 1494 address blocks allocated after November 1997 have not been used. Since some of these blocks were allocated recently, they may get used in near future. Based on our measurement result that over 90% of allocations are used within 75 days, we choose January 1, 2004 as the threshold and consider only the 134 allocations that were made before January 1, 2004 but were never used by August 31, 2004 as *unused* address allocations. We then measure the distribution of these 134 unused allocations in terms of the time period length since the allocation time and the block size, as shown in Figure 5. The most popular size of unused allocated blocks is /24. One conjecture, which is confirmed by some ISPs, is that these allocations may be for localized infrastructure assignments (such as Internet exchange points) which can receive a small public address allocation under the RIRs’ allocation policies, yet they are not globally routed and thus do not need to be advertised to the global routing system.

4.3 Advertisement patterns

We measured the advertisement patterns of all the allocated blocks during [11/08/1997, 08/31/2004] by using BGP routing table logs collected at RouteViews. Our data showed that 45% of them were fragmented when they were advertised in the routing table, 43% of the allocated blocks were advertised as allocated, and the remaining 12% were adver-




Figure 6: Routing table size change

tised as aggregates with address blocks from different allocations. Since the latter two advertisement patterns do not contribute much to the growth of the BGP table size [11], in the following we elaborate on address block fragmentation.

Figure 6 plots BGP table size with and without the new allocations and fragmentation of the new allocations. Instead of taking daily snapshots of the table size as in Figure 1, we plotted Figure 6 by using weekly moving average to reduce sharp spikes in daily snapshots (which were likely caused by BGP log file corruption or BGP misconfigurations). The curve “Actual BGP table” in the figure shows that the table size was increased from about 55K on January 1, 1998 to 160K August 31, 2004. If we exclude all the routing prefixes generated from the new allocations, which are the allocations made during [11/08/1997, 08/31/2004], the table size would be about 62K entries on August 31, 2004, as shown by the curve “No new allocation” in the figure. In other words, while $62K - 55K = 7K$ routing prefixes of the 105K total table size increase are from address blocks allocated before 11/08/1997, the other 98K prefixes are from the new allocations. On the other hand, if we assume all the new allocations were advertised to the routing system without being fragmented, we get the curve “New allocation not fragmented”, and the routing table size would be about 82K on August 31, 2004, which is half of its actual size. In overall, about $\frac{155K - 82K}{105K} \approx 70\%$ of the total 105K table size growth in the past six years is due to fragmentation of the new allocations.

5. EXAMINATION OF BGP ROUTING TABLE GROWTH

In this section, we examine the changes of individual entries in the BGP routing table. Our data show that the routing table growing process is a combined result of both new prefix advertisements and old prefix disappearances.

5.1 Examination of BGP table content changes

We divide the study period into 13 time intervals of 6-month each and compare the routing table entries at the beginning and end of each interval. If a prefix existed at the beginning of the interval but not at the end, we counted it as an *old prefix disappearance*. If a prefix did not exist at the beginning but




Figure 7: Breakdown of BGP table growth at six-month time interval.

showed up at the end of the interval, we count it as a *new prefix appearance*. Due to dynamics in the global routing, some routing prefixes may appear and disappear frequently. To reduce potential measurement bias caused by such short-term dynamics, we collect the routing tables for seven consecutive days at the beginning of each 6-month interval and compare each of them with the corresponding routing table after the interval. Figure 7 shows the average number of new prefix appearances and old prefix disappearances for each of the half year intervals. The lower end of the error bar shows the minimum number of each 7-day measurement, and the upper end the maximum number. This figure shows that the number of new prefix appearances in each interval is much higher than the BGP table size change, which is the difference between the number of new prefix appearances and old prefix disappearances.

A newly advertised prefix does not necessarily increase the size of reachable IP address space in the routing table, because the address space represented by the new prefix may have already been covered by another existing prefix. For example, from January 2004 until July 2004, 36K new prefixes were advertised while 24K old prefixes were removed. The IP addresses in 78% of these new prefixes have already been covered by existing prefixes. Only the other 22% new prefixes brought 59 million (59M) more addresses to the routing table. On the other hand, one might concern whether the disappeared prefixes take away reachable IP addresses. Our data shows that the address space covered by 90% of the aforementioned 24K disappeared prefixes remained reachable through other prefixes, and the other 10% of the disappeared prefixes brought 36M addresses away from the routing table. In overall, the total reachable addresses increased by roughly 23M between January 2004 and July 2004.

5.2 New allocations contribution to BGP table growth

To measure the BGP routing table growth due to new address allocations, we use the same 13 six-month time intervals as mentioned earlier, and for each of the 13 intervals, calculate the fraction of the new prefix advertisement and old prefix disappearance from the address blocks allocated during that interval. We plot the results in Figure 8, where the X-axis is the allocation age, defined as the elapsed time between the allocation and the appearance or disappear-

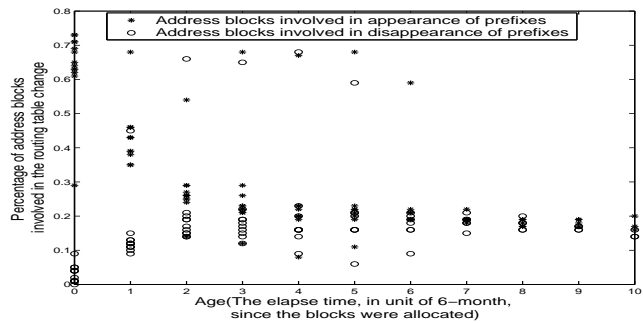


Figure 8: Activity of allocation vs. age.

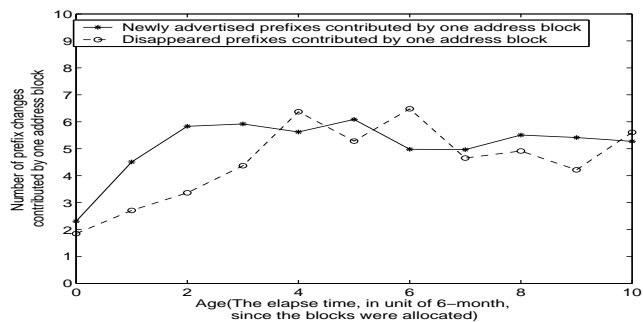


Figure 9: # of newly appeared or disappeared prefixes associated with single allocation (on average).

ance of the corresponding prefixes as measured in unit of six-month, and the Y-axis represents the fraction of allocations that are involved in new prefix advertisement (or old prefix disappearance). Each star (circle) represents the percentage of all the allocations (made within a certain six-month interval) with the corresponding age as indicated by the X-axis.

Figure 8 shows that 60%-70% of newly allocated address blocks bring new routing prefixes to BGP table within the same 6-month interval they were allocated. The percentage decreases gradually over the age of the allocation and gets close to 18% when the age reaches 5 years; because our study period is 6-year long, our data set does not have reasonable number of data samples for observations beyond 5 years. On the other side, few disappeared prefixes in an interval belong to the address blocks allocated during that interval, this the percentage of the allocated blocks associated with disappeared prefixes is low. When the age of the allocations reaches 5 years, this percentage slowly increases to 15%, but always stays below the percentage of allocations associated with new prefix advertisement. We can state this observation in two different ways. (1) an address block allocated some time ago is likely to be associated with new prefix advertisement than with old prefix disappearance; similarly, (2) address blocks allocated several years ago can still contribute new routing prefixes.

we further examine, on average, how many new prefix advertisements (or old prefix disappearances) one allocated ad-

dress block is associated with. The result, shown in Figure 9, is directly obtained by using the similar method for generating Figure 8. Figure 9 shows that, on average, an allocated block generates 5-6 new routing prefixes, and this number does not change much over the age of the allocation. On the other side, an allocated block initially is associated with two or three old prefix disappearances, over 2 years this number grows to 4-6. The curve is slightly more volatile than that of the new advertisements, but again not much affected by the age of the allocation in long term.

Based on the above measurements, we conclude that even recently allocated address blocks are associated with both new prefix advertisements and old prefix disappearances. However, the numbers of the associated prefix advertisements and disappearances are different, both in short and long terms. This difference is the main cause of the BGP table size growth.

6. ANALYSIS OF COVERING AND COVERED PREFIXES

As we mentioned earlier, the global routing table contains both covering prefixes and covered prefixes. More formally, if an IP address block B_1 of size 2^n is fully contained in another address block B_2 of size 2^m , where $n < m$, we call the prefix for B_1 a covered prefix. If B_2 is not contained by any other address block which is represented by some prefix in the routing table, we call the prefix for B_2 a covering prefix. The existence of covered prefixes makes a significant impact on the global routing table size.

In this section, we first measure the number of covering and covered prefixes over time, then study them separately. For covering prefixes, we focus on their advertisement patterns relative to the corresponding address allocations. For covered prefixes, we classify them into 7 classes based on their advertisement paths relative to that of their corresponding covering prefixes. Because no single router peered with RouteViews over our entire study period [11/8/1997, 8/31/2004], the data used in this section were collected from two vantage points, VP1 and VP2, respectively. VP1 peered with RouteViews from October 1998 until December 2001, and VP2 from June 2001 until the time of this writing.

6.1 Evolution of covering and covered prefixes

Figure 10 shows the number of covering and covered prefixes as seen by VP1 and VP2. The *level-1* covered prefix in the figure is defined as a prefix not contained by any other covered ones, and a *level-2+* covered prefix as a prefix contained by at least one other *covered* prefix. Take the address block 61.180.0.0/16 as an example, it was allocated by APNIC to an ISP. On August 1, 2004, the BGP routing table contains 5 routing prefixes which are all related to this single allocated address block: 61.180.0.0/16, 61.180.0.0/17, 61.180.64.0/19, 61.180.96.0/19, and 61.180.128.0/17. Among these 5 prefixes, 61.180.0.0/16 is a covering prefix, 61.180.0.0/17 and 61.180.128.0/17 are level-1 covered prefixes, 61.180.64.0/19 and 61.180.96.0/19 are level-2+ covered prefixes.

From Figure 10 one can estimate that about 50%-65% of the BGP table entries are covering prefixes and the rest are covered prefixes. Among the latter, the majority (over

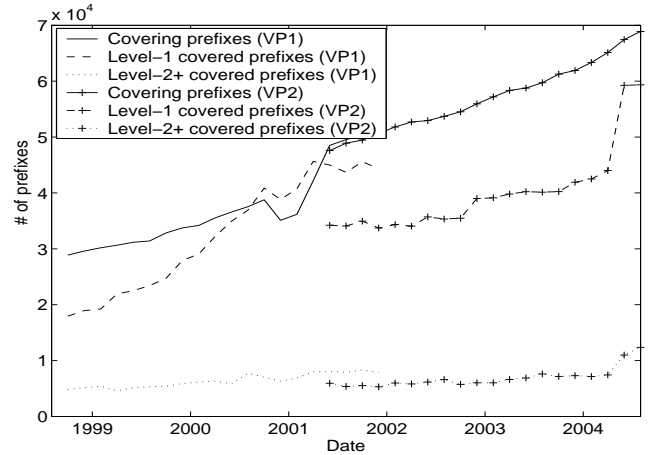


Figure 10: The number of covering and covered prefixes in BGP routing tables.

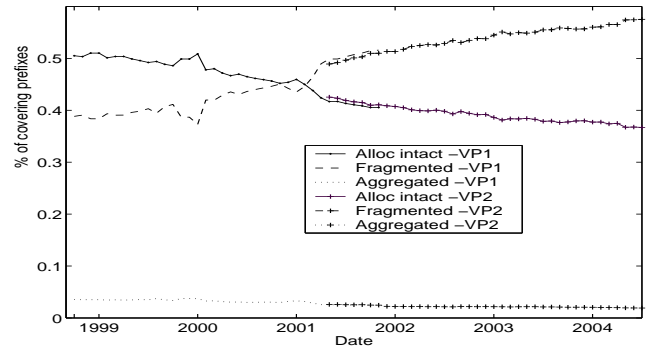


Figure 11: Advertisement pattern of covering prefixes over time.

80%) are level-1 covered prefixes, and the number of level-2+ covered prefixes has grown only slightly over time. Figure 10 also shows that VP1 and VP2 did not see exactly the same numbers of covering and covered prefixes between June 2001 and December 2001 when they both peered with RouteViews. However, the inconsistency is significant only for covered prefixes but not the covering ones.

6.2 Covering prefixes

Based on their relations to the corresponding allocated address blocks, covering prefixes can be categorized into three classes: allocation intact, aggregation over multiple allocations, or fragments from a single allocation. We plot the percentage for these three types of covering prefixes in Figure 11 by using data from VP1 and VP2. The figure shows that, in early 1999, about 50% of covering prefixes represented the allocated address blocks and 40% of them were advertised as fragments of allocations. Over time the former has been steadily decreasing and crossed the latter in early 2001, and by August 2004 about 57% of covering prefixes were fragments of allocated address blocks. Figure 12 plots the number of the covering prefixes which are fragments of allocations together with the number of the corresponding

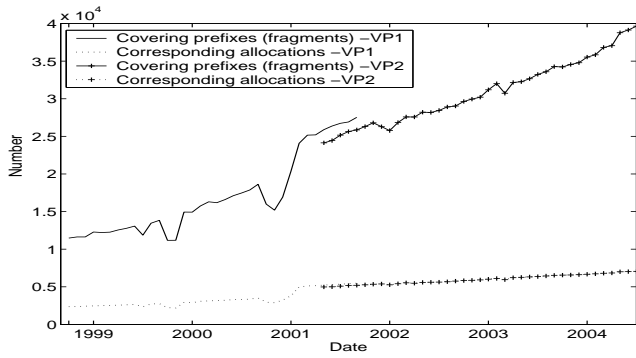




Figure 12: Comparison of the number of covering prefixes (fragmented from allocation) and the number of corresponding allocated address blocks




(a) VP1

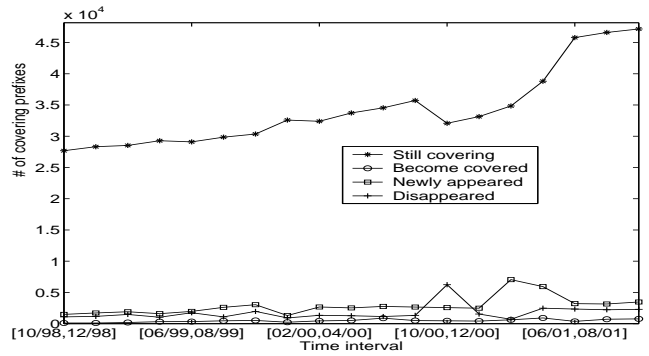


(b) VP2

Figure 13: Evolution of covered prefixes at 2-month time interval.



(a) VP1



(b) VP2

Figure 14: Evolution of covering prefixes at 2-month time interval.

allocations. The figure shows that the ratio of the routing prefixes over their corresponding allocated blocks has been steadily increasing since early 1999. In August 2004, each allocation, if fragmented, contributed 5 to 6 covering prefixes on average.

6.3 Covered prefixes

Our first observation about covered prefixes is that they show up and disappear in the routing table more frequently than the covering prefixes. To show this, we compare the routing prefixes between the beginning and end of each 2-month interval and count the following four events: (1) a covered prefix at the beginning remains unchanged at the end of the interval, (2) a covered prefix at the beginning disappears at the end, but its address space is covered by some other prefix(es), (3) a new covered prefix is advertised at the end, and (4) a covered prefix at the beginning disappears before the end and its address space is no longer covered in the routing table.

We counted the above four events as seen by VP1 and VP2 and plotted the results in Figure 13. It shows that different vantage points may see different degrees of covered prefix dynamics. For comparison purposes, we extend the above method to covering prefixes and plot the results in Figure

14. By comparing Figure 13 with Figure 14, we see that the covered prefixes appear and disappear more frequently in the routing table than the covering ones.

Covered prefixes are usually due to practical purposes such as load sharing or multihoming. To help infer the actual routing practice for covered prefixes, in Table 1 we classify covered prefixes into four classes based on their advertisement paths relative to that of their corresponding covering prefixes, with two of them further classified into sub-classes. The classification is inclusive, that is, every covered prefix falls into one of the classes. Figure 15 shows examples for all the classes and sub-classes, where a solid line with arrow represents a provider to customer AS relationship, which is inferred by using Gao's algorithm [6], a line without arrow shows the cases where the relation cannot be directly determined by the routing table, and a dashed line with arrow indicates prefix advertisement. Note that the line between two ASes represents a logical connection which correspond to more than one physical connections between the two Ases. Also note that this classification does not include level-2+ covered prefixes, which make less than 5% of the routing table entries. Due to space limit, we skip the discussion on the motivations for various prefix advertisement practices; interested readers can find more detail in [9].

According to the properties shown in Figure 15 and listed in Table 1, we classified covered prefixes in VP2's routing table on August 31, 2004. The results are given in Figure 16. This figure shows that the fourth category (DODP) is the top class of covered prefixes (44% of the total), suggesting that among various purposes for advertising covered prefixes, multihoming is the number one reason. The second largest category of covered prefixes is SODP, accounting for 28% of the covered prefixes. The possible motive for SODP is load balancing. The third largest category is SOSP, contributing 16% of covered prefixes. The other category, DOSP, takes the remaining 12%.

In a related work [8], Bu et al. performed a measurement study on the BGP table size growth between November 1997 and November 2000. Their goal was to estimate the contribution to the growth made by factors such as multihoming, failure to aggregate, load balancing and address fragmentation. Different from [8], our focus is solely on the covered prefixes rather than for all the routing prefixes, and on comparing the advertisement paths of covered prefixes with that of the corresponding covering prefixes. Therefore, our study does not directly tell which factors contribute to what percentage of the total routing prefixes, and our classification methodology is more general and tangible than that used in [8].

Besides categorizing covered prefixes, we further categorize the covering prefixes based on the types of their covered prefixes. The result is shown in Figure 16(b). Only 17% of all the covering prefixes have covered prefixes. Figure 16(b) shows that, among all the covering prefixes that contain covered ones, 78% of them have covered prefixes falling into a single category. We conjecture that in reality most ASes have a specific motivation for advertising covered prefixes.

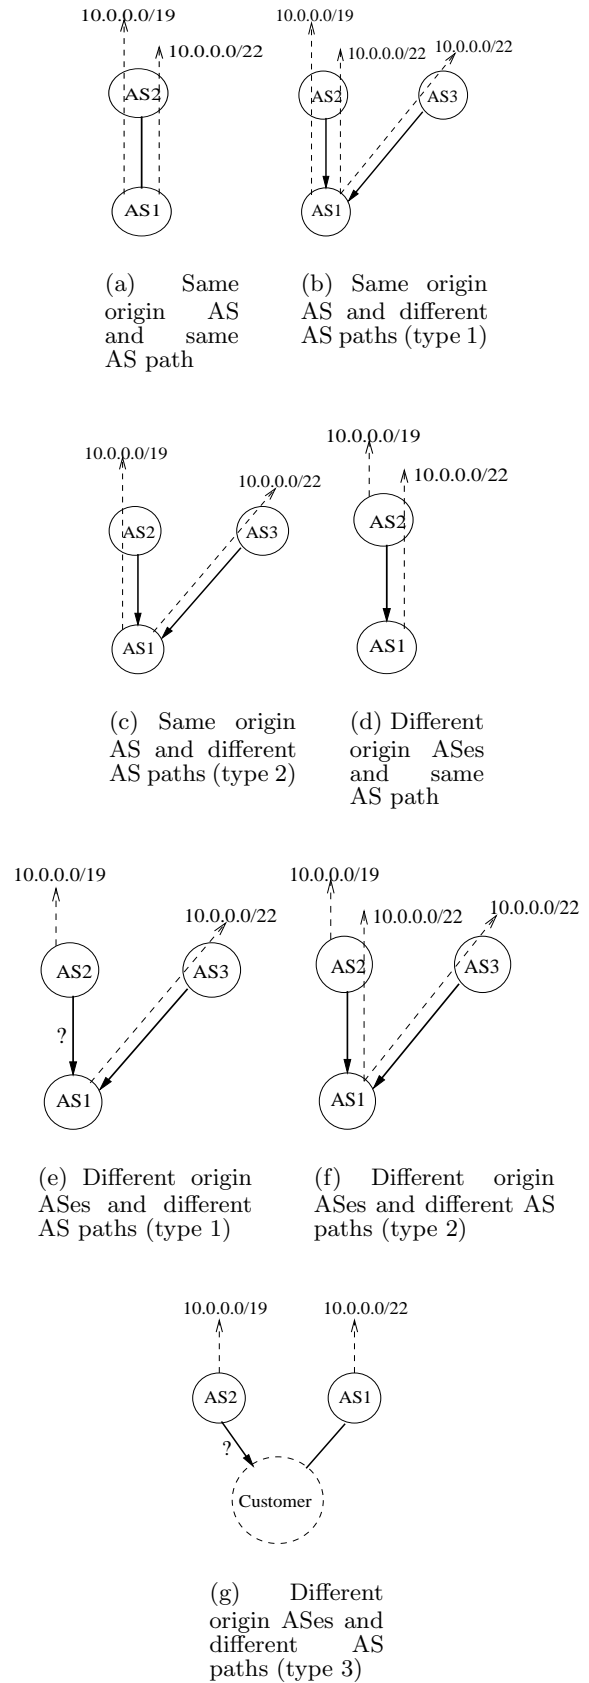
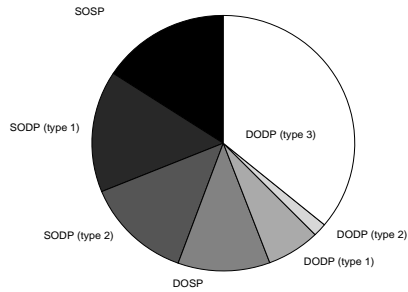



Figure 15: Examples for classifying covered prefixes.

Class	Example	Properties	Possible causes
SOSP	Figure 15(a)	(1) Covered and covering prefixes have same origin AS (2) Covered and covering prefixes have same AS path (3) Covered prefix has single origin AS and AS path	Covered and covering prefixes likely follow different physical paths
SODP	Type 1	(1) Covered and covering prefixes have same origin AS (2) Covered prefix has ≥ 2 AS paths, including the one for covering prefix	Load balancing
	Type 2	(1) Covered and covering prefixes have same origin AS (2) Covered prefix has a single but different AS path from the covering prefix	Covered prefix provides a backup for covered prefix
DOSP	Figure 15(d)	(1) Origin ASes for covering and covered have provider-subscriber relationship (2) Covered prefix has single origin AS and AS path	Fine-grained local routing policies
DODP	Type 1	(1) Covered and covering prefixes have different origin ASes and AS paths (2) Covered prefix has single origin AS and AS path	Multihoming
	Type 2	(1) Covered and covering prefixes have different origin ASes (2) Covered prefix has ≥ 2 AS paths, including the one for the covering prefix	
	Type 3	(1) Covered and covering prefixes have different origin ASes and AS paths (2) Origin AS for covered prefix advertises ≥ 2 prefixes	Multihoming for non-transit end user

Table 1: Classification of covered prefixes.



(a) Classification of covered prefixes



(b) Classification of covering prefixes having covered prefix(es) falling into different classes

Figure 16: Applying covered prefix classification to the BGP views provided by VP2 on August 31, 2004.

7. RELATED WORK

In [3][4], Huston measured the BGP routing table size from multiple aspects and enumerated several operational practices that may have contributed to the table increase. Our work differs from [3] by focusing on the impact of address allocation on the BGP routing table growth. Bu et al. [8] attributed the BGP table growth to four major factors, e.g. multihoming, failure to aggregate, load balancing and fragmentation. The difference between [8] and our work is presented in Section 6. Maennel and Feldmann proposed a realistic BGP traffic generator in [7]. A few results from their work, such as characterizing prefix forest, prefix length distribution, are similar to ours. However their focus is on BGP traffic generation while ours on analyzing the global routing table growth. H.Narayan et al. [5] proposed a model to characterize the impact of address allocation on the lookup tree structure of the routing prefixes. By subject this work may be considered closest to our work. However, [5] is focused on the structure of the routing prefixes for fast routing table lookups while our study is on the dynamics and growth of the routing table.

8. CONCLUSION

In this paper, we quantitatively characterize IPv4 address allocation and its impact on the evolution of BGP table size as well as its content. While about 63K address blocks have been allocated since the beginning of the Internet, the BGP routing table size is much larger, about 160K on August 31, 2004. During the last six years, 18K new allocations were made and the BGP table size was increased from about 55K to 160K. About 75% of the allocated addresses are covered by the routing table, and this ratio has been fairly stable over time. 90% of the new allocations are advertised in the BGP table within 75 days. There are less than 8% of the new allocations that have never been used.

We further found that the routing table growth is a dynamically evolving process with new routing prefixes being

added and existing prefixes disappeared. The routing table dynamics involves address blocks allocated over the entire history, although newly allocated blocks contribute more to the table growth than the old ones. About 70% of the BGP table size increase over the last six years are due to fragmentation of allocations made in the same time period, and 45% allocations show up in the routing table as fragments. Our data shows that the percentage of allocation fragmentation continues to increase over time.

Among all the routing prefixes, about 50%-65% are covering prefixes and the rest are covered prefixes. Covered prefixes in the routing table tend to appear and disappear more frequently than the covering prefixes. We classified covered prefixes by their routing paths as compared with that of the corresponding covering prefixes. This classification may help identify the operational practices that lead to the announcements of covered prefixes.

9. REFERENCES

- [1] IPv4 address allocation records.
<ftp://ftp.arin.net/pub/stats/{arin, ripencc, apnic, lacnic}>.
- [2] Route views project.
<http://www.antc.uoregon.edu/route-views>.
- [3] G.Huston. BGP routing table statistics.
<http://www.telstra.net/ops/bgp/>.
- [4] G.Huston. Analyzing the Internet's BGP routing table. *The Internet Protocol Journal*, 4(1), March 2001.
- [5] H.Narayan, R.Govindan, and G. Varghese. The impact of address allocation and routing on the structure and implementation of routing tables. In *Proceedings of SIGCOMM'03*, 2003.
- [6] L.Gao. On inferring autonomous system relationships in the Internet. *IEEE Global Internet*, November 2000.
- [7] O.Maennel and A.Feldmann. Realistic BGP traffic for test labs. In *Proceedings of SIGCOMM'02*. Pittsburgh, PA, 2002.
- [8] T.Bu, L.Gao, and D.Towsley. On characterizing BGP routing table growth. *Computer Networks*, 45(1):45-54, 2004.
- [9] X.Meng, Z.Xu, S.Lu, and L.Zhang. An analysis of BGP routing table evolution. *Technical Report 030046, CS Department, UCLA*, 2002.
- [10] Y.Rekhter and T.Li. A border gateway protocol 4 (BGP-4). *Internet RFC 1771*.
- [11] Z.Xu, X.Meng, C.J.Wittbrodt, L.Zhang, and S.Lu. IPv4 address allocation and the evolution of the BGP routing table. In *Proceedings of IEEE Computer Communications Workshop (CCW)*, 2003.