IPv4 Address Allocation and the Evolution of the BGP routing table

Zhiguo Xu, Xiaoqiao Meng, Songwu Lu, Lixia Zhang
Department of Computer Science
University of California, Los Angeles
Los Angeles, CA 90095
{zhiguo, xqmeng, slu, lixia}@cs.ucla.edu

Cathy J. Wittbrodt
Packet Design
3400 Hillview ave, Building 3
Palo Alto, CA 94304.
cjw@packetdesign.com

UCLA Computer Science Technical Report CSD-TR03009

Abstract

IP addresses are allocated to Internet service providers (ISPs) by four Regional Internet Registries (RIRs), in turn the ISPs further assign addresses to end users. To understand the relationship, if any exists, between the address allocation and the global routing table growth, we present a quantitative analysis of the IPv4 address allocation and growth of the global BGP routing table over the last four and half years. Our findings show that (1) the distribution of the first-advertisement-delay, which is defined as the time period between the allocation of an IP prefix and its first BGP announcement has a heavy-tail distribution, and a small percentage of the allocated address prefixes have never been used; (2) although up to 50% of the prefixes allocated between Jan.1, 1998 and April 30, 2002 are advertised in the global routing table with the same prefix length as allocated, most of the rest of the prefixes are advertised as longer (more specific) prefixes; (3) the IP prefix set in the global routing table has been evolving over time. More than half of the prefixes existed in the BGP routing table In January 1998 disappeared by January 2002, while 87,941 new prefixes were added. Among the prefixes that disappeared, 77% of them were aggregated into shorter (less specific) prefixes; (4) the impact on routing table size is highly uneven among the allocations. If we take a snapshot of the global routing table dated on April 30, 2002, more than 70% of the routing table prefixes came from 10% of the allocated address blocks.

1 Introduction

In the current Internet, IP addresses are allocated in a hierarchical manner. Four regional Internet registries (RIRs), i.e., ARIN, RIPE NCC, APNIC and LACNIC, are responsible for allocating IPv4 address blocks to Internet service providers (ISPs). These ISPs will further assign IP addresses from their allocated address blocks to the end users. The ISPs run the Border Gateway Protocol (BGP) [3], the Internet's inter-domain routing protocol, to announce all the address prefixes currently in use. Each BGP routing announcement consists of a routing prefix and an AS path: $(P, [AS_k \; AS_{k-1} \; \ldots \; AS_0])$, plus a number of other attributes. The AS path specifies the ordered list of Autonomous Systems (ASes) that data traffic to P would traverse.

The global BGP routing table has been growing rapidly in recent years; our data collection shows that the routing table size has almost doubled from early 1998 to early 2002. The continued growth of the routing table size raises concerns regarding the scalability, stability, management, and increased dynamics of BGP. To keep the BGP table growth in check, it is desirable that each ISP announces as few routing prefixes as possible. Because each allocated address block will eventually be advertised , there is potentially an intimate relationship between IP address allocation and BGP routing table growth. Although a number of earlier studies examined the BGP routing table growth and identified several factors that contributed to the growth, such as multihoming, fragmentation, load balancing, the impact of address allocation on the routing table growth has not been investigated, which is the subject of this paper.

We present the first quantitative study of IPv4 address allocation and its impact on the evolution of the global BGP routing table. We approach the problem in two ways. We first examine the recent address allocations in terms of their advertisement delay and the lengths of advertised prefixes as shown in the global routing table. By analyzing the BGP routing table and the RIR address allocation records between Jan.1, 1998 and April 30, 2002, we characterize the patterns of the prefix usage, including the first-advertisement-delay, the percentage of dead allocations and the persistence of an allocation showing up in the routing table. More specifically, we study how the advertisement patterns of allocations evolve over time, i.e., how much the allocations are getting fragmented, or aggregated, or remain unchanged in the routing table.

Second, we focus on the recent changs in the routing table and examine the changes from an address allocation perspective. We found that the total number of routing prefixes that have appeared or disappeared from the global routing table over the study period is much larger that the overall increase of routing table entries. The size change of BGP routing table is the combined result of advertisements of a large number of new routing prefixes and disappearance of old prefixes of the same order of magnitude. Furthermore, we find that about half of the BGP routing table growth over the last four years came from newly allocated address blocks, and a small number of allocations contributed to a large portion of the routing prefixes.

The rest of the paper is organized as follows. Section 2 provides some basic background information about IP address allocation. Section 3 describes the data source and methodology used in this study. Section 4 presents a brief description of the change in address allocation policies, and Section 5 examines the patterns of routing advertisements of newly allocated address prefixes. In Section 6 we describe changes in the global routing table, then we explore the relationship between these changes and the IPv4 address allocation. Related work is discussed in Section 7 followed by the conclusion in Section 8.

2 Background

The allocation and management of current IPv4 address space are mainly done by the four RIRs. The specific allocation and management policies have been evolving over time to better serve the fast-growing Internet[1]. The IP address allocation is done in prefix-based blocks. The original IP specification divided the address space into three different address classes – Class A, Class B, and Class C. Each class had a fixed boundary between the network-prefix and the host-number at a different point within the 32-bit address. However, because there were a very limited number of Class A's, and because Class C blocks were too small for most organizations who applied for network addresses, Class B address blocks were allocated most of the time. Foreseeing the impending Class B block exhaustion, the Internet Registry started allocating blocks of Class Cs to individual service providers. As a result, the global BGP routing table grew rapidly in a short period.

Classless Inter-Domain Routing (CIDR), which started being deployed during the 1993-1994 timeframe, kept the notion of identifying networks by address prefixes, but allowed a flexible boundary between the network-prefix and the host-number field. For example, CIDR permits the allocation of 64.4.176.0/20, where 64.4.176.0 is the network prefix and 20 is the prefix length. This allocated prefix contains 2^{12} unique IP addresses. CIDR provides more flexible address allocation and enables more efficient utilization of the address space.

CIDR also supports and encourages route aggregation, which can be used to limit the growth of the BGP routing table. For example, if an ISP is allocated a prefix 64.4.160.0/19, it can split it into 64.4.160.0/20, 64.4.176.0/20 and assign them to two customers separately. The ISP can advertise 64.4.160.0/19 instead of the individual /20s. Such route aggregation can effectively decrease the number of route announcements to the global routing table. A recent study [12] shows that the growth of BGP routing table size has slowed down since the deployment of CIDR.

Here we focus on the relationship between the unicast address allocation and the routing table growth. Thus our analysis does not include the IP address ranges reserved for multicast (224.0.0.0/4) and anycast. We also ignore the following address ranges that are reserved for private networks [4]:

10.0.0.0 - 10.255.255.255 (prefix 10.0.0.0/8)
172.16.0.0 - 172.31.255.255 (prefix 172.16.0.0/12)
192.168.0.0 - 192.168.255.255 (prefix 192.168.0.0/16)

3 Data sets

The data sets used in this paper include both up-to-date address allocation records and the BGP routing table records for the past four and half years.

We obtained IPv4 address allocation records (dated up to June 30, 2002) from three Regional Internet Registries (RIRs)[7]: ARIN, RIPE, and APNIC[2]. These records represent a complete set of address allocations over the Internet before June 2002. We also used the WHOIS [8] to further verify the completeness of the data. WHOIS provides a mechanism for finding IP address registration information at the RIRs. 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 located in US on Nov. 15, 1998. Such a block of addresses starts at 24.220.0.0, and the block size is 65,536. Therefore, the allocated address block spans between 24.220.0.0 and 24.220.255.255, and can be denoted by a CIDR prefix 24.220.0.0/16.

[1]RFC2050 [6] sets forth some allocation guidelines. However the allocation policies in practice have not followed them strictly.

[2]The fourth RIR, LACNIC, was approved to function after Nov. 18, 2002. In our study period which is earlier than Nov.18, 2002, all the allocations done by LACNIC were recorded in ARIN.

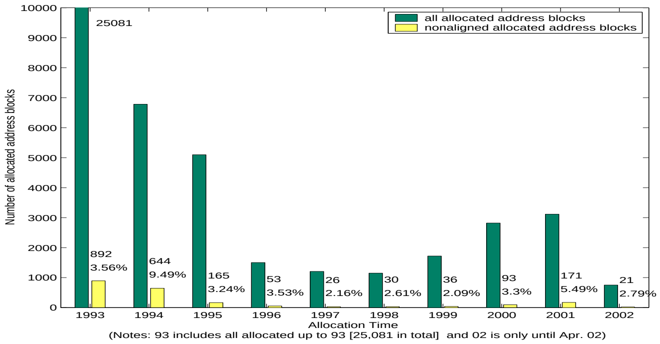

Figure 1. Allocated address blocks over time

We obtained the BGP routing tables from two different sources. One data set is from the Oregon Route-Views project [9], which collects from a number of Internet backbone routers and archives daily routing tables from November 1997 to present. The other data set is obtained from RIPE NCC, Amsterdam [10], which collects the BGP routing logs from a different set of routers from September 1999 to present. We merge the table entries from both sets on a daily basis. Since no magic vantage point exists on the Internet that can capture all the advertised routing prefixes, we hope to construct a global BGP routing table that is as complete as possible by merging the routing records from both sites. As a result, we obtained daily routing tables between November 1997 and Sept. 1999 from Oregon Route-Views, and from Sept. 1999 to present using both sets. Given that both Oregon Route-Views and RIPE RRC peer with a number of major ISPs which provide to the monitoring points a full view of their global routing table, these daily snapshots should provide a BGP table quite close to the complete set of all advertised routing prefixes. To further check the completeness of our accumulated prefixes, we also collected a third set of routing tables from the LINX monitoring site in London [11], whose data was dated from July 2000 to the present. Comparing the routing table records from the LINX with what we collected from Router-Views and RIPE NCC, the difference is less than 1% out of more than 100K routing entries. Therefore, we believe that our routing table record provides a good approximation to the complete set of globally advertised prefixes.

4 Migration of address allocation policy

This section studies how IPv4 address allocation is performed in recent years. We are mainly interested in the following two questions: (1) what is the distribution of address prefix lengths over time as the address allocation policy changed; and (2) what is the percentage of allocated address blocks in each year.

It is expected that IP addresses are allocated in the block sizes of 2^N, where N may change as the allocation policy changes. For example, with classful addresses, a Class A block contains 2^{24} addresses, a Class B block contains 2^{16} addresses, and a Class C block contains 2^8 addresses. After CIDR deployment, although N can take any value between 0 and 32 in theory, RIRs recommended a minimum allocation size of /19 before year 2000, and changed it to /20 after 2000. When addresses are allocated in the block sizes of 2^N, each allocation can be represented by a single CIDR prefix and appear as a single entry in the routing table.

To our surprise, we found that although the majority of the allocations from 1993 to 2002 fit into a single prefix, a number of allocations do not. We plotted the allocations from 1993 to April 2002 in Figure1. The green (dark if printed in black/white) bar of each year represents the number of blocks that fit into a single CIDR prefix, and the yellow (light) bar of each year represents the number of allocations that do not fit into a single prefix. We call such allocations *nonaligned* blocks. The two numbers above each yellow (light) bar are the number of *nonaligned* blocks allocated each year and its percentage out of the yearly allocation, respectively. blocks in the indicated year. A *nonaligned* block must be represented by multiple CIDR prefixes. The following allocation makes a typical example of nonaligned blocks: "`arin|TW|ipv4|192.72.3.0|64000|19900627|allocated`" This block of addresses are represented by 6 CIDR prefixes: /17, /18, /19, /20, /21 and /23. We are currently working with the RIRs to further investigate the causes of the unaligned allocations.

We plotted the prefix length distributions of the allocated address blocks during the period of 1993-2000 in Figures 2–5. Three observations can be made from these figures.

The first address allocation policy change happened when CIDR was introduced in the time frame of 1993-1994. As one can see from Figure 2, classful allocations with prefixes /16 and /24 clearly dominated the allocations of 1993. The prefix length distribution of allocations did not change much until year 1996 when for the first time, the allocated address blocks[3] with prefix length /19 became the largest component in the total allocation. The percentage of /19s among all the allocations increased from 5.57% in 1995 to 25.24% in 1996. Accordingly, the average length of allocated blocks dropped from 22.38 to 19.27.

The second policy change happened when the RIRs changed the minimum allocation size from /19 before year 2000 to /20 after 2000. The number of allocated /20s exceeded that of /19s for the first time in 2000, and this trend continued until this year. Accordingly, the average length of allocated blocks increased slightly, about 0.2 from 1999

[3] When an allocated block is in a CIDR prefix form, we use block and prefix interchangeably.

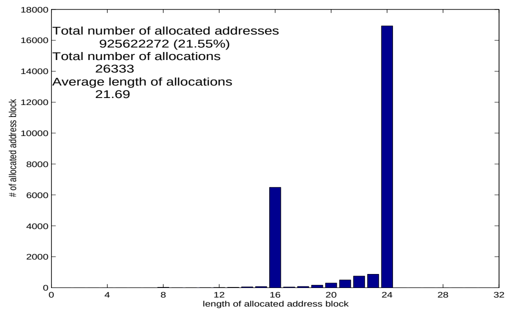

Figure 2. Distribution of address prefix length (allocated before 1993)

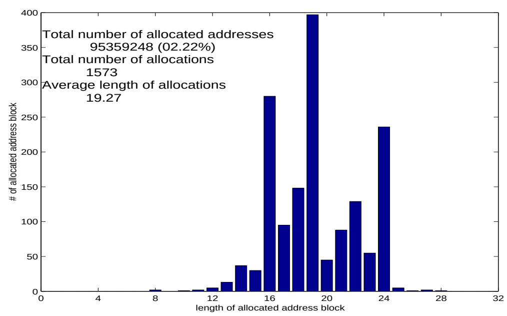

Figure 3. Distribution of address prefix length (allocated in 1996)

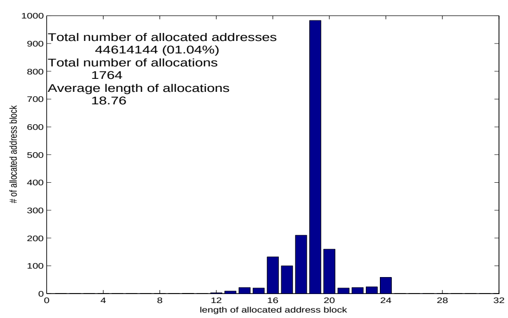

Figure 4. Distribution of address prefix length (allocated in 1999)

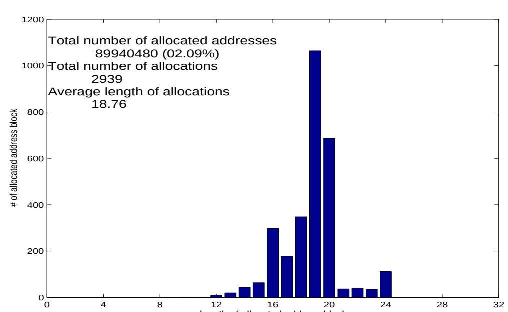

Figure 5. Distribution of address prefix length (allocated in 2000)

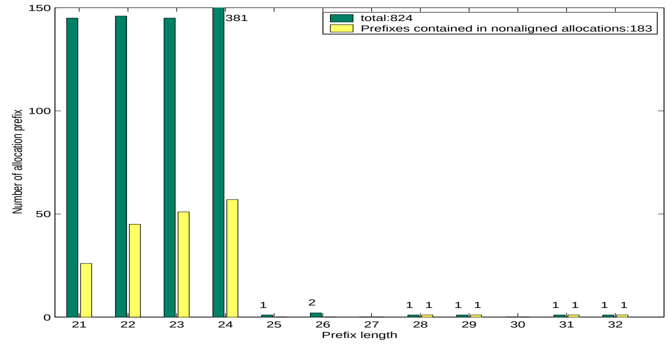

Figure 6. allocated blocks of small size (allocated from 01/01/1998 to 04/30/2002)

to 2001.

Another observation is that address blocks smaller than the recommended minimum allocation size continue to be allocated. As shown in Figure 6, 824 prefixes with length longer than /21 were allocated between 01/01/98 and 04/30/02; we even found 7 allocated blocks of length longer than 24. Out of these 824 prefixes, 623 came from aligned allocations, and the remaining 201 are the breakdown prefixes from nonaligned allocated blocks, as shown by the yellow (or light-colored) bar in Figure 6. One such example is: "`arin|US|ipv4|209.243.9.128|27|19980717|assigned`" According to the ARIN WHOIS database, this allocated block is broken into 4 prefixes: 209.243.9.128/28, 209.243.9.144/29, 209.243.9.152/31 and 209.243.9.154/32.

There were 50887 allocation records in total up to April 2002. Of all the allocations, 829 have no specified allocation dates, and these undated allocations tended to be of large sizes, e.g., of /8's or /16's. Our conjecture is that most of these undated allocations were made before ARIN existed. Between 1994 and 2001, the amount of address space allocated each year are 1.43% (1994), 1.33% (1995), 2.22% (1996), 0.77% (1997), 1.41% (1998), 1.04% (1999), 2.09% (2000), and 3.82% (2001). By April 2002, the total allocated unicast addresses are 1,834,417,891, accounting for 43% of the entire IPv4 address space.

5 Characterizing allocated address blocks

When IP addresses covered by an allocated block are assigned by an ISP to its customers, the ISP needs to advertise necessary routing prefix(es) in the global routing table. Once the prefix appears in the routing table, we say that the allocated block is advertised (or announced). In this section, we address the following issues:

- How long does it take for an allocated block to be advertised in the BGP routing table?

- How persistently has an allocated block been showing up in the routing table since its first advertisement?
- How many allocated blocks have been discarded since their initial advertisement?
- How allocated blocks are advertised: identical, fragmented or aggregated? How do these patterns change over time?

Answers to the above questions help people understand how address prefixes evolve after they are allocated by the RIRs. To tackle these questions, we analyze the 9,554 prefixes that were allocated between 01/01/1998 and 04/30/2002. The reason for choosing this period is that the earliest available BGP routing table was from the end of 1997. Therefore, the chosen study period [01/01/1998,04/30/2002] enables us to characterize the allocated prefixes in terms of their manifestation in the routing table, from the BGP routing table perspective, since the very beginning. During the study period, we also find out that 351 address blocks are not aligned with octet boundary. Since these constitute only 4% of the total 9,554 allocated blocks, we ignore them in the following study.

5.1 Latency for an allocated address block to be advertised in the routing table

We define the *first-advertisement-delay* as the interval between the recorded allocation time and the first time the prefix is advertised in the routing table. Figure 7 plots the histogram of this first-advertisement-delay. The figure shows that the first-advertisement-delay for 82% of the total 9,554 allocated blocks is less than 50 days. We also find out that the first-advertisement-delay follows a heavy-tail distribution, i.e., the number of allocated blocks with first-advertisement-delay larger than x days is well approximated by $ax^{-1.59}$ (where a is a constant). Such a heavy-tail property shows that there always exist some allocated blocks with remarkably large first-advertisement-delay no matter what timescales we use to measure the first-advertisement-delay. Since in our study period the majority of allocated blocks are of prefix /16, /19 and /20 (see Figure 4 and Figure 5), we also compute the first-advertisement-delay for allocated /16s, /19s and /20s separately. The result, given in Table 1, shows that longer prefixes tend to have shorter first-advertisement-delay.

Interestingly, we also discovered that in each month of our study period [01/01/1998, 04/30/2002], there is always a small percentage of allocated prefixes that had never been advertised, at least up to 04/30/2002. These unadvertised prefixes are shown as the light-color bars in Figure 8. From the figure we can see that some of them have not been advertised after more than 4 years since their allocation; we plan to work with ARIN to investigate these cases.

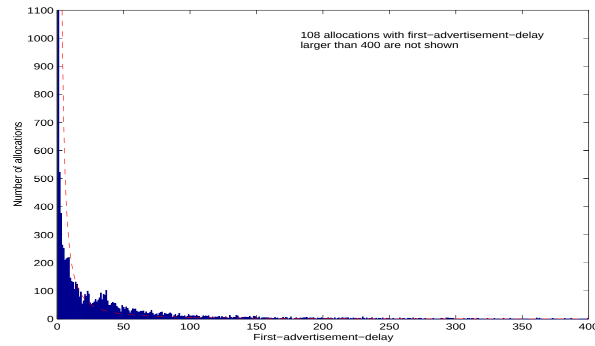

Figure 7. Histogram of the first-advertisement-delay for all the allocated blocks (dashed curve is the approximating heavy tail distribution)

		/16	/19	/20	All
Number of allocated blocks		828	3242	2293	8644
first-advertisement-delay (day)		32	48	50	44
Persistence	Daily	93%	91%	89%	91%
	Monthly	97%	98%	100%	98%

Table 1. Measurement results for prefixes allocated in period [01/01/1998, 04/30/2002] (not including non-aligned and have-not-been-advertised allocated blocks)

5.2 Showup persistence of allocated blocks in the routing table after their first advertisement

We now study how persistently an allocated block shows up in the routing table after it was advertised in the routing table. The persistence of an allocated block is measured by the ratio of the time the prefix shows up in the routing table and the elapsed time since its first advertisement. We use two timescales to carry this measurement. First, we use a day as the time unit and employ the daily routing tables to compute the persistence. If a prefix can be observed in routing table snapshots at any times in that day, the prefix is counted as shownup in that day. Second, we merge each month's routing tables into one single table, and measure the elapsed time in months instead of in days. The results are summarized in Table 1. In most of the time since their first advertisement, 91% on a daily basis and 98% on a monthly basis, the allocated blocks showed up in the routing table. The showup persistence on the monthly basis is quite close to 100% and higher than that on the daily basis. This indicates that, though on average the allocated blocks are inactive during 9% of the days, most of these inactive periods last less than a month.

Time period	Total allocated prefixes	Percentage of discarded prefixes
01/01/1998-06/30/1998	476	27%
07/01/1998-12/31/1998	598	25%
01/01/1999-06/30/1999	757	18%
07/01/1999-12/31/1999	893	18%
01/01/2000-06/30/2000	1219	18%
07/01/2000-12/31/2000	1409	24%
01/01/2001-06/30/2001	1404	15%
07/01/2001-12/31/2001	1317	11%

Table 2. Percentage of discarded prefixes

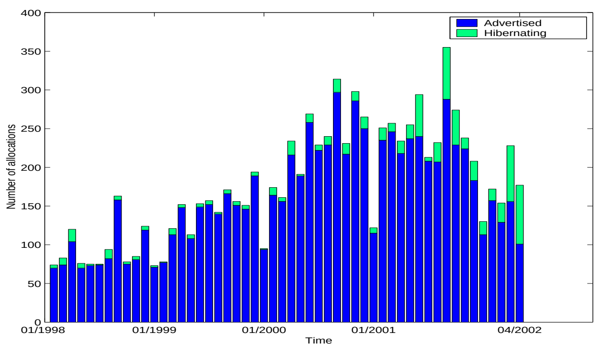

Figure 8. Advertised and have-not-been-advertised allocations (until 04/30/2002)

5.3 Number of discarded prefixes

After an allocated block is advertised initially, it can be withdrawn later and do not show up in the routing tables for a long time. We now study how often this phenomenon happens in the global routing system. Here a discarded prefix is defined as an allocated block that did not show up in the three-month period [01/30/2002,04/30/2002] but it was at least advertised once before 01/30/2002. The three-month threshold is chosen because in the RIR community such a period is generally accepted to allow a network to migrate from an old address space to a new one (usually called *renumbering*). Many reasons may contribute to the existence of discarded prefixes such as bankruptcy of ISPs, or that the ISP may obtain a new, larger address block and discard the old, smaller one.

In our analysis, we first divide the four-year study period into eight shorter intervals each of which lasts half a year, we then compute the percentage of discarded prefixes over each interval. The results are listed in Table 2. The table shows that the percentage of discarded prefixes is quite high, ranging from 11% to 27%. If we look at the entire 4-year study period, 19% of allocated prefixes fall into the discarded category. A general intuition is that the percentage of the discarded prefixes should decrease over time. However, we do find an exception that in the period of [07/01/2000,12/31/2000] such a percentage is much higher than in other periods. A closer examination reveals that this period is the first time when the number of allocated /20s exceeded prefixes with other lengths. Among the 519 /20s allocated in this period, 27% are discarded. So we speculate that this phenomenon is mainly related to the change of the minimum allocation size from /19 to /20 in the year of 2000.

In average, over the 4-year study period there are 19% allocated blocks discarded. Considering the scarcity of IPv4 address space, this percentage is quite high. Therefore, some reclaiming policies seem to be a good option.

5.4 Advertisement patterns for allocated blocks

An allocated block will be advertised in the routing table as one or multiple routing prefixes. Based on the pattern of this (these) routing prefix(es), we classify the advertisement of an allocated block into the following seven modes.

- *Identical* The routing prefix is identical to the allocated address prefix. There is no other routing prefix overlapping with the allocated block.
- *Fragmented* The routing prefix(es) are fragments of the allocated address prefix. Except such fragmented routing prefix(es), no other routing prefix in the BGP table overlaps with the allocated block. Consider the example of address prefix of 64.4.160.0/19. Two routing prefixes, 64.4.176.0/20 and 64.4.176.0/20 may be found in the routing table.
- *Aggregated* The routing prefix is the aggregate of allocated blocks. Except for this(these) aggregated routing prefix(es), no other routing prefix overlaps with the allocated block. In the example of allocated block 64.4.176.0/20, we may find prefix 64.4.160.0/19 in the routing table.
- *Identical+Fragmented* Identical and fragmented routing prefixes co-exist.
- *Identical+Aggregated* Identical and aggregated routing prefixes co-exist.
- *Fragmented+Aggregated* Fragmented and aggregated routing prefix co-exist.
- *Identical+Fragmented+Aggregated* Identical, fragmented and aggregated routing prefixes co-exist.

The above seven modes are mutually exclusive and they enumerate all the possibilities that an allocated block can be advertised.

In the following, we study which of the above modes an allocated block is advertised and how these modes evolve over time. We first define the *age* of an allocated block as the elapsed time since its first BGP advertisement; the age is measured in months. Since the concept of age makes sense only for advertised prefixes, we exclude prefixes that have never been advertised (from 01/01/1998 to 04/30/2002) Section 5.2 shows that over 91% of allocated blocks persisted in the routing table. For any allocated block, its advertisement mode should take one of the aforementioned seven possibilities, and this mode may vary over its age. By averaging over the total allocated blocks, we calculate the percentage of each of the seven advertisement modes. We first do this calculation on all the allocated blocks without considering their prefix length. The result is shown in Figure 9. To make the averaging meaningful, we ensure that the number of allocated blocks involved in the averaging is roughly equal to the number of allocated prefixes in a six-month period.

Figure 9 shows that the percentage of identically advertised allocated blocks is the highest out of the seven possible modes, contributing between 48% and 55%. The second largest percentage happens to the modes *fragmentation* and *identical+fragmentation*, each of which roughly constitutes 20%. We can also see that the percentage of purely aggregation mode roughly remains around 8%.

We next study the modes of the allocated /16s, /19s and /20s, individually. For each prefix length, we ensure that samples used in the averaging process are at least more than prefixes allocated in six months. The results are shown in Figures 10, 11 and 12. Comparing the statistics, the allocated /16s have a higher fragmentation percentage and a lower aggregation percentage. This observation is intuitive because in recent years /16s are usually allocated to large ISPs and organizations. These big ISPs (or organizations) are more likely to perform fragmentation, load balancing or other traffic engineering operations. All these operations will lead to increased fragmentation of address blocks.

Another observation is made on the allocated /20s. Seen from Figure 12, the aggregation percentage of /20s is around 22%, which is much larger than other allocated blocks. Further investigation shows that during the study period, 522 allocated /20s were aggregated. Of these, 193 /20s are aggregated into 62.0.0.0/8 and 66.0.0.0/8, and 190 /20s are aggregated into /19s. Such a comparatively high aggregation percentage of /20s is possibly induced by the policy regarding the default allocation size. Before the default allocation size was changed from /19 to /20 in 2000, /19 was recommended as the shortest globally-routable prefix length. If an announced routing prefix is longer than /19 and does not belong to the address space of traditional Class C network, it might be filtered by some core BGP routers. After the change of the default allocation size in 2000, a large number of /20s were allocated. In accordance, many of the BGP core routers changed their filtering policies to accept /20 as the shortest routable prefix length. However, some ISPs are still sticking to the old strategy of aggregating their /20s. This may explain why the allocated /20s are more likely to be aggregated.

5.4.1 Fragmentation

As we mentioned in Section 5.4, fragmentation of allocated blocks can appear in four modes, i.e., *Fragmented*, *Identical+Fragmented*, *Fragmented+Aggregated* and *Identical+Fragmented+Aggregated*. In total, these four modes apply to about 45% of the allocated blocks. In another words, around 45% of all the allocated blocks are fragmented. Since fragmentation will definitely increase the BGP routing table size, we take a closer examination to see to what extent the allocated blocks are fragmented. To this end, we calculate the total number of routing prefixes that are fragmented from allocated blocks and then divide it by the number of these allocated blocks. We use this value to represent the average number of routing prefixes that are fragmented from a single allocated block. Since this value varies with the age of the allocated blocks, we plot its changes in Figure 13. As seen in the figure, generally the older an allocated block is, the more fragments the block has. We also notice that the values at the age of 28, 29 (months) are extremely high. A further investigation shows that these are caused by a misconfiguration occured between 12/09/2000 and 01/21/2001. During that period, more than 10,000 /24s belonging to the address block 63.0.0.0/10 were announced by AS706, which is due to the (hidden) test-cef command.

We now conduct a coarse calculation to see how many routing prefixes are fragmented from allocated blocks. As we know, 8,644 blocks are allocated and advertised during our study period. Of these, 45% are fragmented (seen in Figure 9). As we observe from Figure 13, the average number of fragmented routing prefixes is about 14. Therefore, the 8,644 allocated blocks will generate routing prefixes as many as $8644 * 45\% * 14 = 54475$, which is roughly half of the routing table size in April 2002. Considering that these 8,644 only constitute 17% of the total 50,436 allocated blocks[4], we conclude that the majority of BGP routing table prefixes are fragmented from allocated blocks.

5.4.2 Aggregation

In Figures 9–12, we can see that about 12% allocated prefixes have been aggregated, and such aggregations appear in modes of *Aggregated*, *Identical+Aggregated*, *Fragmented+Aggregated* and *Identi-*

[4] We do not count non-aligned allocated blocks up to 04/30/2002.

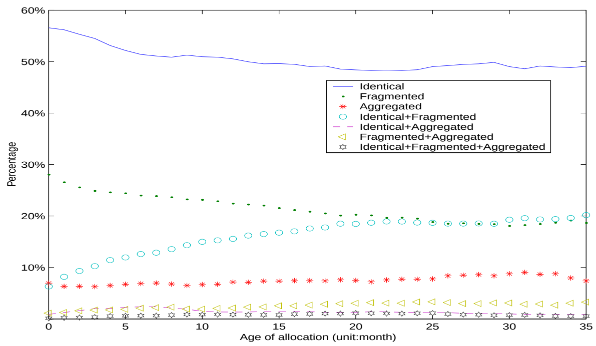

Figure 9. Distribution of advertisement modes for all allocated blocks changes over the prefix age

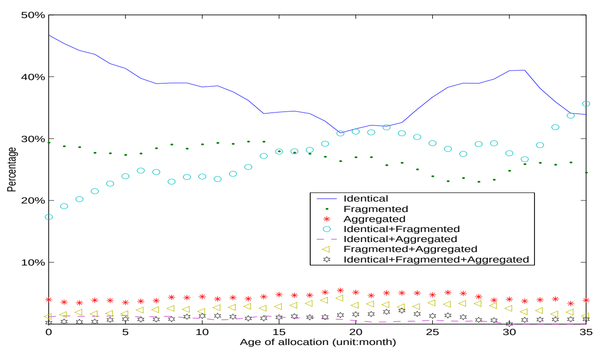

Figure 10. Distribution of advertisement modes for all allocated /16s changes over the prefix age

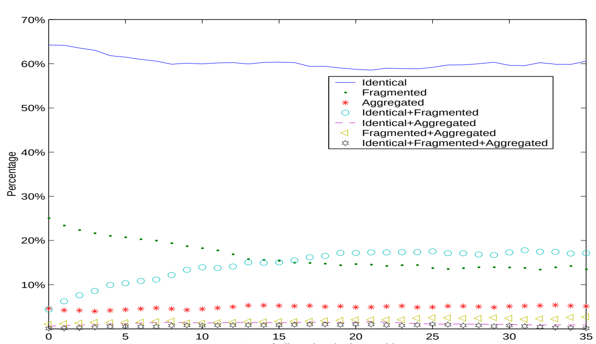

Figure 11. Distribution of advertisement modes for all allocated /19s changes over the prefix age

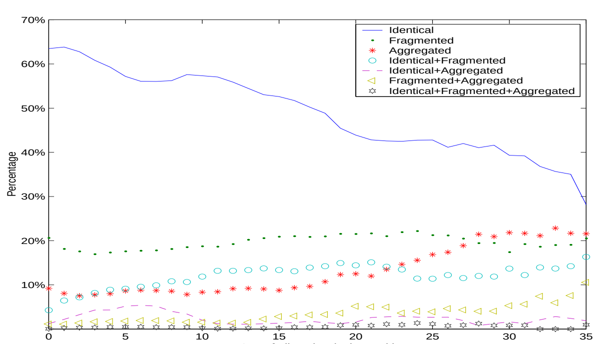

Figure 12. Distribution of advertisement modes for all allocated /20s changes over the prefix age

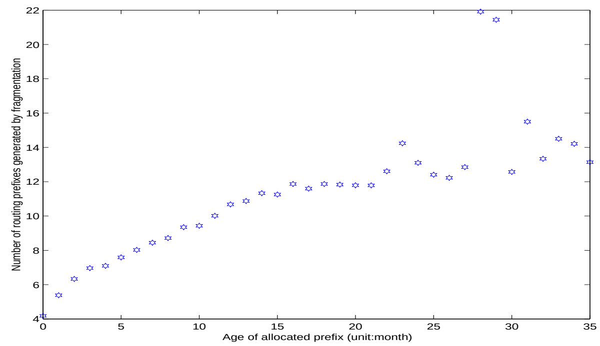

Figure 13. Number of routing prefixes fragmented from an allocated block (in average)

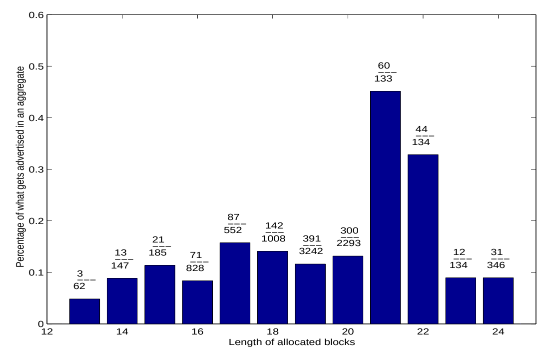

Figure 14. Distribution of prefix length for aggregated allocations(from 01/01/98 to 04/30/02)

cal+Fragmented+Aggregated. Since aggregation is a desirable activity and strongly recommended by the RIRs, we take a closer look and see how the prefix length of these aggregated allocated prefixes is distributed and how aggressively these allocated prefixes are aggregated.

Based on the routing table snapshot in April 2002, Figure 14 depicts the distribution of the prefix length for the aggregated allocations made from 01/01/98 to 04/30/02. Each bar represents a prefix length. The denominator near each bar represents the total number of allocated prefixes with the corresponding prefix length (excluding non-aligned and have-not-been-advertised allocated prefixes), and the numerator denotes the number of aggregated allocations. The figure shows that more /19s and /20s are aggregated compared with other prefixes. However, the percentage of aggregated /21s and /22s is the highest.

We then examine how aggressively the aggregation has been carried, i.e., how much shorter the routing prefix becomes compared with the original allocated prefix after aggregation. To this end, we use the metric of *De-*

gree of aggregation (DOA), which is defined as the difference between the allocated prefix length and the advertised routing prefix length. For example, an allocated prefix 64.4.160.0/20 is announced in an routing prefix 64.4.160.0/19, then its DOA is $20 - 19 = 1$.

We compute the percentage of DOA values for four prefixes of /19, /20, /21 and /22, since /19s, /20s have the largest number of aggregated allocations and /21s, /22s have the highest aggregation percentage. The results are given in Table 3. An interesting observation is that the *preferred* DOA for /20s is 1 in the sense that it happens most frequently. And the preferred DOA for /21s and /22s is 2 and 3, respectively. This indicates that the allocated /20s, /21s and /22s tend to be aggregated into /19s in the routing table. This still can be explained by the policy of minimum allocation size. As mentioned in Section 5.4, /19 has been recommended as the shortest globally-routable prefix length before 2000. To prevent their announced prefixes from being filtered out by BGP core routers, many ISPs tend to aggregate their long prefixes into /19. This operation still exists in practice even after 2000.

DOA	/19	/20	/21	/22
1	42%	54%	20%	23%
2	28%	13%	55%	22%
3	13%	5%	13%	41%
4	5%	3%	3%	12%
≥ 5	12%	25%	9%	2%

Table 3. Degree of aggregation (DOA) for allocated /19s, /20s, /21s and /22s

6 Impact of allocation on the global BGP routing table

Section 5 characterizes the allocated blocks in terms of their manifestation in the routing table. We now study the impact of these address allocation on the evolution of the BGP routing table. The focus here is shifted from address allocation to the evolution of the BGP routing table and the manner how recent allocation affect this evolution.

In the following, we first study aspects of the routing table evolution, which includes the table size change, change of the IP address consumption, and change of the routing prefixes. We then quantify the impact of allocation on both the change of the routing table and the single routing table size.

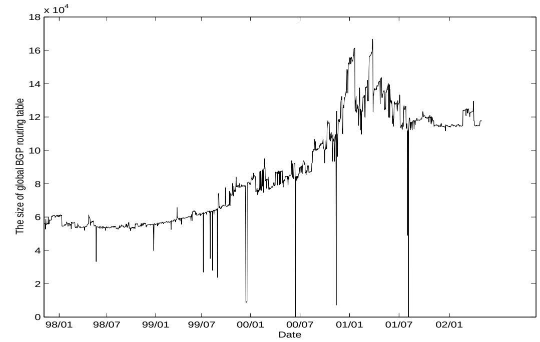

Figure 15. Routing table size

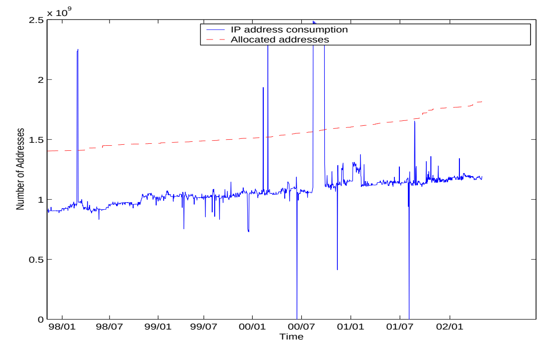

Figure 16. IP address consumption

6.1 Evolution of the global routing table

6.1.1 Change of routing table size

Figure 15 plots the change of the daily routing table size in our study period [01/01/1998, 04/30/2002] (There exists some zero routing table size. This is due to the corruption of the routing table files on a couple of days). The figure shows that the routing table size has almost doubled during the study period. It increases from 60,395 to 117,971 entries. We also notice two spikes in January 2001 and March 2001, respectively. The first spike was caused by a misconfiguration described in Section 5.4.1. The other spike has also been observed by the literature [16] and is due to similar misconfiguration problem.

6.1.2 Change of IP address consumption

IP address consumption is defined as the total number of unique IP addresses spanned by the routing prefixes. It is plotted in Figure 16. Compared with Figure 15, the address consumption grows at a much slower rate than the routing table size. The global BGP routing table roughly doubled over the last four and a half years, while the IP address consumption only increased by 32%. This discrepancy is caused by the fact that multiple routing prefixes overlap and

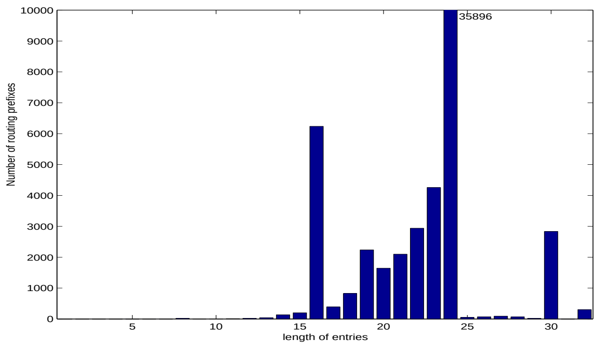

Figure 17. Distribution of routing prefix length in Jan. 1998

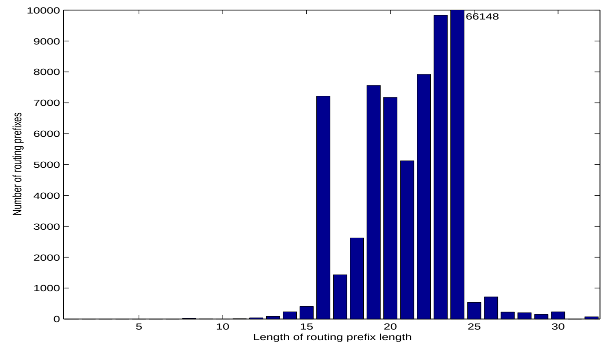

Figure 18. Distribution of routing prefix length in Apr. 2002

cover the same portion of the address space.

We then compare the address assumption to the total allocated IP addresses in the history. The address consumption is found to grow in a way similar to the total allocated addresses. Since Section 5.1 showed that most allocated blocks are advertised within 50 days and tend to persistently exist in the routing table, we speculate that the increase of address consumption is mainly contributed by the simultaneous new allocated blocks.

The figure also shows a sharp spike between [08/15/2000, 09/25/2000], which was caused by the misbehavior of AS7777 that announced 70-80 /8s during that period and the address space contained by these /8s had not been allocated yet. Figure 16 also plots the total number of allocated IP addresses. As we have mentioned, there exist some allocation records without allocation time. Since some of them are in big chunks and thus cannot be ignored, we presume that they were made before 1993 and add them to the allocated IP addresses.

6.1.3 Change of distribution of routing prefix length

Figure 17 and Figure 18 depict the distribution of routing prefix length at two different times of our study period: the starting time 01/01/1998 and the ending time 04/30/2002. During the study period, the majority of routing prefixes have a length ranging between /16 and /24. However, the increase of /16s and /24s is slower than the overall increase, while /19 and /20 increase much faster than the overall growth rate and have at least doubled or even tripled. For example, though /24 is persistently the largest slice in the global BGP routing table, the percentage has decreased slightly, from 59.4% in January 1998 to 56.1% in April 2002. Coincidently, /19 and /20 dominate the new allocation during this period. Due to the rapid growth of routing prefix with medium lengths, the average prefix length of the routing entries has decreased from 22.74 in Jan. 1998 to 22.36 in April 2002.

6.2 Impact of allocation on global BGP routing table change

We have heretofore known that the new increases in address consumption are mainly due to the concurrently allocated address blocks and we have also noticed growth in the routing table size. We now turn to the following question: Is such growth of the routing table mainly caused by the concurrently allocated address blocks? Or does it simply result from fragmentation (or aggregation) of the existing routing prefixes?

To answer this question, we need to examine the relationship between the change of the routing table and the concurrently allocated address blocks. Such a relationship cannot be figured out directly from observing the overall change of the routing table size. Instead, we need to identify those newly-advertised and withdrawn routing prefixes and study the correlation between them and the concurrent new allocations.

Surprisingly, the growth of the routing table actually results from two conflicting events: the advertisements of new prefixes and the withdrawal of existing old prefixes. In other words, when we compare routing table snapshots at two different times, a large number of new routing prefixes show up at time 2 while a large number (in the same order) of existing prefixes that showed up at time 1 no longer show up at time 2. A possible explanation to these disappeared prefixes may be that some ISPs are bankrupted or some ISPs reclaim assigned address blocks from their customers. However, our subsequent analysis shows that this is not always the case.

In the following several subsections, we first give a more detailed study in routing table change, and then branch into two parts. The first part investigates what contributes to

Prefix length	/8	/9	/10	/11	/12	/13
Newly-advertised	4	4	4	5	23	57
Disappeared	6	1	2	4	7	8
Prefix length	/14	/15	/16	/17	/18	/19
Newly-advertised	128	249	1959	1107	1942	5694
Disappeared	34	55	1095	117	279	704
Prefix length	/20	/21	/22	/23	/24	/25
Newly-advertised	5461	3816	6106	7635	50917	582
Disappeared	753	1050	1668	2635	22166	53
Prefix length	/26	/27	/28	/29	/30	/32
Newly-advertised	784	335	279	229	436	155
Disappeared	67	89	66	17	2832	304

Table 4. Distribution of newly-advertised and disappeared prefixes (the total number of newly-advertised prefixes is 87,941, and the total number of disappeared prefixes is 34,012)

Allocation Time	$\leq$ 1993	1994	1995	1996
Number	10117	7405	7290	8139
Allocation Time	1997	1998	1999	2000
Number	4109	5316	10670	14257
Allocation Time	2001	Lack of Time Info		
Number	10363	8907		

Table 5. Distribution of allocation time for newly-advertised routing prefixes

the newly-advertised routing prefixes, and the second part identifies where the disappeared routing prefixes have gone.

6.2.1 Change of routing table = newly-advertised prefix − disappeared prefix

We analyze the difference between the routing tables on 01/01/1998 and 01/31/2002. The chosen time period is three months shorter than the period [01/01/1998, 04/30/2002] used in previous Section 5. The reason is that we need to use the three-month period [02/01/2002, 04/30/2002] to make sure that the newly-advertised or the disappeared routing prefixes are not caused by transient *abnormalities* such as the aforementioned misconfigurations.

During the study period [01/01/1998, 01/31/2002], the routing table size increases from 60,395 to 114,324. However, the actual change includes not only the advertisement of 87,941 new prefixes but also the disappearance of 34,012 existing old prefixes. We further look into the routing tables in the following three-month period [02/01/2002, 04/30/2002] and discover that, among the 87,941 newly advertised prefixes, only 119 are never show up again in the three months. Among the 34,012 disappeared prefixes, only 125 show up again in the three months. This shows that the newly-advertised and the disappeared prefixes are not likely to be generated by abnormal events which typically do not last long.

As a side note, the address consumption during the study period, i.e., the number of unique IP addresses covered by the routing table prefixes, increases from 921,694,960 (21.46% of the entire IPv4 address space) to 1,163,961,392 (27.10% of the entire IPv4 address space).

The distribution of the prefix lengths for the newly-advertised and disappeared prefixes is summarized in Table 4. It shows that /24s contribute most to both the newly-advertised and the disappeared prefixes. Specifically, /24s it constitutes 50,917 (57.5%) of the newly-advertised prefixes and 22,166 (65.2%) of the withdrawn prefixes. Therefore, /24s are the most active components in the routing table. We also notice that the disappearing rate of /24s exceeds the advertisement rate, thereby resulting in a decrease of the percentage of /24s in the routing table. Another information conveyed by Table 4 is that the new advertisement of /19s and /20s is 7 to 8 times faster than their disappearance. This is reminiscent of the fact that most of the allocated blocks in the study period are /19s or /20s.

Therefore, we have observed that the total number of routing prefixes involved in the change of the routing table size is much larger than the seeming change of the table size. A large number of new prefixes are announced to the routing table while at the same time a large number of existing prefixes disappear. The following subsections explore what causes these new advertisements and disappearances.

6.2.2 Newly-advertised prefixes

As we mentioned in Section 5.4, a routing prefix can be an identical advertisement as an allocation block, or it can also be a longer prefix which is fragmented from the allocated address block or a shorter prefix by aggregating several smaller address block. Moreover, these three situations can co-exist for the same routing prefix (see Section 5.4). For simplicity, we classify a routing prefix as a fragmented one whenever it is found to be fragmented from any allocated block, and if it is found to be purely identical or purely aggregated, it will be classified as an identical or aggregated prefix, respectively. Using these three categories, we classify the 87,941 newly-advertised prefixes and find that 77,774 (88.4%) of them are in the fragmented category. The identical category has 8,808 (10.0%) prefixes while the aggregated category only accounts for 588 (0.7%) prefixes.

Allocation time for newly-advertised prefixes The first issue we are interested in is when the address space represented by these newly-advertised prefixes was allocated. The distribution of these allocation time is summarized in Table 5. It turns out that 46.9% of these newly-advertised prefixes [5] are generated by the 7,608 allocated blocks made in our study period [01/01/1998, 01/31/2002], and the other 53.1% are from address block allocated before 01/01/1998.

New IP addresses brought by the newly-advertised routing prefixes A newly-advertised prefix does not necessarily mean that the number of addresses contained by itself are added to the IP address consumption of the routing table. In fact, of the 87,941 newly-advertised prefixes, 40.6% do not contribute new addresses. In other words, these prefixes are just fragmented from existing prefixes. Accordingly, though the newly-advertised prefixes contain 486,719,936 (11.33% of the IPv4 address space) IP addresses, they only bring 374,203,476 (8.71% of the IPv4 address space) new addresses.

6.2.3 Disappeared prefixes

We categorize the 34,012 disappeared prefixes in our study period. Of them, as shown in Fig.19, 26,199 (77%) are aggregated into 2,495 shorter prefixes, 1,322 (4%) are fragmented, and the other 6,491 (19%) are simply discarded. Among the 1,322 fragmented prefixes, 643 are fragmented without losing IP addresses while the fragmentation of the other 679 causes addresses loss. In overall, the number of addresses that are taken away by the disappeared prefixes, are 131,937,044 (3.07% of the IPv4 address space). As shown in Fig.20, the entire address space represented by the "aggregated" prefixes and part of that represented by the "fragmented" prefixes is still in use but represented in different set of prefixes, shorter or longer. But the address space covered by the "discard" prefixes and part of that covered by the "fragmented" prefixes simply goes out of use. From these two figures, we find the prefixes that get aggregated dominate the disappeared prefixes but the address space covered by them is much smaller than that covered by others. The main reason is there are several /8 in other two types of prefixes.

6.3 Impact of address allocation on routing table size

Our final goal is to examine the impact of address allocations on the static routing table size. We first describe our method to measure this impact: for each routing prefix, the addresses it contains must belong to one or multiple allocated blocks P_1, P_2, …, P_n ($n \geq 1$). In this case, we say that each of the allocated blocks $P_1, P_2, \ldots, P_n$ contributes $\frac{1}{n}$ to the routing table size. We then check all the routing

[5] A single routing prefix that are aggregated from several allocated blocks may correspond to multiple allocation time. To avoid this situation and consider their small percentage (0.7%), we ignore them in the calculation.

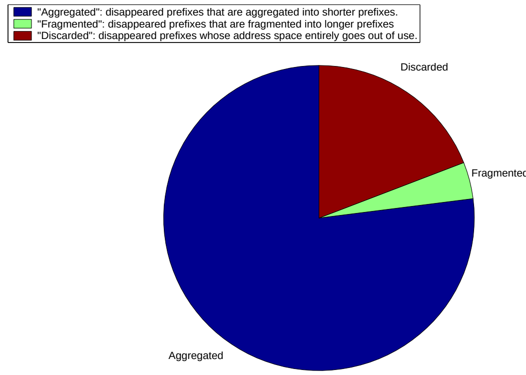

Figure 19. The distribution of the disappeared prefixes

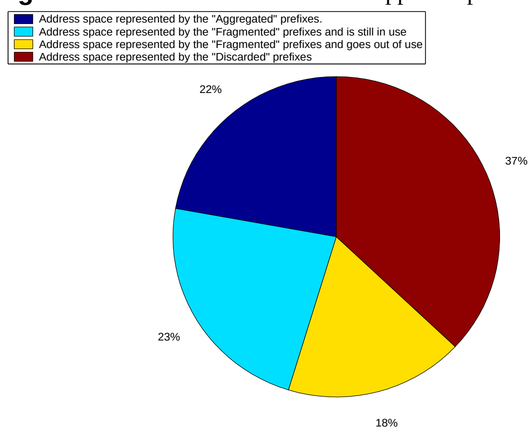

Figure 20. The distribution of the address space represented by the disappeared prefixes

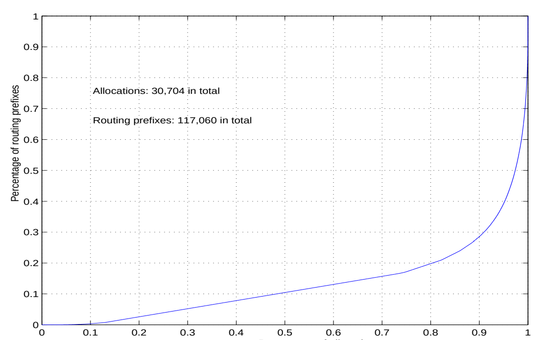

Figure 21. CDF of allocated blocks' contribution to the routing table size

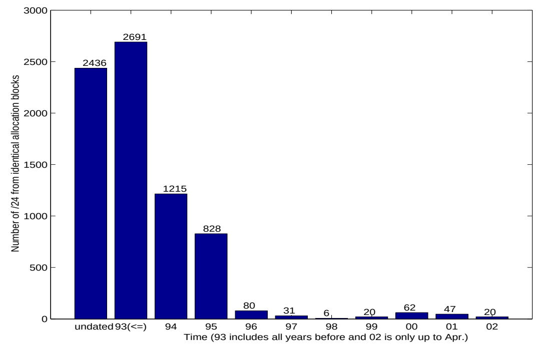

Figure 22. Number of routing prefix /24s that are identical to allocated blocks in different time periods

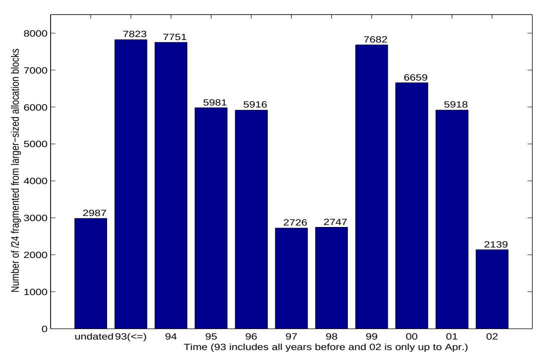

Figure 23. Number of routing prefix /24s that are fragmented from allocated blocks in different time periods

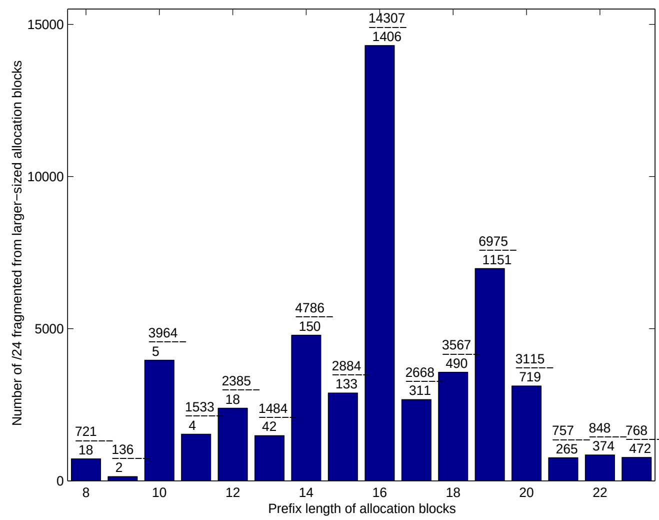

Figure 24. Distribution of routing prefix /24s in terms of the size of allocated blocks from which they are fragmented

prefixes and for each involved allocated blocks its attribution to the routing table size is accumulated. Obviously, the sum of the accumulated contribution of all the involved allocated blocks should be equal to the total number of routing prefixes.

We perform the above measurement on the routing table snapshot on April 30, 2002. The routing table has 117,060 routing prefixes which are generated from 30,705 allocated address blocks. Figure 21 plots the CDF for all the involved allocated blocks' contribution to the routing table size. It shows that 90% allocated blocks only account for less than 30% routing prefixes. In other words, 10% allocated blocks contribute to as high as 70% routing prefixes. This indicates that the impact on the routing table size is highly skewed among the allocated blocks.

6.4 Impact of address allocation on routing prefix /24s

In our study period [01/01/1998, 04/30/2002], /24s are persistently the largest component of the routing table. Though the percentage is slightly decreasing, it is still as high as 56.1% until April 30, 2002. Out of the 66,148 /24s in the routing table on April 30, 2002, 11.2% are identical to allocated blocks while 88.2% are fragmented from larger allocated blocks. Figures 22 and 23 show the distribution of routing prefixes /24 in terms of the allocation time of their address space. Specifically, Fig.22 is for /24 that are identical to the allocated blocks, and Fig. 23 is for /24 that that are fragmented from larger allocated blocks. Before 1996, most allocated address blocks are of prefix length /24. Coincidentally, as shown in Fig. 22, most routing prefixes /24 that are matched with identical allocated blocks come from allocated blocks made before 1996. Also in this figure, there is a big bar labeled with "undated", which means the corresponding allocation records have no specified date. To our conjecture, they should have also been allocated well before 1993.

For these /24, 58329 in total, that are fragmented from larger-sized allocation blocks, we further investigate the size of their matched allocated blocks. Of these, 50898 come from aligned allocated blocks with prefix lengths ranging from 8 to 23. In Fig.24, we present the distribution of these 50898 routing prefix /24s in terms of prefix length of allocated blocks from which they are fragmented. The numerator near each bar denotes the total number of /24 that are fragmented from allocation blocks with the corresponding prefix length and the denominator represents the number of the matched allocated blocks. From the figure, we notice allocation blocks with prefix length 16 contribute more /24 than allocated blocks with any other prefix lengths.

Other 7431 /24 are fragmented from 1031 non-aligned allocated blocks. Since the non-aligned allocated blocks

	Allocated Prefixes	Allocated addresses	Average number of routing entries fragmented (aggregated) from single allocated block
U.S.	3,194	214,809,088	6.8 (0.3)
Canada	293	7,461,632	2.4 (0.4)
China	165	21,010,432	6.7 (0.5)
Japan	105	25,952,512	6.2 (0.3)

Table 6. Average number of routing entries fragmented/aggregated from single allocated block

must be split in nature and the number of /24 that are fragmented from them is small, we do not characterize them in detail. An interesting property is the contribution to /24 from these non-aligned allocated blocks is highly uneven. Of these 7431 /24, 38.4% are broken from less than 1% of these non-aligned allocated blocks.

6.5 Level of fragmentation/aggregation for regions with less address space

One intuitive conjecture is that for regions with tighter allocated address space, the level of fragmentation is higher and the level of aggregation is lower compared to regions with more plentiful allocated address address. To check this conjecture, we choose U.S., Canada, China and Japan to conduct a case study. US is typically considered as a region with plentiful address space while the latter three regions have much tighter address space. For each region, we use the routing table on April 30, 2002 to calculate the total number of routing prefixes that are fragmented from the original prefixes allocated in the study period [01/01/1998, 04/30/2002]. This number is then divided by the number of original allocated blocks to get the average number of routing prefixes that are fragmented from each allocated block. The same calculation is also applied to aggregation and the results are summarized in Table 6. The table shows that the four regions we have chosen have quite different fragmentation/aggregation phenomenon no matter how many addresses they have been allocated. There is no obvious clue to support the aforementioned conjecture.

7 Related work

Two recent works [12][13] have studied the growth of the global BGP routing table. [12] quantified the growth while the focus of this paper is to understand the pontential relationship between address allocations and the growth. [13] identified four relevant factors that contribute to the rapid growth of the routing table size. These four factors are multi-homing, failure of aggregation, load balancing, and address fragmentation. Based on analysis of the routing tables, they concluded that 75% of the newly increased routing table prefixes are caused by fragmentation. Our main goal is to understand the routing table growth from the address allocation perspective. The data sets used include both routing tables and address allocation records. We also obtained several new findings as described in the previous sections.

There have been a few recent proposals to control the rapid growth of the BGP routing table size. [14] proposes a prefix-based filtering policy to remove routes that have mask lengths longer than their assigned prefixes. The paper quantitatively shows that the most aggressive filtering policy only affects the reachability of about 0.3% of the total address space. [5] recommends to use techniques such as address lending to improve the scalability of BGP routing table. Our focus is to better understand how the current allocation policy affects the routing table growth; we do not indend to propose a new address allocation or route management policy.

As this paper being developed, we learned that other researchers [17] have also been collecting the statistics of allocated address blocks. Our current study can benefit from such data collection and analysis. However, our focus is not limited to the characteristics of allocated IPv4 address blocks, but instead our goal is to understand the relationship between address allocations and the global routing table growth.

8 Conclusions

The global BGP routing table has almost doubled since early 1998. The rapid growth of the routing table has stimulated several recent studies on what factors contribute to such growth. While earlier research efforts have examined factors such as multihoming, fragmentation and load balancing, the relationship between the address allocation and the evolution of the routing table has not been studied. This paper presents the first quantitative study of the IPv4 address allocation and its impact on the global BGP routing table based on the allocation records and the routing table data in the last four and a half years. Our findings can be summarized as follows.

We first characterize the recent address allocations in terms of their advertisement patterns in the global routing table. The typical life phases of an allocated address prefix include: after certain delay it is first advertised and appeared in the BGP routing table in the form of a routing prefix. The allocated address prefix may appear in the routing table in several patterns, including unchanged, fragmented, aggregated, or some combination of these three. We discovered that the first-advertisement-delay takes a heavy-tail distribution: a small percentage of the allocated prefixes exhibits exceedingly long first-advertisement-delays. After their first advertisement, a majority of the allocated prefixes will remain in the routing table, but 19% of them stopped being advertised. We also find out that about 48%-55% of the recent allocations are identically advertised to the routing table. However, about 45% of the recent allocations are fragmented and these fragments contribute to about half of the routing table size of April 2002.

We also studied the changes in the global routing table from the perspective of the allocated blocks. We examined the evolution of the routing table over the past four years and discovered that the increase in the routing table size is due to the combined effect of two events: More than 50% (i.e., 34,012) of the routing prefixes disappeared, and at the same time 87,941 new prefixes were added to the routing table. Our study also shows that the impact on routing table size is highly uneven among the allocations. By taking a snapshot of the global routing table on April 30, 2002, we found that more than 70% of the routing table prefixes came from 10% of the allocated address blocks.

References

[1] Y.Rekhter and T.Li, Architecture for IP Addressp Allocation with CIDR, Internet RFC 1518

[2] V.Fuller, T.Li, J.Yu, and K.Varadhan, Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy, Internet RFC 1519

[3] Y.Rekhter and T.Li, A Border Gateway Protocol 4 (BGP-4), Internet RFC 1771

[4] Y.Rekhter, B.Moskowitz, D.Karrenberg, G.J.de Groot and E.Lear, Address Allocation for Private Internets, Internet RFC 1918

[5] Y.Rekhter and T.Li, Implications of Various Address Allocation Policies for Internet Routing, Internet RFC 2008

[6] VK. Hubbard and et. al., Internet Registry IP Allocation Guidelines, Internet RFC 2050

[7] ftp://ftp.arin.net/pub/stats/{arin, ripencc, apnic}

[8] whois.{arin, apnic, ripe}.net

[9] http://www.antc.uoregon.edu/route-views

[10] http://data.ris.ripe.net/rrc00

[11] http://data.ris.ripe.net/rrc01

[12] G.Huston, Analyzing the Internet's BGP Routing Table, *Internet Protocol Journal*, January 2001.

[13] Tian Bu, Lixin Gao and Don Towsley, On Routing Table Growth, *Proceedings of Globe Internet 2002*. http://www-net.cs.umass.edu/ tbu/

[14] S.Bellovin, R.Bush, T.G.Griffin and J.Rexford, Slowing Routing Table Growth by Filtering Based on Address Allocation Policies. http://citeseer.nj.nec.com/493410.html

[15] CIDR Working Group, Implications of Various Address Allocation Policies for Internet Routing, *Internet Draft*, draft-ietf-cidrd-addr-ownership-07.txt

[16] G. Huston, Commentary on Inter-Domain Routing in the Internet, Internet RFC3221

[17] Phillip Smith, http://www.apnic.net/stats/bgp