Policy Proposal: Limit Address Allocation to Extend the Lifetime of IPv4 in the APNIC Region

Zheng Dong
School of Informatics and Computing
Indiana University Bloomington
zhdong@indiana.edu

L. Jean Camp
School of Informatics and Computing
Indiana University Bloomington
ljcamp@indiana.edu

ABSTRACT

The fourth revision of the Internet protocol (IPv4) has been so widely implemented, that the IPv4 address pool approaches full allocation. Currently, only 8% of the IPv4 address blocks are not yet allocated by IANA. Based on the allocation history, IPv4 addresses will be fully allocated in 2012 [11]. In this work, we analyze the historical allocation records of APNIC, the organization which manages IP address allocations in the Asia-Pacific region. We propose three policies that can be implemented by APNIC. We then validate our policies by simulating the APNIC allocations. The experiments show that the lifetime of IPv4 could be significantly extended.

1. INTRODUCTION

The pool of unallocated IPv4 addresses is small and rapidly decreasing. With all /8 blocks allocated, the IANA address pool is predicted to be depleted on October 2011. The most recent prediction as of this writing indicates that all Internet Registries (IR) would have their address resources fully allocated by August 2012. This is only two years from now [11]. IPv6 is proposed as one long term solution. Others argue that there is no effective limit to address sharing, with multi-level NATS. Regardless of the long term solution, what is required are effective policies to slow full allocation or mitigate the implications of the shortage following full allocation. Previously, we examined a set of policy options that could be considered by ARIN [4]. Here we expand our research to the Asia-Pacific region, the area with the largest population, among the shortest Internet histories, and home to the most rapid Internet growth.

We examined the following three allocation policies in the case of APNIC.

1. Allocate addresses only if the historical total allocation for an organization doesn't exceed a threshold.
2. Select a maximum size allocation for any organization.
3. Set annual allocation upper bounds based on the currently available resources and use this to select an exhaustion year.

The remainder of the report is organized as follows: Section 2 introduces the related work on IPv4 address pool exhaustion. Section 3 describes the data processing and experimental design. Section 4 reports potential results for the three policies we proposed. Section 5 summarizes our findings, and concludes the work.

2. RELATED WORK

The risk of IPv4 address depletion was first recognized and addressed in the early 1990's. Class B address resources were predicted to be fully allocated in 1994 based on then-current policies [14]. By changing the class allocation policies, the availability of unallocated IPv4 addresses has been vastly extended. Regardless, full allocation looms.

After the initial Class B crisis had been managed, there were early one-time predictions. One model predicted an exhaustion date of 2037 [10], yet today that is clearly optimistic. A later projection suggested that exhaustion would be in 2019 [9]. Now, there is a prediction model that is both regularly updated and arguably canonical [12], and the most recent estimate date for the depletion of the IANA pool is in 2011,[1] with all Regional Internet Registries (RIR) having allocated their resources in 2012. Observing the urgent allocation situation, many researchers have proposed new policies to delay the allocation of IPv4 or to accelerate the implementation rate of IPv6. IPv6 is one long-term solution for address scarcity. Hain [8] applied several mathematical models to the historical allocation data, and pointed out that since there may be many factors that affect the consumption rate of IPv4 addresses, the resource may be depleted much sooner than previously predicted. Elmore [7] identified valid measures of IPv6 diffusion and used classic diffusion models to bound the uncertainty in those measures. Even given the order of magnitude uncertainty, the work concluded that there is no reasonable case for diffusion of IPv6 before IPv4 full allocation. Edelman [6] argued that network incentives impede transition to IPv6, and the rational choice for many organizations is the intensive long-term use of IPv4. The author also introduced a paid-transfer system that would enhance the adoption of the IPv6, given strict assumptions about the demand curve.

In the APNIC region, one proposal [2] recommended that account holders for the final /8 space of IPv4 addresses demonstrate IPv6 deployment or a transition plan in order to stipulate the adoptions of IPv6. Another proposal [1]

[1]Clearly most recent to the date which we write this draft.

GTIP 2010 Dec. 7, 2010, Austin, Texas USA

concentrated on processing and recording the address transfers between current members of APNIC. The author argued that this would help APNIC maintain an accurate record of resources, and incentivize the return or transfer of unused addresses.

Our proposal is different from the previous work in two aspects. First, we concentrate only on Asia and the Pacific region. Our analysis includes only historical allocation data and Whois database directly from the APNIC. Therefore, our policy and experimental results are applicable only to this particular region. Secondly, our policies are simple and could be implemented immediately without the construction of a market, for example. (We acknowledge the existence of stakeholders who may object to these policies.) In addition, each policy is independent but could combine with other policies.

3. EXPERIMENT SETUP

3.1 Data Sources

In this section we describe our data compilation, filtering, and modeling. Our main data source is the public allocation history statistical files stored on the APNIC FTP server. Allocations of three types of resources were recorded: Autonomous System Numbers (ASN), IPv4 addresses, and IPv6 addresses. Allocation status was updated on a monthly basis between May 2001 and May 2003, and a daily basis starting on May 3rd, 2003. In each file, resources that had been allocated were listed following the format shown in Table 1.

Table 1: Allocation Record Format

Field name	Description	Value
registry	registry that allocates the resource	{apnic} for the Asia-Pacific region
cc	country code	e.g. cn, jp, in
type	type of Internet number resource	{asn, ipv4, ipv6}
start	first address of the range	e.g. 202.197.0.0
value	addresses count in the range	e.g. 1024
date	allocation date	e.g. 20010201
status	allocation status	{allocated, assigned}

Table 2: Selected Fields for Whois Database

Field name	Description	Value
inetnum	start and end address of the block	e.g. 121.96.18.48 - 121.96.18.63
netname	block identification	e.g. BAYAN_SBOAAP
country	country of the assigned block	e.g. cn, jp, in
descr	description of the assigned block	name and physical address of the end user
status	current allocation status of the block	{assigned/allocated portable/non-portable}

In some early record files, the field "value" referred to the CIDR prefix size instead of the host count. For example, the record "*apnic*|*nz*|*ipv*4|202.0.48.0|22|19930101|*allocated*" indicates the IP range 202.0.48.0 - 202.0.51.255. Values in the field "status" can be "allocated" or "assigned". For entries that are marked "assigned", the given block is allocated to an identified end user, and that block is not expected to be sub-allocated or transferred. For those entries marked "allocated", the resource range is assigned to an IR for sub-allocations. Such an IR might be a national registry, Internet service provider (ISP), or a large organizationally diverse final user (e.g. a multinational corporation). APNIC doesn't maintain the sub-allocation information in their public statistical files. Therefore, a single block in the statistical file may correspond to multiple end users. There is no information provided for the organizations to which the blocks were assigned, nor is sub-allocation recorded. Therefore, accurate analysis required another data source.

The second important data source was the bulk Whois database file provided by the APNIC. The bulk Whois data are updated daily, and previous files are deleted. Compared to the public allocation history, the Whois database provides more details about end users for each assigned block. User data includes organization name, description, contact person, managed organization, and several additional useful fields as listed in Table 2. Most important for our purposes is the field "netname" which we initially treated as a unique organizational identifier as described below.

The only missing part for our research is the exact allocation date for an address block. In some Whois entries, however, there are "date-changed" columns indicating the time a contact person or managed organization record is updated, but it is obviously not a reliable source of an allocation date. As a consequence, we need additional data to perform historical statistics.

3.2 Assumptions

The experimental design requires making the assumptions described in this section.

Allocation statistical files became available on the public FTP server on May 1st 2001. For allocations prior to that date, we assume they were completely recorded in the allocation history statistical file *20010501*. In other words, no IPv4 addresses were returned to APNIC before May 2001.

We assume sub-allocations take place on the same day a block is allocated to an IR by APNIC, since there is no reliable time stamp for sub-allocations recorded in the Whois database. That is to say, we used the same date recorded in allocation history files for the sub-allocation records as for allocations in the Whois database.

We may distinguish the ownership of blocks by comparing values in the field "description". Initially, we considered the "netname" field in the Whois entries, which appeared to be a reasonable method to identify different organizations. However, large corporations may be assigned different netnames for each branch, while it is also possible that different methods of abbreviation were used in each allocation. As advised by the APNIC helpdesk, detailed organizational information for most allocations is contained in "description". In practice, however, the description field may take up to 8 lines. Our approach takes the first line as an organization identifier; if the first line doesn't contain English words (e.g. starts with "∼"), we then consider the second line of description as the identifier.

3.3 Data Processing

Data Sources Combination. As mentioned above, two

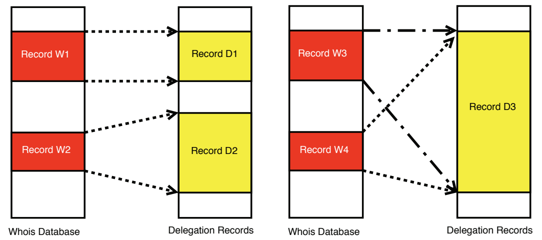

(a) (b)

Figure 1: Whois Database Entries Join Public Allocation Records

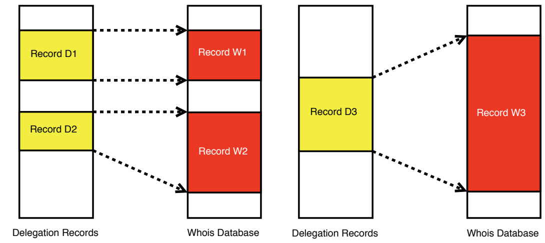

(a) Step 1 (b) Step 2

Figure 2: Public Allocation Records Join Whois Database Entries

data sources were utilized in our experiments: the Whois database entries and public historical allocation records. Data from both sources were obtained on March 1, 2010 from the FTP server of APNIC, and then exported to a MySql database.

Since both data sources contain part of required information, linking two of them is required to generate a complete record. Considering sub-allocations happen after APNIC allocates a block to an IR, the Whois database appears to contain more information about "end users". One possible approach would be to join the two data sources starting with Whois. In this approach, we identify matches in the public allocation records for each Whois entry. Figure 1 illustrates the three approaches we took towards the two data sources. As one example, the record W1 block in the Whois database and record D1 block in the allocation history are exactly the same in both their start points and end points. We link the two records with a clear match, and this is considered to be our best case. Most records, however, are linked based on the previously described assumption that sub-allocation occurs simultaneously with allocation; we also leverage the fact that blocks recorded in allocation history are normally larger than the size of blocks in Whois. For another example, record W2 in the Whois database has the same start point as the record D2 in the allocation history, but the size of the Whois entry is smaller than the allocation history. We may also link multiple Whois entries using one allocation that contains multiple sub-allocations as shown in Figure 1(b). For a third example, both record W3 and W4 in the Whois database are within the block D3 in allocation history, so we simply link W3 and W4 with the allocation record D3.

The benefit of this approach is clearly that detailed information for sub-allocations is created in a consistent manner. However, this approach doesn't provide reliable estimates on the total allocated addresses. For years prior to 2003, the difference between officially published data [3] and the experimental result becomes larger than half of the allocation amount, which is unacceptable. One possible reason is that records in allocation history date back to the year 1985, while the block Whois data is always the most recent.

We then switched to an alternative method, to join the databases starting with the public allocation records. The main idea of this method is to make the best use of allocation history, especially the size of allocation history blocks. In this case, we simply assume that a Whois block is large enough to accommodate allocation histories. In other words, the size of a Whois block is greater than or equal to the size of an allocation block.

We implemented this approach in two ways: First, for each allocation block, we searched for a Whois entry with the same start point and of a larger or equal size. As shown in Figure 2(a), Record D1 and W1 are a perfect match since both start and end point of the records are the same. Record D2 in the allocation history and W2 in the Whois database start at the same address, but the Whois block is larger than the block recorded in the allocation history. These two records are linked based on the observation that one Whois entry may be associated with multiple IR allocations. In other words, an allocation record may contribute only to part of a Whois entry. Suppose we have an allocation block "$apnic|AU|ipv4|58.6.0.0|32768|20050128|allocated$" and two entries in Whois: "58.0.0.0 − 58.255.255.255" and "58.6.0.0 − 58.6.128.255". It is obvious that the latter Whois entry would provide the correct organizational information although the allocation block is within the range of both Whois entries.

For unlinked records, we used a second approach. We linked two records if the allocation record was fully within a Whois record. As mentioned above, a Whois entry may be formed after multiple consecutive allocations, it is therefore reasonable to link blocks like D3 and W3 as illustrated in Figure 2(b).

Combining the two options described here ensures a closer estimate of the total address allocations, and therefore results in a more accurate prediction. Admittedly, although the total number of allocations is much closer to the officially published values in the first approach, there would be some sub-allocations made by IRs that are not included in the count for each organization; for example, branches of corporations or small allocations may also overlooked. However, based on the resources we have in hand, the second approach is preferable.

4. DATA ANALYSIS

4.1 Data Features

Compared to the IPv4 address allocation data we collected from ARIN, the statistical allocation data recorded by APNIC showed three distinctive features.

First, APNIC is quite diverse in terms of jurisdiction. The allocation history in APNIC includes 58 countries, which is twice as many as that of ARIN. Nations in ARIN fall into two discrete categories. There are the large NAFTA nations and then the Caribbean nations. In contrast to this bifurcation, the APNIC countries cannot be so easily classified. Different countries bring diverse policies, levels of network readiness, wealth, distribution of wealth and diverse corporations, organizations and policies to govern IPv4 allocations.

Second, National Internet Registries (NIR) play a more important role in allocations. As mentioned in the previous section, there were two possible statuses for a block in allocation history files, "assigned" and "allocated". Instead of applying resources directly from APNIC, most of the blocks were first allocated to an NIR, and then sub-allocated to the end-users, and therefore showed "allocated" in allocation files. Comparing the block status in the allocation files created on March 3 2010, ARIN showed 29% of its total IPv4 blocks "allocated", while for APNIC, the ratio was 93%. Recall that RIRs such as ARIN and APNIC do not keep sub-allocation information in their allocation history files. Therefore, the Whois database file, which is our second data source, is especially important in the Asia-Pacific region.

Third, Internet history in the Asia-Pacific region is different from that in North America. Owing to the uneven development of the economy and historical reasons, the Internet development in the Asia-Pacific region was not as early as that of the North America region. However, allocation history after 2003 in APNIC showed that the speed of Internet implementation dramatically increased. Therefore, in the allocation analysis for APNIC, we shall expect a different trend than was depicted in the figure of ARIN.

4.2 Estimate of Addresses Pool for APNIC

According to current statistics [11], IANA has only 22 /8 blocks unallocated in its address pool. If we assume that IANA allocates these blocks evenly to the five RIRs, each one of them gets 4.4 /8 [11]. Roughly 2.2 /8 blocks that are currently administered by APNIC haven't been allocated [12]. According to the report of IANA [13], the APNIC allocated 36 /8 blocks, and administered 6 /8 blocks. Therefore, we make our estimation on the total number of available IPv4 addresses for the Asia-Pacific region by simply combining the two sources, i.e. the APNIC will have 46.4 /8 blocks (778,462,822 addresses) in total. Like some former studies [11][8], historical allocations of APNIC have been used in the calculation of the exhaustion date of the APNIC address pool. However, we included confidence interval (CI) [5] for both the projection of current policy and our proposed model. Typically, the confidence interval describes the possibility of a certain result in the future. For example, a 95% confidence interval indicates that the real data would fall into the designated area with a probability of 95%.

Based on historical records from two data sources, we plotted figures with respect to the number of allocations, organizations, and addresses. The number of allocations indicates the number of records that existed in allocation history (our first data source) as depicted in Figure 4. When we evaluated allocations per year and netname (which we initially used to identify different organizations), a similar trend resulted; this is shown in Figure 3. Note there is a slight difference between the above graphs, since multiple allocations may be allocated to a single organization. For example, if ten blocks were allocated to an organization in 1995, ten would be shown in Figure 4, while in Figure 3 it would show as one. According to the two figures, there is a significant decline starting from the year 1997, and after that both the number of organizations and allocations increase

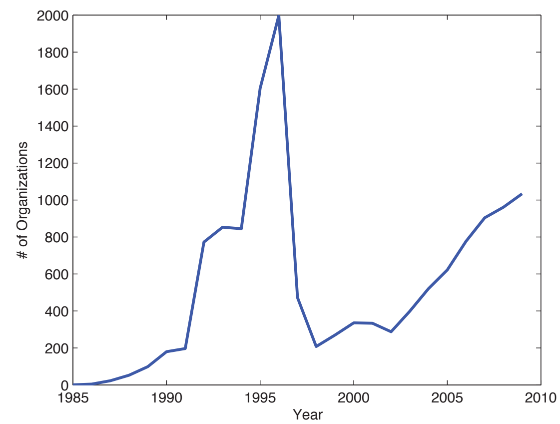

Figure 3: Number of Organizations Per Year

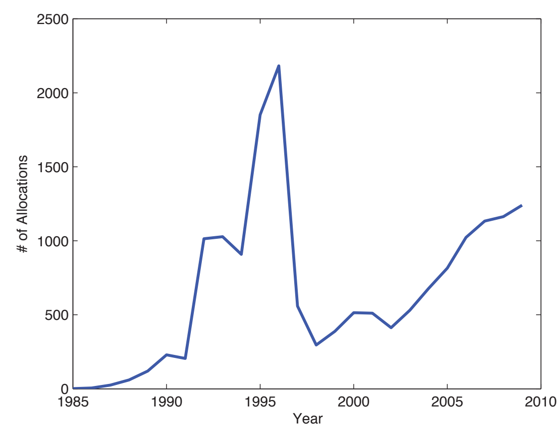

Figure 4: Number of Allocations Per Year

monotonically. As a result, our linear projection on address allocation starts from 1998. Figure 5 illustrates the number of addresses allocated by APNIC in the past 25 years. This figure illustrates the number of addresses allocated each year, and doesn't reflect a cumulative situation. It is therefore clear that though the allocations and organizations have decreased during part of our history, the actual summation of allocated addresses is increasing monotonically.

In addition, we used a cumulative method to predict the exhaustion date of the IPv4 address pool. In other words, allocations that occurred in previous months need to be included in the total for all following months. In order to achieve more accurate fits for the historical allocation data, we performed a cubic polynomial fit first, since we can expect a fit that has only a minor difference from the entire original allocation data between June 1985 and February 2010. Results for the cubic polynomial model are shown in Figure 7. As shown in Figure 6, a quadratic polynomial fit was also performed to validate the predicted exhaustion date. In this case, however, we need to avoid significant changes in the number of organizations. Therefore, our quadratic fit is based only on the historical data from 1998 to 2010. Interestingly, two sets of polynomial fit indi-

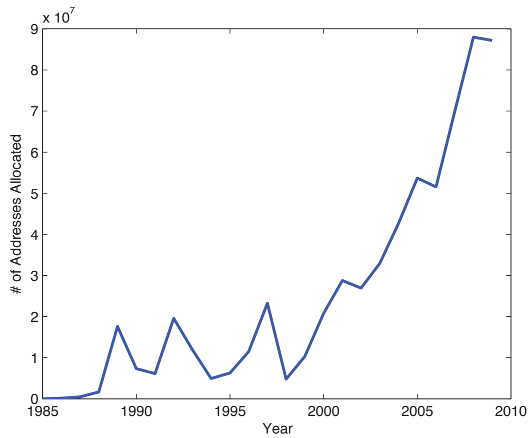

Figure 5: Number of Addresses Per Year

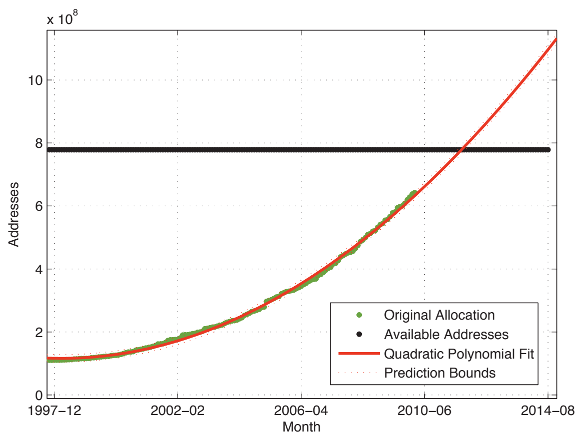

Figure 6: Exhaustion Quadratic Polynomial Projection with Current Trend

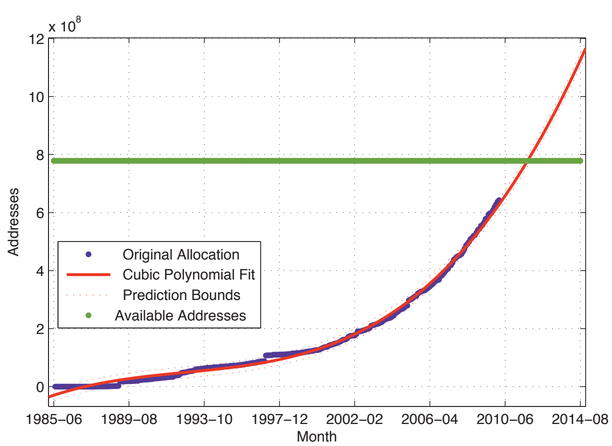

Figure 7: Exhaustion Cubic Polynomial Projection with Current Trend

Table 3: Summary of Exhaustion Projections

Polynomial Estimate	Lower Bound	Upper Bound
2011-09	2011-07	2011-11

Table 4: Organization Statistics

Organization Category	Number of Organizations	Average of Allocated Addresses
IP $\geq$ /8 block	6	673 /16s
IP in [/12 block, /8 block]	60	54 /16s
IP in [/16 block, /12 block]	1004	542 /24s
IP in [/20 block, /16 block]	1697	50 /24s
IP $\leq$ /20 block	6639	2 /24s

cated the same exhaustion date, as listed in Table 3. Given that it has been only one and half year since the depletion date, effective allocation policies need to be enforced.

4.3 Organizational Distribution

We divided organizations into five categories based on the number of addresses delegated in allocation history. The first category consists of 6 organizations that own more than /8 blocks (16,777,216 addresses), and each of them got 673 /16s on average. A similar pattern emerged among organizations that were formerly allocated between a /12 address block and a /8 address block. There were 60 members in this category, while the average number of blocks was 54 /16s. However, the last three categories have had relatively small allocations. There were 9340 organizations in these categories. Based on this observation, it is possible that organizations in the last three categories have more urgent need for address resources, and have also lower total organizational network expertise. Our policies, which benefit the last three categories, could solve the urgent needs of applying for new address resources for most organizations. In addition, no prediction can be done at this time for future innovators who now have zero organizations.

4.4 Analysis of Policy 1: Addresses Reserved for those with Smaller Allocations

Recall that in Table 4, a small number of organizations received a large portion of address resources, while the majority of them only get few allocations. Based on this observation, we propose a policy that allows a future allocation only when the historical total number of addresses of an organization is below a certain threshold. This policy may lead large organizations to revocate previously allocated addresses, or to register new blocks by applying addresses directly from their branches or subsidiaries. When evaluating the cost for these countermeasures, it applies only to the organizations with large address allocation.

We examine the performance of this policy by setting up three experiments with the following thresholds: /12 block (1,048,576 addresses), /16 block (65,536 addresses), and /20 block (4,096 addresses). Figure 8 shows the historical allocations that would be prevented by this policy. According to the figure, at least 3×10^8 addresses could be considered as "residues" meaning they would not have been distributed.

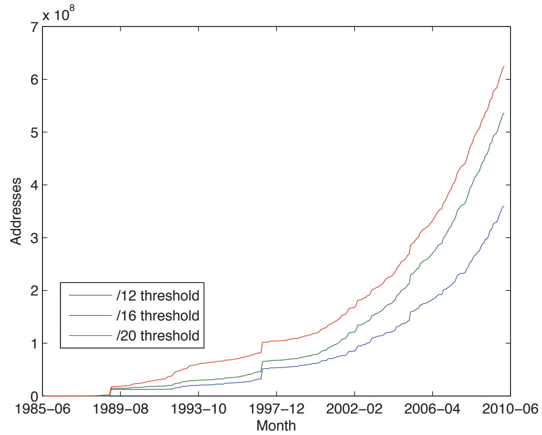

Figure 8: Residue Addresses Based on Thresholds

We implemented experiments by following process.
First, we subtract these monthly residues from original allocation data.
Second, we calculate a linear and a polynomial fit on the data we get from the first step.
Third, we add the total number of residues back to the final result.

Figure 9 shows a predicted exhaustion date when a /12 threshold is enforced. The polynomial result indicates a depletion date in 2012, which extends our previous prediction by one year; while the linear fit projected the depletion happens in 2018, which is eight years from now. Similarly, Figure 10 illustrates the polynomial and linear projection with a /16 threshold, and predicts the exhaustion date to be in the year 2019 and 2037 respectively. When we switch to a threshold of /20, the expected depletion dates are extended to even the year 2083, as shown in Figure 11. We summarized results from different approaches in Table 5.

The substantive argument for this policy is that organizations with large number of addresses allocated are often prominent in both the number of skilled technology specialists and IT-related experience. In addition, considering the possible costs implementing IPv6, these organizations are more likely to benefit from implementing new version of the IP protocol compared to small organizations. Furthermore, we would expect a much quicker implementation rate for IPv6 if large corporations migrate first, since the majority of address resources were allocated to them, and small organizations can learn experience from large companies and modify their migration plan accordingly.

4.5 Analysis of Policy 2: Fixed Addresses Per Allocation

In the previous section, we proposed that only a given threshold of IPv4 addresses may be allocated to each organization "lifetime". This policy, however, may be unfair to large-scale organizations. Therefore, we offer a second proposed policy that allows a fixed number of addresses to be allocated to each request. To avoid over-allocating resources, Selecting a reasonable threshold for allocation size is extremely important. For requests that exceed the limit, APNIC allocates only the threshold, otherwise provides the number of addresses requested. For example, suppose the

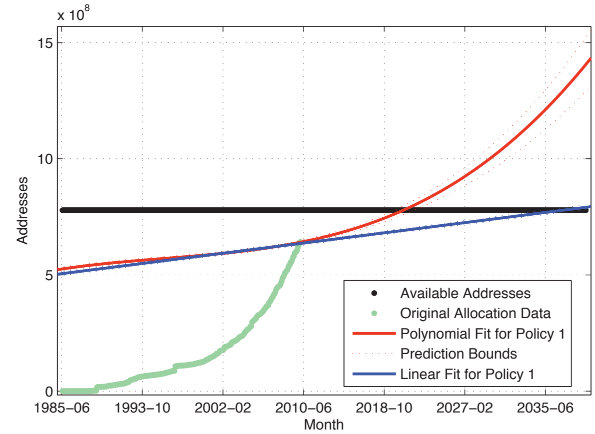

Figure 10: Allocation Projection if Organizations with /16 Blocks No Longer Receive Addresses

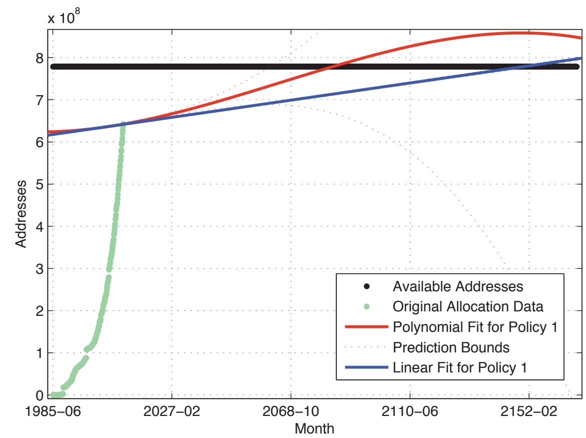

Figure 11: Allocation Projection if Organizations with /20 Blocks No Longer Receive Addresses

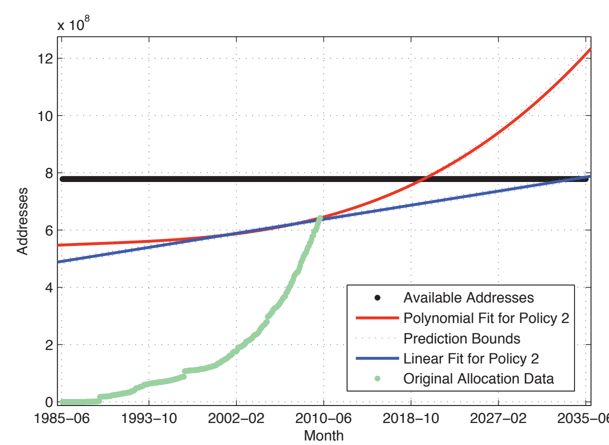

Figure 12: Allocation Projection Each Event Allocates at Most /18 Block

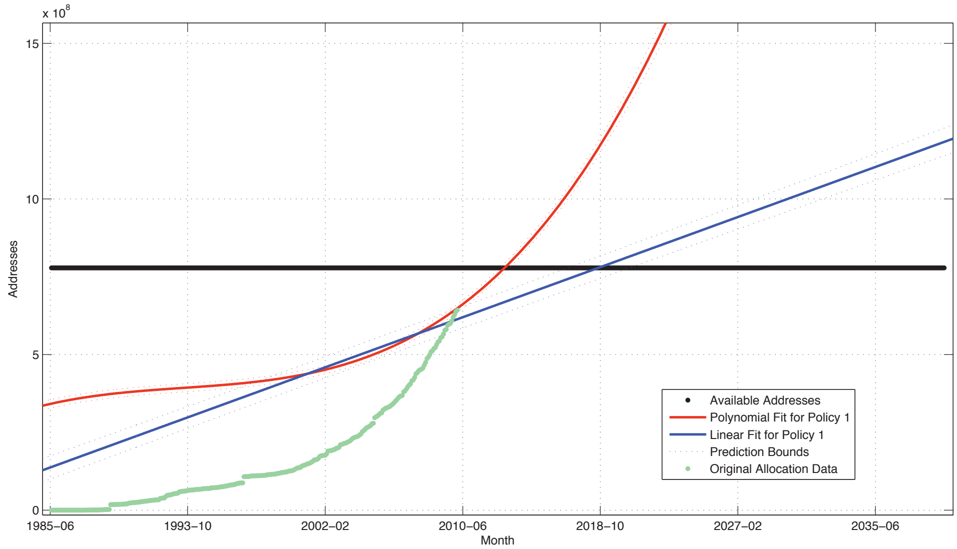

Figure 9: Allocation Projection if Organizations with /12 Blocks No Longer Receive Addresses

Table 5: Projected Exhaustion Date for Policy 1

Projection	Threshold /12			Threshold /16			Threshold /20		
	Projected	-95%CI	+95%CI	Projected	-95%CI	+95%CI	Projected	-95%CI	+95%CI
Linear	2018-09	2016-11	2020-07	2037-03	2035-05	2039-01	2149-11	2146-02	2153-10
Polynomial	2012-12	2012-09	2013-03	2020-10	2019-10	2021-11	2083-05	2061-08	-

threshold has been set to a /20 block (4,096 addresses). Organization A previously received no allocation, and requests 1,024 addresses; organization B, a large address resource holder, requests 65,536 addresses. In this case, A would get 1,024 addresses as it requested. In comparison, organization B would only get 4,096. Note that we don't consider historical allocations in this policy, so the result doesn't change whether A or B already controlled a large number of address resources or not.

We validate the effectiveness of this policy by designing two experiments with thresholds of a /18 block (16,384 addresses) and a /20 block (4,096 addresses) respectively. The experiments steps are similar to those in policy 1, and they follow here.

1. First, we calculate the number of historical addresses that would not have been allocated based on a threshold, and subtract monthly residues from the original allocation data.
2. Second, we calculate a linear and a polynomial fit on the data we get from the first step.
3. Third, we add the total historical residues back to the final result for correcting historical data.

Figure 12 illustrates the projection when we set the fixed allocation size to be a /18 block. The polynomial fitting result shows an exhaustion month of January 2020, while the linear result gives a depletion date of more than 20 years from now. Experiments with a /20 block threshold give out an even longer extension, 2031 for the polynomial and 2070 for linear fit, as depicted in Figure 13. We summarized our results in Table 6.

The argument for this policy is that it is impartial to all organizations. It is possible that large corporations need relatively more address resources to maintain their everyday operations, and support their business. Therefore, we allow a certain amount of resources to be allocated per request, so that needs from large organizations that may be urgent are not ignored. According to our experiments, we can expect a long extension of the exhaustion date, to even more than 20 years.

4.6 Analysis of Policy 3: Fixed Address Allocation Pool Per Year

We analyzed the APNIC allocation events for the year 2009, and plotted a distribution graph in Figure 14 which uses the X-axis to denote allocation events, and Y-axis for the number of addresses allocated in each event. According to that figure, only about 10% of total allocations achieved 10^8 addresses per event, while the majority of allocations show less than 10^7 addresses. Based on this observation, APNIC can indicate a desired exhaustion date, and calculate an upper bound of allocations made each year according to the exhaustion date. Specifically, we may simplify the problem by calculating the annual allocation upper bound from a division of available addresses by the designated years. For example, suppose 3,626 /16 blocks currently present in the APNIC address pool, and we expect the depletion to happen in 10 years, then the annual allocation limit would be 362.6 /16 blocks.

We examined this policy by calculating percents of requests that can be accommodated with thresholds from 10

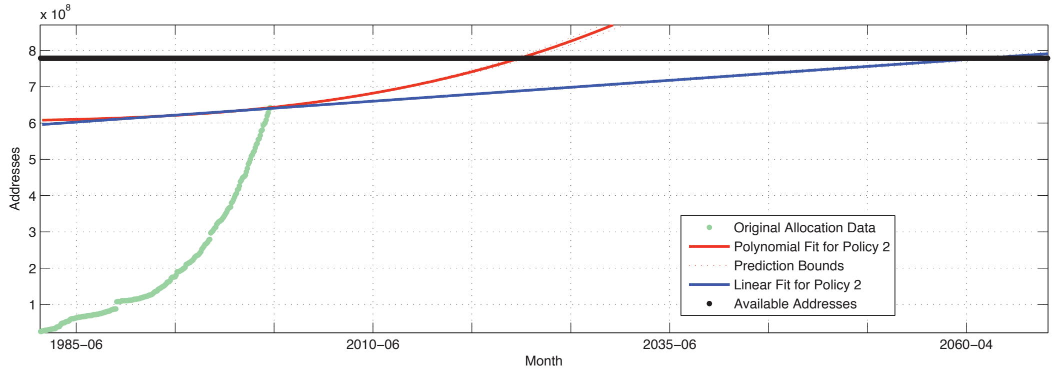

Figure 13: Allocation Projection Each Event Allocates at Most /20 Block

Table 6: Projected Exhaustion Date for Policy 2

Projection	Threshold /18			Threshold /20		
	Projected	-95%CI	+95%CI	Projected	-95%CI	+95%CI
Linear	2034-05	2032-12	2035-10	2070-03	2068-03	2072-06
Polynomial	2020-01	2019-09	2020-04	2031-03	2030-06	2031-11

Table 7: Requests Can be Met Under Policy 3

Lifetime of IPv4 (Year)	Upper Bound for Annual Allocation	Percent of Fulfilled Requests
10	363 /16s	94.5%
20	181 /16s	89.0%
30	121 /16s	85.2%
40	91 /16s	82.7%
50	73 /16s	80.6%

to 50, and results were summarized in Table 7. According to the table, nearly 94.5% of requests would be met given the annual allocations below 363 /16 blocks. Even if we extend the lifetime to be 50 years, it is still possible to fulfil 80% of the requests.

The argument for this policy is also organizational non-discrimination. In order to make sure most requests are satisfied, APNIC could combine this policy with some other policies, so that small allocations might be met first. Alternatively, IRs could use auctions, to deal with this scarce resource, thus applying a more limited scope to the proposed v4 markets as proposed by Edelman [6].

5. CONCLUSION

Our experiments were based on two data sources: the APNIC historical allocation history files, and bulk Whois data entries. We combined the two data sources to generate complete records for our statistics. We then calculated organizational distribution according to the historical allocation data, and discovered that a small number of organizations were allocated the majority of resources.

Based on this observation, we proposed three allocation policies. The first policy sets up a lifetime limit for each organization. In other words, addresses allocated to an organization can not exceed a certain threshold. Considering the possible negative effect on large-scale organizations, we proposed the second policy, which restricts only on the size of address blocks in each allocation. The third policy provides APNIC an option to control the predicted exhaustion date. By dividing the number of available addresses by the desired allocation years, we can easily calculate an annual allocation upper bound. For each allocation policy, we performed several experiments to validate the effectiveness. According to the results, our policies show a significant extension on the depletion date of the IPv4 address pool. Some policies even achieve dates that are 20 years away from the previous predictions.

Our result is preliminary, and future work is still needed to better understand the status of address allocations. Firstly, we have only limited resources in our data analysis. There is no information recorded in the APNIC allocation history files if sub-allocations are made by other IRs within the APNIC. Although the Whois database contains a complete list of different allocations to end users, a significant amount of data is missing if we simply sum up blocks that are recorded in the Whois database. By joining the two data sources, we unavoidably ignored a certain number of sub-allocations. However, since our approach may have enlarged the size of organizations, we would expect an even better outcome given were our proposals enforced. Since some part of our proposals may ignore certain factors (e.g. Policy 1 may not be fair to large organizations), the regional registry would need to implement multiple policies to get a better outcome. Thirdly, all predictions were made based on the historical allocation data, and we did not consider much about the possible divergence that may occur in the future.

The greatest risk for this model is the assumption that IANA will allocate its available address blocks evenly to the five RIRs. We understand that due to the rapid growth of Internet in the Asia-Pacific region, more addresses have been allocated to this area. For example, APNIC received four /8 blocks from IANA in 2009 [13], half of the total allocations IANA made in that year. However, since the available address pool for APNIC may be larger than our calculation, the implementation results of our policies should also be better than what we predicted in this proposal.

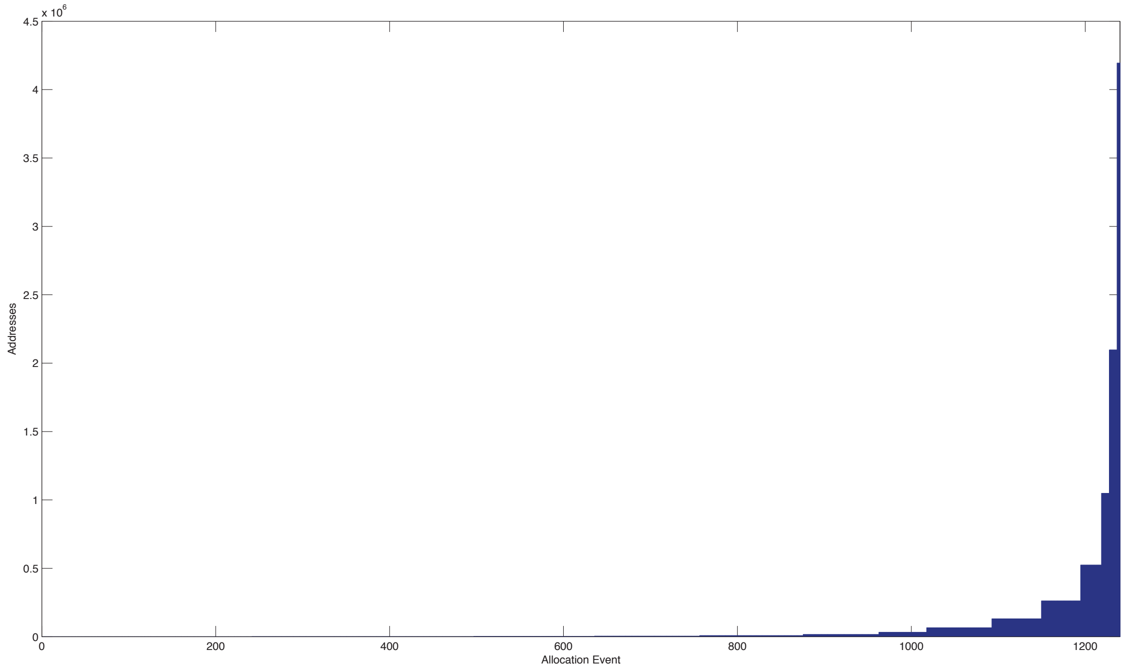

Figure 14: Addresses Distribution in 2009

6. REFERENCES

[1] Ipv4 address transfers. http://www.apnic.net/__data/assets/text_file/0009/12420/prop-050-v005.txt, Last Checked: 04/15/2010.

[2] Ipv6 deployment criteria for ipv4 final /8 delegations. http://www.apnic.net/__data/assets/text_file/0017/17009/prop-078-v002.txt, Last Checked: 04/15/2010.

[3] APNIC. Ipv4 distribution. http://www.apnic.net/publications/research-and-insights/stats/ipv4-distribution, Last Checked: 04/15/2010.

[4] L. J. Camp, R. Wang, and X. Zhou. Policy proposal: Limit the address allocation to extend the lifetime of ipv4. *ARIN*, 2009.

[5] C. Chatfield. Calculating interval forecasts. *Journal of Business & Economic Statistics*, 11(2):121–135, 1993.

[6] B. Edelman. Running out of numbers? the impending scarcity of ipv4 addresses and what to do about it. http://www.benedelman.org/publications/runningout-draft.pdf, Last Checked: 04/15/2010.

[7] H. Elmore, L. J. Camp, and B. Stephens. Diffusion and adoption of ipv6 in the arin region. *Workshop on the Economics of. Information Security (WEIS)*, 2008.

[8] T. Hain. A pragmatic report on ipv4 address space consumption. *The Internet Protocol Journal*, 8(3), 2005.

[9] G. Huston. Ipv4 - how long have we got? http://www.potaroo.net/ispcol/2003-08/ale.pdf, 2003.

[10] G. Huston. Ipv4 address lifetime expectancy revisited. hhttp://www.iepg.org/november2003/, Last Checked: 04/15/2010.

[11] G. Huston. Ipv4 address report. http://www.potaroo.net/tools/ipv4/index.html, Last Checked: 04/15/2010.

[12] G. Huston. Ipv4 all allocation report. http://bgp.potaroo.net/ipv4-stats/allocated-all.html, Last Checked: 04/15/2010.

[13] IANA. Iana ipv4 address space registry. http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml, Last Checked: 04/15/2010.

[14] M. W. Murhammer, O. Atakan, S. Bretz, L. R. Pugh, K. Suzuki, and D. H. Wood. *TCP/IP Tutorial and Technical Overview*. IBM Corporation, Research Triangle Park, NC, 1998.