

Active Probing to Classify Internet Address Blocks

(poster abstract)

Xue Cai John Heidemann
USC/Information Sciences Institute, {xuecai,johnh}@isi.edu

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Network topology*; C.2.3 [Computer-Communication Networks]: Network Operations—*Network management*

General Terms: Measurement

Keywords: Internet address allocation, survey, pattern analysis, clustering, classification, availability, volatility

1. INTRODUCTION

Previous Internet topology studies mostly focused on AS- and router-level topologies [1,2,4,8,9], providing insight into AS relationships and interdomain routing. However, relatively little is known today about the demographics of Internet edge hosts and the use of the IPv4 address space. Since the transition to classless routing [3], external observers have only limited view into how IP address blocks are allocated, managed, and used. Only recently have researchers tried to characterize address usage [11], with a methodology limited to addresses used by clients.

In this poster we begin to explore the potential of *active probing* and *external classification* of address block usage. Our work makes three assumptions: a significant number of Internet addresses will respond to external probes, some patterns of repeated probes correspond can indicate different address usage, and contiguous addresses are often used for similar purposes. While there are cases where these assumptions do not hold, in recent work we study probing effectiveness with ICMP and TCP, showing they detect the majority of used addresses [5]. In this poster we expand on the second hypothesis, using probe results to study address allocation strategies and to infer address usage (always be occupied or not in use, use by stable or frequently inaccessible hosts, etc.).

In ongoing work we are evaluating the accuracy of this approach and these assumptions; in this poster we suggest this as a promising new direction for research. We expand on the goals of this research, the applications of classification, and our preliminary results below.

2. GOALS AND APPLICATIONS

Our goal is to classify Internet address blocks based how addresses respond to frequent probes over the source of one week. By developing and validating this technique, we hope to answer specific questions, including:

Do groups of blocks show consistent patterns? How many vary? Groups of consistent patterns will allow clustering and classification of address usage with only external probes.

What are the sizes of blocks that show consistent usage? These patterns will suggest how address blocks are managed at fine granularities (smaller than /24s).

Can we map consistent, popular patterns to operational network usage? Can we identify groups of servers, cable or DSL customers, dial-up users, etc., by ping responses?

While these questions may seem academic, they enable answering important operational and policy questions. We expect that our approach will be approximate, yet even *approximate* answers to these new questions would be a huge step forward from today's understanding based primarily on anecdote or occasional surveys of providers.

A first important policy question is that understanding the address allocation and usage can give the academic community and the ISPs a better vision of current IPv4 address usage. Best estimates suggest ICANN will allocate its last free block of IPv4 addresses as soon as 2011 [6], so understanding address utilization is essential to Internet governance. For example, in each block, is addresses utilization high or low? Would new methods of address assignment improve efficiency? We expect that answering these questions and quantifying the costs of IPv4 address management will motivate IPv6 adoption and guide methods of IPv6 allocation.

Secondly, understanding dynamic usage of IP addresses will be helpful to the network research community concerning security. Xie et al. have begun to explore this question with a goal of identifying dynamic blocks to assist spam prevention [11]. Our method uses a completely different approach and so can extend and corroborate their findings.

Finally, understanding IP address usage patterns may help in locating the servers and clients, thus allows the service providers to better distribute their services on the Internet. Combining block usage with address geolocation can identify geographic regions of clients that are distant from services.

3. METHODOLOGY

Our methodology consist of four steps: surveying addresses with probes, pattern analysis of individual addresses, clustering to generalize these results, and broad classification.

Survey: We frequently probe a fraction of Internet address space [5]. We selected 24,000 /24 blocks (1% of the allocated address space), used ICMP Echo Request packets, and actively probed each IP addresses in these blocks every 11 minutes for around a week. While ICMP can be blocked, studies of a university and random addresses confirm that it is the most accurate method of active probing (more accurate than TCP) and solicits relatively few complaints. Our results here are based on data from June 2007 [10].

Pattern Analysis: We analyze the usage of each IP address based on the survey replies. We define several metrics to characterize addresses. First, for each address, we consider it to be up when it responds to pings, and we define *up durations* (U_i , for each uptime i) as the time from the first response until the next non-response. (Precision of up durations estimates is limited by 11-minute probe interval.)

We define *availability* as the sum of all up durations, normalized by total survey time. We define *volatility* as the number of up durations, normalized by the maximum number of possible state changes (half the number of probes). These metrics define an (A, V) plain of address blocks.

We also consider the distribution of lengths of up durations, including mean, median and maximum.

Clustering: We cluster adjacent IP addresses with similar usage (based on the above metrics) into blocks. We are currently experimenting with different clustering algorithms to explore all possible blocks with prefix lengths from /29 through /24 (groups of 8 to 256 addresses).

Classification: We classify the blocks we got from *clustering* into five categories: servers, and blocks that are stable, intermittent, underutilized, or unclassifiable. We trained our classification by manual examination of hostnames and services in several hundred blocks, and then we verify that it works by evaluating it on hundreds of other, randomly chosen blocks.

Server block: highly available and stable (high A , low V)

Stable block: stable, usually continuously up for more than 6 hours (low V , and high median U). Statically assigned addresses usually fall into this category.

Intermittent block: short up durations (low median U). For example, many DSL addresses are reassigned every few hours.

Underutilized block: low A values. Most wireless and some dial-up addresses are severely underutilized.

Unclassifiable block: unclassifiable due to too few responders. We frequently have difficulty classifying blocks of size /29 or smaller.

4. PRELIMINARY OBSERVATIONS

Figure 1 shows our preliminary results: a density plot of mean (A, V) values for all /24 block. Blocks that cluster to smaller block sizes. Darker dots indicate multiple blocks sharing the same availability and volatility. We have a preliminary classification of four categories, as indicated in the plane. The right rectangle around $(A, V) = (0.95, 0.0016)$ are blocks of servers with high availability and a low volatility. The next group, with a high A/V ratio (more than 16.4) are blocks of stable addresses. The A/V ratio indicates the typical duration of address occupation, with a higher ratio suggesting greater stability. The largest group, those with a low A/V ratio (less than 16.4), we call intermittent blocks. They show greater volatility than servers or stable blocks. Finally, the left side we call underutilized, where most addresses are sparsely occupied with low availability.

We developed these classifications by training on a subset of blocks, based on DNS names and occasional service discovery for many addresses. We get DNS data both from manual lookup and an ISC Internet Domain Survey [7]. As an example, we call out eight blocks in the figure. The hostnames in Block A contain “dns”, a common type of server. Blocks B, C have hostnames with the word “static”. Names of addresses in blocks D, E, and F contain the terms “dy-

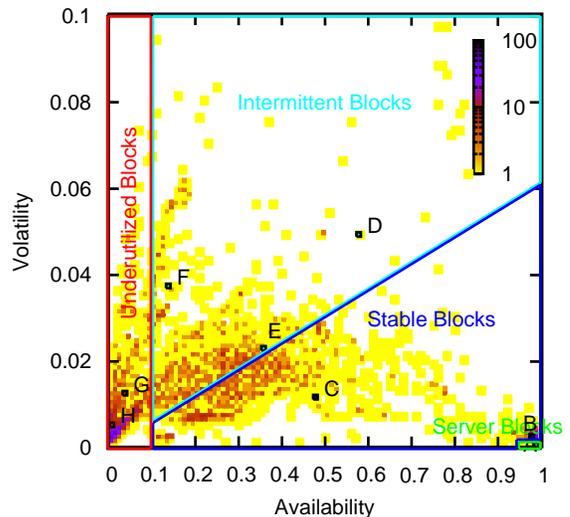


Figure 1: Volatility vs Availability of /24 blocks

amic” or “dsl”. Finally, names of addresses in blocks G and H include the term “dialup”.

Our accompanying poster provides more detail and analysis of this classification scheme, and validation of our classification on a random sample of blocks. As ongoing work we are evaluating clustering algorithms and block size accuracy.

5. REFERENCES

- [1] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley. AS relationships: Inference and validation. *ACM Computer Communication Review*, 37(1):29–40, Jan. 2007.
- [2] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *Proc. of the ACM SIGCOMM Conference*, pages 251–262, Cambridge, MA, USA, Sept. 1999. ACM.
- [3] V. Fuller, T. Li, J. Yu, and K. Varadhan. Classless inter-domain routing (CIDR): an address assignment and aggregation strategy. RFC 1519, Internet Request For Comments, Sept. 1993.
- [4] L. Gao. On inferring autonomous system relationships in the internet. *ACM/IEEE Transactions on Networking*, 9(6):733–745, Dec. 2001.
- [5] J. Heidemann, Y. Pradkin, R. Govindan, C. Papadopoulos, G. Bartlett, and J. Bannister. Census and survey of the visible internet (extended). Technical Report ISI-TR-2008-649, USC/Information Sciences Institute, February 2008.
- [6] G. Huston. IPv4 address report. <http://bgp.potaroo.net/ipv4/>, June 2006.
- [7] Internet Software Consortium. Internet Domain Survey, ISC_DS-2007JAN. web page <http://www.isc.org/ds>, Jan. 2007.
- [8] W. Mühlbauer, O. Maennel, S. Uhlig, A. Feldmann, and M. Roughan. Building an AS-topology model that captures route diversity. In *Proc. of the ACM SIGCOMM Conference*, pages 195–204, Sept. 2006.
- [9] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *Proc. of the IEEE Infocom*, pages 618–627, June 2002.
- [10] USC/LANDER project. Scrambled Internet Trace Measurement, PREDICT ID USC-LANDER/internet_address_survey_reprobing_it17w-20070601. web page <http://www.isi.edu/ant/lander>, June 2007.
- [11] Y. Xie, F. Yu, K. Achan, E. Gillum, M. Goldszmidt, and T. Wobber. How dynamic are IP addresses? In *Proc. of the ACM SIGCOMM Conference*, Kyoto, Japan, Aug. 2007. ACM.