An Analysis of Network Reachability Using BGP Data

Javier Salido
Electrical Engineering
University of Washington
Seattle, WA, 98195-2350
javiers@ee.washington.edu

Masanori Nakahara
Platform Technology Dev. Ctr.
Canon Inc.
Tokyo, Japan.
nakahara.masanori@canon.co.jp

Yinhai Wang
Civil Engineering
University of Washington.
Seattle, WA, 98195-2350
yinhai@u.washington.edu

Abstract

Though there have been extensive studies on AS and ISP level topologies, BGP path stability and overall path variations and churn, there is comparatively little data on the behavior of overall reachability in the Internet. We can explain reachability as a measure of path robustness over time, and thus as a significant measure of the general quality of the infrastructure. Our results show that BGP views across different collector route servers, through our period of study, are very homogeneous, and that variations in overall reachability are relatively small. We interpret these observations as indications that overall reachability is robust, and found it to be particularly good in G7 countries, perhaps due to a better communications infrastructure.
Keywords - BGP, availability, reachability

1. Introduction

Over the past several years, there have been a series of studies that have focused on analyzing different aspects of the Internet's topology, stability and reliability. These studies have done much to obtain an accurate representation of Internet topology at the Autonomous System (AS) [14,15], and Internet Service Provider (ISP) [16] levels, as well as to improve our understanding of some of the potential sources of instability and lack of availability [2,3,4] in it.

There is however, a much more modest and simple matter that is also relevant: reachability. It is our belief that there are some relevant questions that need to be answered. Questions like, how often do significant portions of the available IP address space become unreachable to a significant portion of Internet users? To what extent does origin AS or geographic location have an impact on reachability? Are there any other factors that affect reachability over the Internet?

These answers are relevant because they may improve our understanding of the potential sources of instability and availability problems on the Internet, and they may also be helpful to researchers working on protocol design at the transport and routing levels. In this paper we make an initial attempt at answering these questions, first by examining BGP tables from 3 different, well connected locations in the Internet: the University of Oregon (Route Views Project [10]), and the RIPE [11] servers in London and Amsterdam. Second, by examining how origin AS and geographic location relate to incidents during which reachability was affected.

We believe this data to be appropriate for this type of analysis, because the study of reachability does not require detailed knowledge of all the links in all possible paths to a certain prefix. Knowledge of the existence of a single path is enough, and a series of studies have shown these sources to be reliable in that sense. Chang et al established in their study [14] that while Route Views tables do not provide a full description of the number and distribution of existing links on the Internet, they do provide a fairly accurate count, and view, of the number of existing ASes. A number of other important studies have also used this data as the point of departure, to implement a *directed probing* approach to begin the discovery of ISP topologies [16], identify BGP misconfigurations [3], and analyze aspects of AS topology [15].

Our results show the information in the selected BGP tables to be very consistent, that overall reachability of the Internet is very stable over time, and that this consistency extends throughout the connected world, being particularly good in G7 countries. We concluded that overall reachability, during our study period, was more affected by individual incidents, than it was by spatial locality in general.

2. Background

2.1. BGP, Border Gateway Protocol

The Internet is a loosely organized international collaboration of interconnected networks. It can be divided into a large number of different regions of administrative control commonly called Autonomous Systems (ASes). At the boundary of each AS, peer border routers exchange reachability information to destination IP address blocks, or prefixes, for both transit networks and networks originating in that routing domain. Currently, most ASes exchange routing information using a path vector routing protocol called BGP (Border Gateway Protocol).

Path information exchanged by peer ASes is stored in BGP tables that compile all paths known by a particular AS, to existing available prefixes. Each AS then selects the best paths, using policies established by it based on business, technical and other criteria, and advertises them to neighboring ASes.

BGP plays an important role in Internet connectivity and not surprisingly it brings its own share of challenges. Different studies have identified some of these challenges; implementation may introduce excessive churn [4], and delay convergence [5], policy-interaction-caused persistent oscillations [6], and worm generated instability [7]. All these studies show that the reliability of BGP is crucial to the overall performance of the Internet. A simple BGP configuration error may disrupt Internet connectivity and cause reachability problems to large portions of the Internet.

2.2. Prefix Aggregation

Given the need to optimize the use of IPv4 addresses on the Internet, the old class system was replaced with CIDR (Classless Inter-Domain Routing). CIDR substitutes the old process of assigning class addresses with a generalized, variable, network "prefix". Instead of being limited to network identifiers of 8, 16 or 24 bits, CIDR uses prefixes of practically all sizes (in fact, in both the Route Views and RIPE tables we found prefixes of all sizes between 8 and 31 bits). For example, in the CIDR address 206.13.01.48/25, the "/25" indicates that the first 25 bits are used to identify a unique network, leaving the remaining bits to identify the specific host in that network. This scheme allows for address assignments that much more closely fit an organization's specific needs.

This solution however, may seriously increase the storage requirements at the routers. If a single AS has 20 network numbers with 27-bit long prefixes, for example, then every Internet backbone router needs 20 entries in its routing tables for that AS. To minimize the number of routing table entries, the CIDR addressing scheme also enables “hierarchical routing aggregation”, in which a single high-level route entry can represent many lower-level routes in the global routing tables. Such aggregation is critical to the survivability of the Internet’s routing system and BGP-4 takes advantage of it. ASes use BGP-4 to inform each peer AS of decisions it has made with respect to overlapping routes [12]. For example, if Router A learns about prefixes 128.95.0.0/17 and 128.95.1.0/17 from Router B, it can then decide to aggregate the two prefixes into a single 16-bit long prefix 128.95.0.0/16. Router A advertises the single, aggregated prefix 128.95.0.0/16 to its neighbors. By doing this, routing table sizes can be significantly reduced.

It is important however to note that, though aggregation works well to solve the routing table expansion problem, it introduces an additional complication on our reachability analysis because any given announcement or withdrawal update may cause reachability changes at other prefix levels. For example, an update message that withdraws prefix 128.95.0.0/16 does not necessarily mean that all networks under 128.95.0.0/16 become unreachable. Alternate routes to a more specific prefix (i.e. 128.95.204.0/24) might exist in the routing table, that were not announced before due to aggregation. Therefore, special attention is required when processing update messages from peer servers.

3. Methodology

3.1. BGP Data

There are some well known ASes that have “collector route servers”. That means that BGP tables that are built from the interaction between these collector route servers, and a large number of other peer ASes, are made available over the Internet as a collection of table snapshots, along with additional update files that provide the history of all updates received by these collector sites from their neighbors. Such is the case of the University of Oregon’s Route Views project [10] and RIPE (Réseaux IP Européens) servers [11].

The information extracted from BGP snapshots and regular updates is very useful to analyze Internet connectivity and, compared to active probing data, is easier to parse, process and comprehend [9].

It is important to point out that even though BGP tables show only the selected best paths, rather than all possible paths known to an AS, they provide enough information about prefix reachability since we are only

concerned about the availability of a path to a specific prefix, rather than the characteristics of that path.

Each of the collector route servers acts as a BGP listener with an AS number (e.g AS6447 for Route Views), that builds and stores a full snapshot of its own table every 2-8 hours. In addition to that, each collector site stores the updates that it has received from its neighboring ASes in between snapshots. Our study periods span from Jan. 10 through Jan. 16, 2002 (seven days) for the Oregon and Amsterdam sites, and from Jan. 1 through Jan. 16 for the London server. The London data set is used to investigate some discrepancies between it and the data from the other two.

Table 1. Descriptive statistics of each BGP collector route server

Server location	Route Views (Oregon)	RIPE (London)	RIPE (Amsterdam)
Number of neighbor ASes	20	40	12
Num. peers	23	55	13
Number of reachable ASes	12456	12239	12486
Number of reachable prefixes	112807	105159	114405

Table 1 shows the basic characteristics of each collector route server, obtained from a BGP snapshot on Jan 10. Over the length of our study period, the number of peers for each site remained unchanged. Two ASes (Verio. Inc and Tiscali Intl Network) contribute directly to all the three sites: Oregon, Amsterdam and London. Major Tier 1 ISPs and European/Asian ISPs contribute data to Oregon and Amsterdam while mostly UK focused ISPs connect to the London site. Some of these ASes provide several peering points for each site. Through its peer routers, each site can reach more than 110,000 prefixes attached to about 12,200 ASes.

The first step in this study was to establish a baseline for each site, by taking the first snapshot in our study period and looking for what we called "top level prefixes". To do this, all existing prefixes in each of the three tables were ordered sequentially, and a search was conducted to extract redundant prefixes. This process resulted in a table with a one-to-one relationship between reachable prefixes of any size, and the number of available paths to those prefixes. This is the table of top level prefixes.

Since subsequent update messages may affect these top-level prefixes and therefore the reachable IP ranges, (i.e. a update message may add or remove a top level prefix), these updates had to be processed in order to maintain our table of top level prefixes up to date at every step.

Figure 1 illustrates an example of this behavior. Suppose that the BGP table in AS100 tells us that this AS can reach prefix 128.95.0.0/16, using a path through AS150, as indicated by the solid lines. Suppose that the same table also tells us that AS100 can reach the more specific prefixes 128.95.200.0/24 and 128.95.235.0/24, through AS250, as indicated by the segmented arrows.

The "top level prefix" in our analysis will be the less specific prefix 128.95.200.0/16, indicating that this prefix is reachable from AS100. Now assume that an update message arrives from AS150, withdrawing prefix 128.95.0.0/16. This prefix will be withdrawn from our top level prefix table, but since we still have access through AS250 to the two more specific prefixes then these become "top level". The net result will reflect correctly which portion of the IP address space is still reachable from AS100.

This view, for each of the collector route servers, is what Chang et al [14] define as the *known address space.* For this study, the snapshots and updates from each of the three BGP views is analyzed, and an estimate of the size of the known address space is calculated with every update. This estimate is what we will later use to construct Figure 2 in section 4.

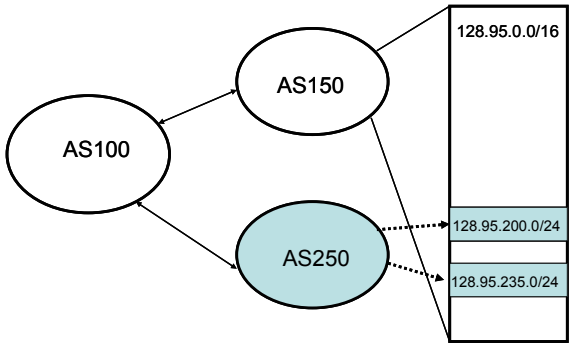

Figure 1. Top level prefixes.

Finally, we developed a computer program that took the information we downloaded and processed from the Route Views data, and then queried the NetGeo server from CAIDA in San Diego [13], for the geographic location of each of the prefixes that were determined to be unreachable during a portion the analyzed period of time. This is the information that is used to support our conclusions in section 4.3.

3.2. Generality of the Reachability Analysis

It is important to mention that reachability to and from any prefix may be seriously affected by local network conditions at a particular time. To avoid such bias, we use 7 day's worth of data from three BGP data collection sites located at different places (Oregon, London, and Amsterdam). By comparing results from analyses using data from these three sites, we expect to be able to draw some general conclusions.

We also felt that some clarification was needed given the way in which we liberally use the term "IP addresses" in the following sections. It is a known fact within the Internet community that though in theory there are 2,404,160,623 IP addresses, or 57% of the entire IPv4 space available for distribution as unicast addresses, only a small portion of these addresses, as is the case with telephone numbers, is actually in use. For comparison, around the same time of our study period, the Internet Domain Survey [17] estimated a lower bound of just over 162 million active virtual hosts in the Internet, and as of the time of writing of this paper, the estimate has gone up to 170 million.[1]

As we have mentioned, in this study we base our observations on the *known address space*, and use the term "*n* IP addresses are reachable" to refer to a range of IP addresses, or a group of prefixes. It is not our intention to imply that *n* computers are reachable at some prefix, since for our analysis we do not have any information that would allow us to make this assumption. Thus, the number of virtual hosts that is actually affected by the incidents that we refer to is considerably smaller than the numbers shown in our graphs.

4. Analysis

4.1. Known Address Space

Our analysis showed that our know address space comprises just over 27% of the total IPv4 address space. That is, approximately 27% of the total IPv4 address space is actually reachable through our top level prefix tables. Roughly speaking, that represents an average of 1.145 million individual unicast addresses that may be assigned to network nodes anywhere in the world, and to which we can find a workable path using BGP tables. Over the same period of time, we observed variations in this number as high as 10 million individual addresses.

[1] A discussion, and a proposed measure, on address assignment efficiency can be found in [18,19].

4.2. Temporal Locality Reachability Analysis

We began our comparison of the data from the different collectors by looking at variations in the known address space. Results are shown in Figure 2.

The *y* axis in 2.a and 2.b, represents the size of the known address space as seen from Route Views and Amsterdam (2.a), and from Route Views and London (2.b). The time axis has been shifted for the Amsterdam and London data, to account for time zone differences between them and Oregon, and so that the graphs can be more easily compared. The two graphs on the right represent the absolute value of the difference between Route Views and Amsterdam (2.c) and between Route Views and London (2.d), as a percentage of the total number of addresses in the known address space. We observe that there is barely any difference between the different collectors. Even in the case of London, that shows some visible differences in 2.b, as compared to Oregon, the total difference in the size of the known address space for each location is less than two hundredths of a percent, and averages much less. It is important to point out the fact that the difference between London and the other two sites in the original data was about four times larger, but we discovered that both Oregon and Amsterdam are advertising paths to prefix 39.0.0.0/8, which is a prefix listed by IANA as reserved. Thus, we subtracted the number of nodes covered by this prefix from all the Amsterdam and Oregon data points, to allow for an adequate comparison.

Even though the remaining differences between London and the other two locations are very small, we decided to analyze them. We found that they are distributed among at least 130 valid prefixes that show up in the Amsterdam and Oregon data, but do not show up in the London data. Practically all the missing prefixes have numbers above 80.0.0.0/8, with the overwhelming majority located on or above 192.0.0.0/8. We found no clear pattern or reason for this, a close examination of 30+ individual prefixes showed no correlation in geographical location (Europe, Asia, Australia, Americas, even one in the UK), type of AS (ISPs, NASA, SAP AG, Hewlett-Packard, a Mexican bank, etc.) or other. These prefixes also vary in size, but are not very large (the largest one been 150.229.0.0/19, an unidentified company in Australia). We found no explanation for this. To further complicate the matter, we can observe that both Oregon and Amsterdam experienced a sudden growth in their known address space around the second day, and that this growth accounts almost entirely for this discrepancy. We suspect that an analysis of the following week of London data would probably show the return of the missing prefixes to this table, but were not able to confirm this.

To expose any time pattern the reachability may follow, we also analyzed a week's worth of additional data for London, the previous week, and found that average reachability does not vary significantly in that period of time. An analysis of daily trends for each of the three sites revealed graphs that were mostly flat, with the same sudden spikes that can be observed in the weekly graph.

There are some instances in the graphs on Figure 2 where there are reductions in the size of the known address space, most of them lasting for short periods of time. However, the majority of the spikes we can see point upwards. They represent *increases* in the size of the known address space. Mahajan et al [3] show that most of the increases that last for a short period of time (less than a day), can be attributed to possible misconfiguration errors in BGP tables, which can be introduced mistakenly by the administrator of the affected AS.

Our study period was in fact selected to coincide with the last of the three weeks of [3], so that we could compare our observations to their results. According to [3], they identified and confirmed 123 individual incidents that affected up to 270 prefixes of varying sizes, over a period of three weeks, and this represented only a small portion of the total number of observed possible incidents.

4.3. Spacial Locality Reachability Analysis

Figure 3, shows plots (3.a, 3.b and 3.c) of all recorded instances in which an individual prefix became unreachable, and for how long, at each of the three collector sites. The graphs on the right (3.d, 3.e and 3.f) show the number of oscillations that account for the total unreachable time for each prefix. As in the case of the graphs in Figure 2, we can see that there is considerable similarity in the results from each site.

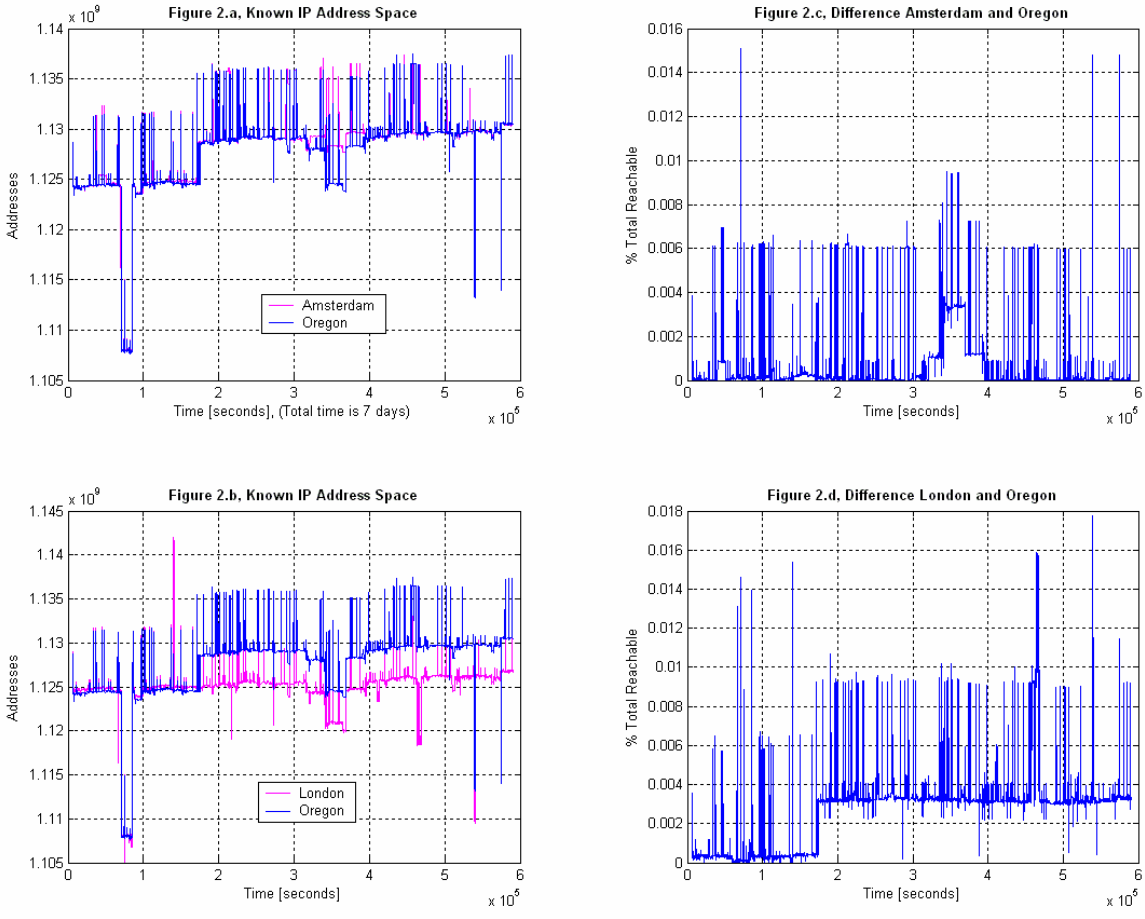

Figure 2, Size of the *known address space* at all three sites, and their differences.

The wide empty spaces show those portions of the IP range that are not assigned for use, and we thought that an interesting feature in these graphs is the fact that the density of the points increases as we move from left to right over the IP address range. A look at the assignment of IP ranges in the IANA web site [22] reveals that most of the address ranges for which the first octet is less than 50, are assigned to U.S. government entities, Universities, research centers, and large private corporations that integrate a small group of well connected ASes

As we move to the right of the range however, the number of ASes that have assigned prefixes increases, and their relative size decreases. The allocation of prefixes in this area however, particularly in the right most cluster of points above 190.0.0.0/8, can not be directly associated with any geographic location, since it is distributed all over the world. Therefore, a more detailed analysis of this data was required for the following step.

As previously mentioned, we used the data points from Figure 3 to query the NetGeo server in San Diego and obtain the geographic location of the prefix involved in each of these incidents.

In order to avoid bias introduced by single, very large incidents, we first looked for specific cases in which a single AS failure could account for a large number of unreachable prefixes for a large portion of time. We found 5 such cases, accounting for 1614 unreachable prefixes, in which a single AS was responsible for over 150 prefixes being unreachable, all the ASes involved were ISPs (2 in the U.S., one in Canada, one in Japan and one in Taiwan).

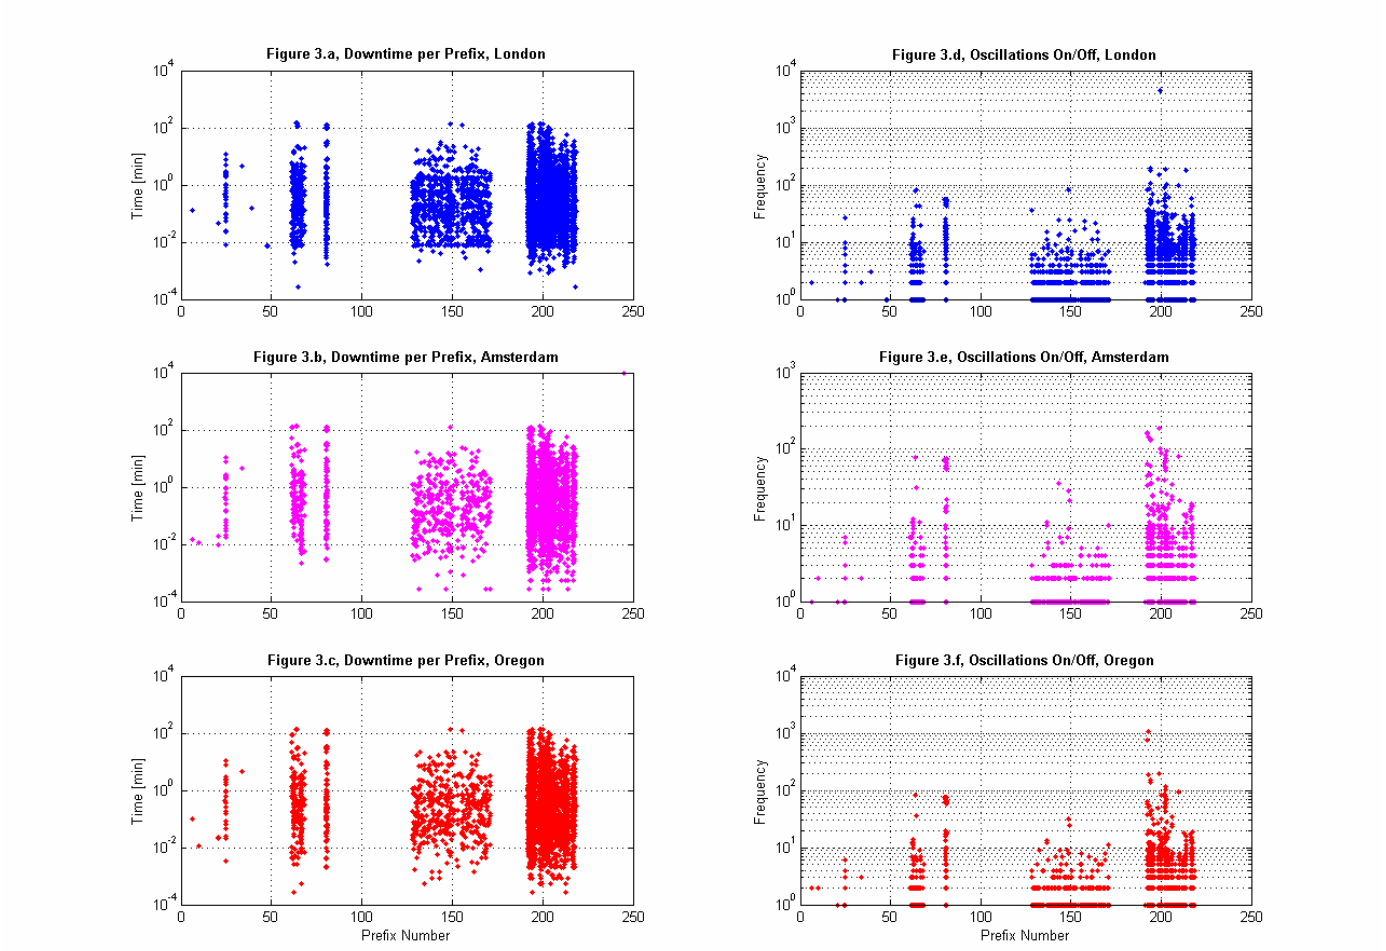

Figure 3, Observed incidents in which individual prefixes became unreachable, and their oscillations. The *X* axis represents prefix number in all 4 figures.

One of these cases was responsible for 954 prefixes with an average prefix size of 20 bits, becoming unreachable for an average of over three days. Finally, we also removed all the incidents in which a prefix located in the Unites States was unreachable, 1277 prefixes. Having done this, we were left with 3539 data points involving prefixes that had become unreachable at some point in time during our study period.

Our initial analysis did not show correlation between the unreachable prefixes and geographic location. Table 2 shows the results for greatest average downtime, and greatest number of incidents. We can see that the only country that shows up in both groups is Brazil. In fact, we found that there is very little correlation between average downtime and number of incidents, between average downtime and average oscillations, or between number of incidents and oscillations. The average size of the involved prefixes was very similar in all these incidents, close to 20 bits.

We made a more specific comparison of reachability between countries with different GNP levels. For this, we separated the G7 countries (the U.S., Japan, United Kingdom, Germany, France, Canada, and Italy) from all the remaining countries, and analyzed if the difference in reachability is significant between the two groups. Table 3 shows the statistics of the analysis.

Table 2. Countries with most incidents in which a prefix became unreachable (left), and countries with largest avgerage downtime per incident (right).

Incidents		Average Downtime Per Incident [min]	
Australia	322	Monaco	124.81
Mexico	295	Croatia	98.58
Rumania	283	Jordan	36.30
Canada	199	Macedonia	31.77
Taiwan	143	Netherlands	29.45
Brazil	142	Brazil	14.63
India	137	Zaire	14.21
South Africa	105	Iran	13.74
Germany	85	Czech Rep.	13.62
Russia	85	Thailand	12.23
Pakistan	81	U.K.	12.16

Table 3 shows that the mean down time (unreachable time), for G7 countries is only about 39% of that for the remaining countries. This difference is significant at 0.01 level as indicated by the t-ratio. Very significantly, we compared number of route oscillations and number of hops in a route between the two groups. We find that the mean of oscillations for G7 countries is only 5.465, much less than 11.703, the mean of oscillations for the remaining countries.

We also found that the average path length to reach a prefix in G7 countries (6.206) is about 0.7 hops shorter than that of the remaining countries (6.882).

This may be an indirect reflection of the difference in ISP size between the two country groups. Both analyses are significant at 0.01 levels.

Table 3. Down Time comparison between G7 countries and the rest of the world. All units are in seconds.

	G7 Countries	Other Countries
Number of cases	2365	4064
Mean	29386.30	75294.40
Standard deviation	82558.74	144446.49
Difference in mean	45908.10	
t-ratio	16.22	
Significance level	0.000	

4.4. Individual Incidents

While our study focuses on general trends rather than an individual analysis of reachability incidents, it is important to point out the fact that we observed several major incidents affecting ISPs and Telecommunications companies that must have had considerable impact.

The incident involving over 900 individual prefixes that is described in section 4.3. heavily biased our initial results. We could not explain why the Netherlands, a very well connected country, consistently showed up in all our negative categories until we discovered that a North American ISP was in fact having problems that affected up to 5 different ASes assigned to it and covered 6 different countries, including Holland.

Monaco shows up at the top of our average downtime per incident table, yet its appearance involves a single incident involving 3 different prefixes 21, 22 and 23 bits in size.

Single incidents play an important role in reachability, and studies like [3] can help ISPs, Telecom companies and equipment manufacturers to improve it.

5. Discussion

Our analyses have exposed reachability changes both temporally and spatially. As aforementioned, reachability is a comprehensive reflection of working condition and policies for the entire network. Any components that are related to physical connections and routing processes can impact prefix reachability. In this section, we discuss factors that may disrupt prefix reachability.

5.1. Software Effects

Routing software at each switching node must support all the routing related functions, such as communication, route calculation, and packet forwarding. Design options and implementation bugs can seriously affect routing performance. Labovitz et al [4] described that after a router maker improved its router operating system software, observed update messages dropped abruptly from two million to below ten thousand pathological withdrawals per day.

5.2. Policy Effects

BGP assumes that Internet is an arbitrarily interconnected set of ASes. Hence, it allows each AS to independently formulate its own routing policies, and it allows these to override distance metrics in favor of policy concerns. BGP regards issues such as, which routes to accept from a neighbor and the preference with which those routes should be treated, as a local decision based on its routing policy. An important part of this routing policy is to decide which set of paths should be advertised to each BGP neighbor.

The decision on which routes to accept from and advertise to various BGP neighbors, has a profound impact on what traffic crosses a network [12], and hence affects reachability of prefixes. Varadhan et al [20] found that some routing policies are unsafe in the sense that they may cause BGP to diverge. BGP policy can also be actively used to reduce instability. For example, Cisco IOS version 11.0 introduced "bgp dampening" command to minimize the instability caused by route flapping and oscillation over the network. To accomplish this, criteria are defined to identify poorly behaved routes, and take consequent action. A route that is flapping receives a penalty of 1000 for each flap. When the accumulated penalty reaches a configurable limit, BGP suppresses advertisements of the route even if the route is up. Therefore, routing policy plays an important role on reachability determination.

5.3. Human Mistakes

Human introduced errors in routing system software or configuration file can be disastrous for prefix reachability. To use the aforementioned AS7007 example to show how serious the damage can be caused by a simple routing misconfiguration. In April 1997, a small ISP in Florida, AS7007, made a mistake in configuring the router that joined its small network to Sprint. It allowed all the routes it learned from Sprint using BGP to be exported back to Sprint as its own routes. If the Sprint BGP speaker had done the policy-based filtering properly, loops should have been detected and filtered out. Unfortunately, the Sprint BGP speaker wasn't filtering properly either and began sending out updates that added AS7007 as the correct route for a portion of every CIDR block. This misinformation spread through Sprint's network, and further propagated into neighboring NSPs, including ANS, MCI, UUNet, and others. Many routers crashed because their routing tables suddenly doubled in size (an additional route was added for each CIDR block), and the routing instability spread throughout the Internet. The crashed routers dropped their BGP connections with their peers and made the networks they spoke for, unreachable.

6. Conclusions

Our main findings from both the temporal and spatial analyses are:

BGP table data across all three different sites are remarkably consistent, with one significant exception, the case of slightly smaller total number of reachable addresses in the London data set.

The size of the *known address space* was remarkably stable. Most of the observed variations are sudden, short lived increases that have been explained by Mahajan et al in [3], as BGP table misconfigurations.

While overall trends are very stable, it is important to point out that individual incidents were responsible for a good deal of the major observed reductions in reachability. While small in percentage, these incidents are significant in the sense that they affected mostly ISPs and Telecom companies.

Human induced errors play a significant role in reachability changes, both in terms of BGP miscon-

figurations [3] and, we suspect, in the most relevant incidents in which reachability was seriously reduced during our period of study.

There is a significant difference, in every measurement we observed (oscillations, average downtime per incident, and route length), in reachability between G7 countries and the rest of the world.

In some cases, perhaps deliberately or perhaps due to human error, large portions of the reserved IP space can be found advertised as available, as was the case with the prefix 39.0.0.0/8, that covers over 16 million addresses in the known address space.

7. Future Work

This analysis covers a brief period, one week, of all the data available from Route Views and RIPE. The same analysis over a longer period of time might expose monthly and yearly trends in temporal locality.

While it is tempting to dismiss individual incidents of lost reachability as infrastructure or administrative problems in ISPs, a study of their real causes and longer term impact might show some surprises, as was the case with [3]. However, it is difficult for us to ases the feasibility of such a study at this point in time, given the fact that it would have to be conducted with some level of cooperation from ISPs.

8. References

[1] V. Fuller, T. Li, J. Yu, and K. Varadhan. Classless "Inter Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", IETF, *RFC 1519*, Sept. 1993.

[2] C. Labovitz, G. R. Malan, and F. Jahanian. "Internet Routing Instability.", *ACM SIGCOMM*, September 1997.

[3] R. Mahajan, D. Wetherall, and T. Anderson. "Understanding BGP Misconfiguration.", *ACM SIGCOMM,* August 2002.

[4] C. Labovitz, G. R. Malan, and F. Jahanian. "Origins of Internet Routing Instability.", *IEEE INFOCOM*, June 1999.

[5] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. "Delayed Internet Routing Convergence*.", ACM SIGCOMM*, September 2000.

[6] T. Griffin, and G.T. Wilfong. "An Analysis of BGP Convergence Properties.", *ACM SIGCOMM*, August 1999.

[7] J. Cowie, A. Ogielski, B. Premore, and Y. Yuan. "Global Routing Instabilities during Code Red II and Nimda Worm Propagation.", September 2001. http://www.cc.gatech.edu/~ncb/bgp-storm.pdf

[8] Y. Rekhter, and T. Li. A "Border Gateway Protocol 4 (BGP-4)", IETF, *RFC 1771*, 1995.

[9] A. Broido, and K. Claffy. "Internet Topology: Connectivity of IP Graphs.", *San Diego Proceedings of SPIE International symposium on Convergence of IT and Communication*. Denver, CO. 2001.

[10] University of Oregon, Route Views Project. http://www.antc.uoregon.edu/route-views/

[11] Réseaux IP Européens (RIPE) Network Coordination Center. http://www.ripe.net/

[12] J. W. Stewart III. *BGP4: Inter-Domain Routing in the Internet.* The Addison-Wesley Networking Basics Series., New Jersey, December 1998.

[13] NetGeo Project, Cooperative Association for Internet Data Analysis (CAIDA). http://www.caida.org/tools/utilities/netgeo/

[14] H. Chang, R. Govindan, S. Jamin, S.J. Shenker and W. Willinger. "Towards Capturing Representative AS-Level Internet Topologies.", *SIGMETRICS*, November 2002.

[15] D. Andersen, N. Feamster, S. Bauer, and H. Balakrishnan. "Topology Inference from BGP Routing Dynamics", *Internet Measurement Workshop,* November 2002.

[16] N. Spring, R. Mahajan, and D. Wetherall. "Measuring ISP Topologies with Rocketfuel.", *ACM SIGCOMM*, August 2002.

[17] Internet Domain Survey, Internet Software Consortium (ISC), Jan 2003. http://www.isc.org/ds/

[18] C. Huitema. "The H Ratio for Address Assignment Efficiency.", IETF, *RFC 1715*, Nov. 1994.

[19] C. Huitema and A. Durand. "The Host Density Ratio for Address Assignment Efficiency: An Update on the H Ratio", IETF, *RFC 3194*, Nov. 2001.

[20] K. Varadhan, R. Govindan and D. Estrin. "Persistent Route Oscillations in Inter-Domain Routing.", *ISI Technical Report* pp. 96-631, USC/Information Sciences Institute, 1996.

[22] "IPv4 Address Space", Internet Assigned Numbers Authority (IANA), April 2003. http://www.iana.org/assignments/ipv4-address-space