Experimental Study of Internet Stability and Backbone Failures *

Craig Labovitz, Abha Ahuja, Farnam Jahanian
University of Michigan
Department of Electrical Engineering and Computer Science
1301 Beal Ave.
Ann Arbor, Michigan 48109-2122
{labovit, ahuja, farnam}@umich.edu

Abstract

In this paper, we describe an experimental study of Internet topological stability and the origins of failure in Internet protocol backbones. The stability of end-to-end Internet paths is dependent both on the underlying telecommunication switching system, as well as the higher level software and hardware components specific to the Internet's packet-switched forwarding and routing architecture. Although a number of earlier studies have examined failures in the public telecommunication system, little attention has been given to the characterization of Internet stability. We provide analysis of the stability of major paths between Internet Service Providers based on the experimental instrumentation of key portions of the Internet infrastructure. We describe unexpectedly high levels of path fluctuation and an aggregate low mean time between failures for individual Internet paths. We also provide a case study of the network failures observed in a large regional Internet backbone. We characterize the type, origin, frequency and duration of these failures.

1. Introduction

In a brief number of years, the Internet has evolved from a relatively obscure, experimental research and academic network to a commodity, mission-critical component of the public telecommunication infrastructure. Internet backbone failures that previously only impacted a handful of academic researchers and computer scientists, may now as easily generate millions of dollars of losses in e-commerce revenue and interrupt the daily routine of hundreds of thousands of end-users

*Supported by National Science Foundation Grant NCR-971017, and gifts from both Intel and Hewlett Packard.

The computer engineering literature contains a large body of work on both computer fault analysis, and the analysis of failures in the Public Switched Telephone Network (PSTN) [17, 1, 8]. Studies including [6, 18] have examined call blocking and call failure rates for both telephony and circuit switched data networks. Although a number of researchers have applied graph theoretic approaches to the study of faults in simulated, or theoretical networks [2], the topological stability and dynamics of deployed wide-area Internet Protocol (IP) backbones has gone virtually without formal study, with the exception of [9, 4, 3, 15].

In this paper, we describe an experimental study of Internet stability and the origins of failure in Internet protocol backbones. Unlike telephony networks, the stability of end-to-end Internet paths is dependent both on the underlying telecommunication switching system, as well as the higher level software and hardware components specific to the Internet's packet-switched forwarding, name resolution and routing architecture. Although a number of vendors provide mean-time to failure statistics for specific hardware components used in the construction of wide-area networks (e.g. power supplies, switches, etc.), estimations of the failure rates for IP backbones at a systemic level remain problematic.

The Internet exhibits a number of engineering and operational challenges distinct from those associated with telephony networks and applications. Most significantly, unlike switched telephony networks, the Internet is a conglomeration of thousands of heterogeneous dynamically packet switched IP backbones. Internet hosts segment application level streams into one or more independently routed IP datagrams. At the edge of every Internet backbone, routers forward these datagrams to the appropriate next-hop router in adjacent networks. Internet routers build next-hop routing tables based topological information exchanged in con-

trol messages with other routers.

The most common inter-domain (exterior) routing protocol used between Internet providers is the Border Gateway Protocol (*BGP*) [5]. BGP route information includes a record of the inter-domain path the route has followed through different providers. We refer to this path record of as the route's *ASPath*.

Backbone service providers participating in the Internet core must maintain a complete map, or "*default-free*" routing table, of all globally visible network-layer addresses reachable throughout the Internet. At the boundary of each Internet Service Provider *(ISP)* backbone, peer border routers exchange reachability information to destination IP address blocks, or *prefixes*. A prefix may represent a single network, or a number of customer network addresses grouped into one larger, "supernet" advertisement. Providers commonly aggregate large numbers of customer networks into a single supernet announcement at their borders.

A number of studies, including [12, 15], have examined the stability of both Internet end-to-end paths and end-systems. We approach the analysis from a complimentary direction – by analyzing the internal routing information that gives rise to all end-to-end paths. Our study of the "default-free" routing information from the major Internet provides analysis of a supserset of all end-to-end Internet paths. For example, a single /8 route described in Section 3 may describe the availability of more than 16 million Internet end-systems. Our measurement infrastructure also allows the observation of higher frequency failures than described in [12, 15]. Overall, the significant findings of our work include:

- The Internet backbone infrastructure exhibit significantly less availability and a lower mean-time to failure than the Public Switched Telephone Network (PSTN).
- The majority of Internet backbone paths exhibit a mean-time to failure of 25 days or less, and a mean-time to repair of twenty minutes or less. Internet backbones are rerouted (either due to failure or policy changes) on the average of once every three days or less.
- Routing instability inside of an autonomous network does not exhibit the same daily and weekly cyclic trends as previously reported for routing between Inter provider backbones, suggesting that most inter-provider path failures stem from congestion collapse.
- A small fraction of network paths in the Internet contribute disproportionately to the number of long-term outages and backbone unavailability.

The remainder of this paper is organized as follows: Section 2 describes the infrastructure used in our characterization of backbone failures and the analysis of both inter and intra-domain path stability. Section 3 includes our analysis of the rate of failure and repair for both inter-domain Internet paths and intra-domain routes from a case study of a regional network. We also categorize the origins of failures during a one year study of this regional network. Finally, we compare the frequency and temporal properties of BGP and intra-domain routing data.

2. Methodology

Our analysis in this paper focuses on two categories of Internet failures: faults in the connections between service provider backbones, and failures occurring within provider backbones. Our data is based both on experimental measurements of deployed wide-area networks and data obtained from the operational records of a large regional Internet service provider. We use a number of tools developed by the MRT [14] and IPMA [7] projects for the collection, analysis and post-processing of our data.

We base our analysis of failures between service providers on data recorded by a central route collection probe, named RouteViews, located on the University of Michigan campus. We configured RouteViews to participate in remote BGP peering sessions with a number of cooperating regional and national backbone providers. Each of these backbone routers provided RouteViews with a continuous stream of BGP updates on the current state of the provider's default-free routing table between January 1997 and November 1998.

We base our analysis of intra-domain failures on a case study of a medium size regional network. The regional backbone interconnects educational and commercial customers in 132 cities via high speed serial lines and frame-relay links at speeds up to OC3 (155 MB). The network includes 33 backbone routers connected via multiple paths with links to several hundred customer routers. We use both recorded routing data and failure logs from this provider to categorize the type and frequency of different sources of failure.

We use a single provider case study due to the significant challenges of a more complete survey of internal failures across multiple providers. Factors limiting a more complete survey include the scale of the Internet, difficulties in the correlation of failure data amongst providers with different backbone infrastructure and fault monitoring practices, and the highly proprietary nature with which most provider's regard their failure data. As Paxson observed in [16], no single backbone,

or snapshot of the Internet provides a valid representation of the heterogeneous and rapidly changing Internet. As a result, we do not claim our case study is representative of all providers. Instead, our focus in this paper is on comparing a source of intra-domain failure data with faults observed in the connectivity between providers.

For our intra-domain analysis, we first study the frequency and duration of failures using the operational monitoring logs from our case study provider. The monitoring system used by this provider includes a centralized network management station (CNMS) which periodically monitors all of router interfaces throughout the network using SNMP queries and the transmission/receipt of "ping" packets. We base our analysis on twelve months of CNMS logs from November 1997 to November 1998.

Our characterization of network failures used data culled from the trouble ticket tracking system managed by our case study provider's Network Operations Center (NOC). The NOC staff uses the trouble ticket information for tracking, troubleshooting and coordinating the resolution of detected network failures. During the course of normal operations, network operations staff manually create trouble tickets upon either the automated detection of a fault by the CNMS, or upon receipt of customer complaints.

3. Analysis

We divide our analysis in this section into three areas. We first examine the frequency and duration of failures observed in inter-provider backbone paths. Repeating the standard method of analysis used in computer systems, we examine the availability, mean-time to failure, and mean-time to repair for Internet routes. In the second subsection of our analysis, we explore the source, frequency and duration of internal backbone failures using the failure logs and routing data from our case-study provider. Finally, we discuss the relationship between the frequency of intra-domain failures and the behavior of inter-domain routing changes.

3.1. Analysis of Inter-domain Path Stability

In this section, we first turn our attention to failures observed in the inter-domain routing paths exchanged between core backbone providers. Specifically, we examine nine months of default-free BGP routing information recorded from three remote Internet Service Provider (*ISP*) backbone routers (ISP1, ISP2, ISP3). As noted in Section 2, the three providers represent a spectrum of different ISP sizes, network architecture and underlying transmission technology.

Our logs of routings updates from the three ISP routers provide BGP transition information about both the provider's own customer and transit routes, as well as routes received from other ISPs.

In our analysis, we examine the routing activity of each ISP independently. By this, we mean that if an ISP lacks a route to a given prefix destination, we consider that destination unreachable from that ISP even if other providers maintain a route to that destination. We define an inter-domain *fault* as the loss of an ISP's route to a previously reachable prefix.

In the taxonomy below, we distinguish between three classes of BGP routing table events observed from each provider:

Route Failure: A route is explicitly withdrawn and no alternative path to the prefix destination, or to a less specific aggregate network address, is available.

Route Repair: A previously failed route to a network prefix is announced as reachable. This also may include the addition of new customer routes, or the announcement of secondary, backup paths due to policy or network failures.

Route Fail-Over: A route is implicitly withdrawn and replaced by an alternative route with differing next-hop or ASPath attributes to the prefix destination. Route Fail-over represents the rerouting of traffic to a given prefix destination after a network failure. Recall from Section 1 that the ASPath represents the routing path of the prefix through different inter-connected autonomous systems.

Inter-domain Route Failures generally reflect faults in the connectivity between providers, or the internal loss of a provider's connectivity to multiple customer routers. Lacking internal knowledge of the policies and design of provider backbones, we cannot always distinguish between "legitimate" network failures, and certain classes of policy changes, consolidation amongst provider networks, or the migration of customers between providers.

We first look at the availability of inter-domain routes. We define the *availability* of a given default-free route from a provider as the period of time that a path to the network destination, or a less specific prefix, was present in the provider's routing table. We include less specific prefixes in our definition since as described in Section 1, provider's regularly aggregate multiple more

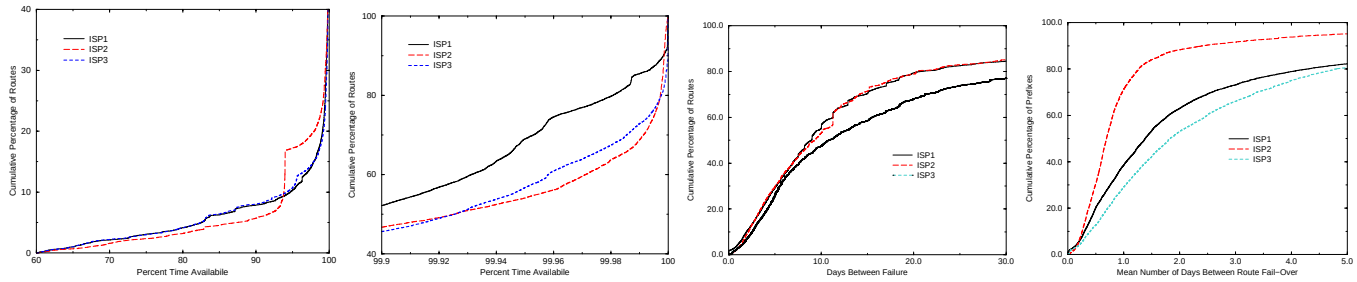

(a) Availability (b) Availability (c) Failure (d) Fail-Over

Figure 1. Cumulative distribution of the route availability of three service providers.

specific network addresses into a single supernet advertisement. We make several modifications to our data, described in [11], to more accurately reflect outages.

The graphs in Figure 1(a)(b) show the cumulative percentage of time default-free routes were available from each provider during our ten month study. The horizontal axis shows the percent time available; the vertical shows the cumulative percentage of routes with such availability. Both graphs in Figure 1(a)(b) represent the same data, but Figure 1(b) provides an expanded view of route availability above 99.9 percent.

A recent study study [8] found that the PSTN averaged an availability rate better than 99.999 percent during a one year period. From the graph in Figure 1(b), we see that the majority of Internet routes (65 percent) from all three providers exhibited an order of magnitude less availability. Only between 30 and 35 percent of routes from ISP3 and ISP2, and 25 percent of routes from ISP1 had availability higher that 99.99 percent of study period. Further, a significant 10 percent of the routes from all three providers exhibited under 95 percent availability. The availability of the three providers exhibit similar curves for most of Figure 1(a). The step in the curve for ISP3 at 95 percent availability represents a multi-hour loss of inter-provider connectivity due to an outage described in [11]. ISP1 exhibits significant less availability above 99.9 than ISP2 and ISP3 as evinced by the higher curve in Figure 1(b).

In addition to availability, we examine the rate of failure and fail-over in inter-domain paths. We define an inter-domain route *failure* as the loss of a previously available routing table path to a given network, or a less specific, prefix destination. A *fail-over* of a route represents a change in the inter-domain path (ASPath or NextHop) reachability of that route.

The two graphs in Figure 1 show the cumulative distribution of the mean number of days between route failures (c), and route fail-over (d) for routes from ISP1, ISP2 and ISP3. The horizontal axes represent the mean number of days between failures/fail-over; the vertical axes show the cumulative proportion of the ISP's routing table entries for all such events. Examining the graph in Figure 1(c), we see that the majority of routes (greater than 50 percent) from all three providers exhibit a mean-time to failure of fifteen days or more. By the end of thirty days, the majority (75 percent) of routes from all three providers had failed at least once. The distribution graphs for ISP1, ISP2 and ISP2 share a similar curve, with ISP1 exhibiting a slightly lower cumulative MTTF curve starting at ten days.

As noted earlier, most Internet providers maintain multiple, redundant connections to other providers. In the case of a single link or provider failure, routers will dynamically reroute around faults. Since not all Internet routes enjoy redundant connectivity, we focus our analysis on fail-over by modifying the vertical axis in Figure 1(d) to reflect a cumulative subset of interdomain routes – only those routes that exhibit multiple paths. Examining this graph, we see that majority of routes with redundant paths fail-over within two days. Further, only 20 percent of these routes from ISP1 and ISP3, and five percent from ISP2 do not fail over within five days. Both these mean-time to failure and fail-over results suggest a slightly higher incidence of failure in today's Internet than described in Paxson's 1994 study [15] which found 2/3's of Internet paths persisted for either days or weeks.

The graph in Figure 2(a) shows the cumulative distribution of the mean number of minutes between a route failure and repair. The horizontal axis shows the average time a route was unavailable; the vertical shows the cumulative percentage of all routes experiencing

such an event. Since default-free routes announced by each ISP include routes transiting other providers, the mean-time to repair reflects both the time for fault resolution as well as the propagation delay of routing information through the Internet.

From Figure 2(a), we see that 40 percent of failures are repaired in under ten minutes. The majority (60 percent) are resolved within a half hour. After thirty minutes, the cumulative MTTR curves for all three providers demonstrates a heavy-tailed distribution, with slow asymptotic growth towards 100 percent. We can see the relationship between availability, MTTF and MTTR by examining the data for ISP1. The MTTF curve for ISP1 rose faster than ISP2 and ISP3 in Figure 1(c), but at a slower rate in the Figure 2(a) MTTR graph. The lower average mean-time to failure, but slower mean-time to repair contributes to ISP1's overall lower availability in Figure 1(a).

Overall, analysis of our MTTR data agrees with our qualitative findings in Section 3.2 that repairs not resolved within an hour usually represent more serious outages requiring significant engineering effort for problem diagnosis, or the replacement of faulty hardware. Our data also corroborates Paxson's findings [15] that most Internet outages are short-lived – lasting on the order seconds or minutes.

The above mean-time to repair data provides an indication of the average unavailability of a route, but it does not provide insight into the overall distribution of outage durations. In Figure 2(b) we show the cumulative distribution of outage durations for all three providers. The horizontal axis represents the duration of outages in hours on a logarithmic scale; the vertical axis represents the cumulative percentage of outages lasting the given duration or less. During the course of our study, we observed over six million outages. From Figure 2(b), we see that only 25 to 35 percent of outages from the three providers are repaired in under an hour. This data is in marked contrast to Figure 2(a) where the average repair time for a route failure is under a half hour. Analysis of the relationship between our failure duration data with the graph Figure 2(a) indicates that a small number of routes disproportionately contribute to overall unavailability. Or, more specifically, forty percent of routes exhibit multiple failures lasting between one hour and several days during our study. This result agrees with our findings in [9] that a small fraction routes are responsible for the majority of network instability.

3.2. Analysis of Intra-Domain Network Stability

In the last section, we examined the stability of inter-domain paths. We now focus on intra-domain failures using a case study of a regional provide described in Section 1. Intra-domain routing serves as the basis for much of the information exchanged in inter-domain routing and analysis of the faults associated with an intra-domain network also provides insight into failures in other areas of the Internet.

The graph in figure 2(c) shows the cumulative distribution of the mean-time to failure for two categories of router interfaces: backbone nodes and customer-sites. The horizontal axis represents the mean-time between interface failures; the vertical axis shows the cumulative percentage of interface failures at each mean-time. We define *backbone nodes* as router interfaces connected to other backbone routers via multiple physical paths. *Customer connections* represent router interfaces attached to the regional backbone via a single physical connection. As critical elements of the network infrastructure, backbone routers are closely monitored, and housed in telco-grade facilities with redundant power. In contrast, routers at customers nodes often are maintained under less ideal physical conditions and administration.

From Figure 2(c), we see that 40 percent of all interfaces experienced some failure within an average of 40 days, and five percent failed within a mean time of five days. Overall, the majority of interfaces (more than 50 percent) exhibit a mean-time to failure of forty days or more. This differs from our earlier analysis of BGP paths, which found the majority of inter-domain failures occur within 30 days. The curve of the better equipped and management backbone interfaces exhibits significantly lower MTTF than customer routers.

The step discontinuities in Figure 2(c) represent both the relationship between interfaces and an artifact of our data collection architecture. Specifically, interface failures tend to occur in groups due to power, maintenance and related outages simultaneously affecting all interfaces on a router. In addition, rare simultaneous failures of multiple redundant paths through the network may lead to a network partition and a disconnect between multiple router interfaces and the central data collection host.

The graph in Figure 2(d) shows the cumulative mean-time to repair for the two different categories of router interfaces described earlier. The horizontal axis shows the mean number of minutes to repair; the vertical shows the cumulative percentage of all interfaces averaging such repair duration. From the graph, we see that 80 percent of all failures are resolved in un-

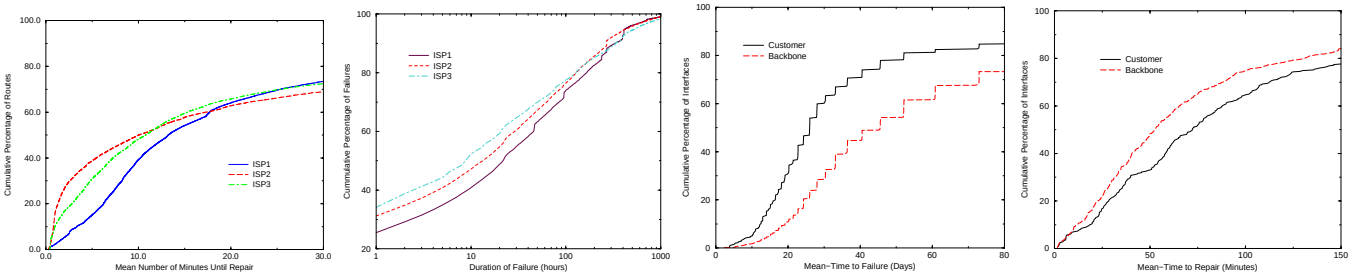

(a) Inter-domain MTTR (b) Inter-domain Duration (c) Intra-domain MTTF (d) Intra-domain MTTR

Figure 2. Mean time to repair and mean failure duration for inter and intra-domain routes.

der two hours. Further analysis of the data indicates that outages lasting longer than two hours usually represent long-term (several hours) outages which require significant engineering effort for problem diagnosis or the replacement of hardware or circuits.

3.3. Network Failures

In this section, we categorize the origins of the hardware, software and operational faults that gave rise to the intra and inter-domain failures described in the previous two subsections. As discussed in Section 2, we base our characterization of network failures on the operational trouble logs of a regional ISP.

Figure 3(a) shows a breakdown of all the outages recorded during our one-year case study (November 1997 to November 1998). As the diagnosis and categorization of outages remains an inexact science, several of the categories overlap and a few include some degree of ambiguity. The largest category at 16.2 percent, maintenance, refers to either a scheduled, or unscheduled emergency upgrade of software or hardware, or router configuration changes. A power outage (16 percent) includes either loss of power to a router, or a power failure in a PSTN facility which impacts one or more ISP circuits. Fiber or carrier failures (15.3 percent) usually result from a severed fiber optics link or a PSTN facility problem. Unreachable includes intermittent failures which mysteriously resolve themselves before an engineer investigates the outages. These unreachable outages usually result from PSTN maintenance or failures. A hardware problem (9 percent) includes a router, switch or power supply failure. Congestion refers to sluggishness, or poor connectivity between sites and usually represents link/router congestion on links, or router software configuration errors. A routing problem designation reflects errors with the configuration or interaction of routing protocols (OSPF, BGP, RIP). Most routing problems stem from human error and misconfiguration of equipment. Finally, the software problem category includes router software bugs.

From Figure 3, we see that majority of outages stem from maintenance, power outages and PSTN failures. Specifically, over 15 percent of all outages were due to sources outside of the provider's immediate control, including carrier and frame-relay failures. These percentages reiterate the observation in Section 1 that the reliability of IP backbones shares a significant dependence with the reliability of the underlying PSTN infrastructure. Approximately 16 percent of the outages were due to power outages. Power failures generally affect only customer routers which lack the same redundant power supplies as housed in backbone router facilities. Another 16 percent of the outages were planned maintenance outages. Overall, we note that most of these observed outages were not specifically related to regional IP backbone infrastructure (e.g. routers and software).

Further analysis of the data represented in Figure 3(a) shows the majority of outages were associated with individual customer sites rather than backbone nodes. This result is somewhat intuitive as backbone nodes tend to have backup power (UPS), more experienced engineers and controlled maintenance and upgrades.

Figure 3(b) shows number of interfaces, minutes down, and average number of interface failures for each backbone router monitored during our case study. From the table, we see that the overall uptime for all backbone routers averaged above 99.0 percent for the year. Further analysis of the raw data shows that these

Outage Category	Number of Occurrences	Percentage
Maintenance	272	16.2
Power Outage	270	16
Fiber Cut/Circuit/Carrier Problem	261	15.3
Unreachable	215	12.6
Hardware problem	154	9
Interface down	105	6.2
Routing Problems	104	6.1
Miscellaneous	86	5.9
Unknown/Undetermined/No problem	32	5.6
Congestion/Sluggish	65	4.6
Malicious Attack	26	1.5
Software problem	23	1.3

(a) Failure Categories

Router Name	# Interfaces	Percent Time Available	Average Number of Interface Failures	Average Minutes Down per Interface
bspop	14	99.74	17.79	1360.79
cmu	137	99.12	7.52	776.01
flint	16	99.88	7.94	625.5
flpop	20	99.9	4.45	506.7
grpop	49	99.67	11.8	1733.18
ironmt	9	99.82	18.33	955.11
jackson	19	99.82	9.26	926
lssu	3	99.69	68.33	1635.33
ltupop	36	99.81	10	1014.97
michnet1	17	99.96	3.76	210.65
michnet5	142	99.82	10.23	964.87
msu	49	99.87	8.55	686.8
mtu	15	99.71	15.93	1538.67
muskpop	43	99.7	12.77	1572.19
nmu	12	99.85	24.75	788.08
oakland	44	99.82	14.57	932.89
oakland3	8	99.9	10.88	520.38
saginaw	24	99.96	4.63	213.33
tcity	20	99.68	11.4	1697.45
umd	19	99.79	8.26	1098.74
wmu	60	99.88	7.55	617.58
wsu	36	99.84	10.69	824.75
wsu1	23	99.85	9.39	767.17

(b) Node Failures

Figure 3. Source and frequency of regional backbone failures.

averages are biased towards less availability by individual interfaces which exhibit a disproportionate number of failures. Specifically, the failure logs reveal a number of persistent circuit or hardware faults which repeatedly disrupt service on a given interface.

Since the trouble ticket system used in our study does not maintain outage duration statistics, we could not relate the duration of outages in Figure 3(b) with the source of outages in Figure 3(a). However, discussions with operations staff and empirical observations indicate that the duration of the most backbone outages tends be small – on the order of several minutes. Customer outages generally persist a bit longer – on the order of several hours. Specifically, most power outages and hardware failures tend to be resolved in four hours or less, and faults stemming from routing problems usually last under two hours. Carrier problems tend to be harder to estimate as the length of time down is independent of the regional provider.

3.4. Frequency

In this section, we examine frequency components of intra and inter-domain routing data. For this analysis, we define a routing update's *frequency* as the inverse of the inter-arrival time between routing updates; a high frequency corresponds to a short inter-arrival time. Other work has been able to capture the lower frequencies through both routing table snapshots [4] and end-to-end techniques [15]. Our measurement apparatus allowed a unique opportunity to examine the high frequency components of network failures.

Normally one would expect an exponential distribution for the inter-arrival time of routing updates, as they might reflect exogenous events, such as power outages, fiber cuts and other natural and human events. In our earlier analysis [9], we found a strong correlation between North American network usage and the level of inter-domain routing information at the major IXPs. Specifically, the graph of inter-domain route failures exhibited the same bell curve centered on 1pm EST as shown on most graphs of network traffic volume [11].

In this section, we repeat the analysis in [9] to identify frequency components in the inter-arrival internal routing updates exchanged within the backbone of our case study provider. We generated a correlogram, shown in [11], of both datasets generated by a traditional fast Fourier transform (FFT) of the autocorrelation function of the data. The graph of BGP data exhibits significant frequencies at seven days, and 24 hours. In marked contrast, the correlogram of intradomain routing information does not exhibit any significant frequency components. The absence of intradomain frequency components suggests much of BGP instability stems from a different class of failures than the hardware and software faults we described in the previous section. In particular, the lack frequency components supports the supposition in [9, 13] that significant levels BGP instability stem from congestion collapse.

As a mechanism for the detection of link-level or host failures, BGP uses the periodic TCP exchange of incremental routing updates and KeepAlives to test and maintain the peering session. If KeepAlives or routing updates are not received within a bounded time period (the router's Hold Timer), the peering session is severed, causing the withdrawal of all the peer's routes

– making them unreachable through the autonomous system and its downstream networks.

Because TCP end-stations adapt to network congestion by reducing the amount of available bandwidth, KeepAlive packets may be delayed during periods of peak network usage. Under these conditions, a KeepAlive may not be received before the remote BGP hold timer expires. This would cause peering sessions to fail at precisely those times when network load was greatest. The effect is most pronounced in internal BGP communication.

4. Conclusion

Our analysis confirms the widely held belief that the Internet exhibits significantly less availability and reliability than the telephony network. The detection of Internet failures is often far less problematic than identification of the failures' origins. Our characterization and analysis of backbone faults was hampered by the lack of standard fault reporting and measurement mechanisms across providers. A number of Internet engineering associations have called for the development of a uniform trouble ticket system schema and mechanisms for inter-provider sharing of the trouble ticket data. Based on our limited case-study of a regional provider, we found that most faults stemmed hardware and software not unique to the Internet's routing infrastructure.

In contrast to our analysis of the routing between providers, we did not find daily or weekly frequency components in our case-study of the internal routing of a regional provider. This absence supports our earlier findings [13] that Internet failures may stem from congestion collapse. Validation of this theory and correlation of faults amongst multiple providers remains an area for future research.

References

[1] R. Becker, L. Clark, D. Lambert, "Events Defined by Duration and Severity with an Application to Network Reliability", Technometrics, 1998.

[2] K. Calvert, M.B. Doar, E.W. Zegura, "Modeling Internet Topology," in *IEEE Communications Magazine*, June 1997.

[3] B. Chinoy, "Dynamics of Internet Routing Information," in *Proceedings of ACM SIGCOMM '93*, pp. 45-52, September 1993.

[4] R. Govindan and A. Reddy, "An Analysis of Inter-Domain Topology and Route Stability," in *Proceedings of the IEEE INFOCOM '97*, Kobe, Japan. April 1997.

[5] B. Halabi, "Internet Routing Architectures." New Riders Publishing, Indianapolis, 1997.

[6] Inverse Network Technology home page, http://www.inverse.net.

[7] Internet Performance Measurement and Analysis project (IPMA), http://www.merit.edu/ipma.

[8] R. Kuhn, "Sources of Failure in the Public Switched Telephone Network." IEEE Computer, Vol. 30, No. 4, April 1997.

[9] C. Labovitz, G.R. Malan, and F. Jahanian, "Internet Routing Instability," in *Proceedings of the ACM SIGCOMM '97*, Cannes, France, August, 1997.

[10] C. Labovitz, G.R. Malan, and F. Jahaniam, "Origins of Pathological Internet Routing Instability," in *Proceedings of the IEEE INFOCOM '98*, New York, March 1999.

[11] C. Labovitz, A. Ahuja, F. Jahanian, "Experimental Study of Internet Stability and Wide-Area Backbone Failures," University of Michigan CSE-TR-382-98, November 1998.

[12] D. Long, A. Muir, R. Golding, "A Longitudinal Survey of Internet Host Reliability," Hewlett Packard Technical Report, HPL-CCD-95-4, February, 1995.

[13] G.R. Malan, and F. Jahanian, "An Extensible Probe Architecture for Network Protocol Performance Measurement," in *Proceedings of the ACM SIGCOMM '98*, Vancouver, Canada, September 1998.

[14] Multi-Threaded Routing Toolkit, National Science Foundation Project (NCR-9318902).

[15] V. Paxson, "End-to-End Routing Behavior in the Internet," in *Proceedings of the ACM SIGCOMM '96*, Stanford, C.A., August 1996.

[16] V. Paxson, S. Floyd, "Why We Don't Know How to Simulate the Internet," in *Proceedings of the 1997 Winter Simulation Conference*," Atlanta, GA, 1997.

[17] D. Pradhan, "Fault Tolerant Computer System Design", Prentice Hall, New Jersey, 1996.

[18] Vital Signs, home page http://www.vitalsigns.com