Experimental Study of Internet Stability and Wide-Area Backbone Failures *

Craig Labovitz and Abha Ahuja
Merit Network, Inc.
4251 Plymouth Road
Ann Arbor, Michigan 48105-2785
{labovit, ahuja}@merit.edu

Farnam Jahanian
University of Michigan
Department of Electrical Engineering and Computer Science
1301 Beal Ave.
Ann Arbor, Michigan 48109-2122
farnam@umich.edu

Abstract

In this paper, we describe an experimental study of Internet 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. Our paper analyzes Internet failures from three different perspectives.

We first examine several recent major Internet failures and their probable origins. These empirical observations illustrate the complexity of the Internet and show that unlike commercial transaction systems, the interactions of the underlying components of the Internet are poorly understood. Next, our examination focuses on the stability of paths between Internet Service Providers. Our analysis is based on the experimental instrumentation of key portions of the Internet infrastructure. Specifically, we logged all of the routing control traffic at five of the largest U.S. Internet exchange points over a three year period. This study of network reachability information found unexpectedly high levels of path fluctuation and an aggregate low mean time between failures for individual Internet paths. These results point to a high level of instability in the global Internet backbone.

While our study of the Internet backbone identifies major trends in the level of path instability between different service providers, these results do not characterize failures inside the network of service provider. The final portion of our paper focuses on a case study of the network failures observed in a large regional Internet backbone. This examination of the internal stability of a network includes twelve months of operational failure logs and a review of the internal routing communication data collected between regional backbone routers. We characterize the type and frequency of failures in twenty categories, and describe the failure properties of the regional backbone as a whole.

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

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. Several recent wide spread Internet failures have led the popular press to predict the imminent "death of the Internet" [21]. Although the predicted Internet collapse has yet to materialize, further analysis of the behavior and characteristics of wide-area network faults is critical for the continued evolution of the Internet.

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) [30, 4, 17]. Studies including [14, 31] 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 [5], the topological stability and dynamics of deployed wide-area Internet Protocol (IP) backbones has gone virtually without formal study, with the exception of [18, 10, 6, 27].

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. As we describe below, the interactions between the underlying components of the Internet are poorly understood [28].

Typical analysis of faults in telephony networks has focused on the number of customers affected by an outage [1]. The US Federal Communication Commission requires service providers to report all outages lasting 30 minutes or more and affecting 30,000 customers or more [17]. No such reporting requirements yet exist for Internet providers. And, if such requirements did exist, the same estimations of the impact of failures would be problematic for Internet providers. Both the definition of failure and even "end-user" are somewhat ambiguous on the Internet. In contrast to the fixed bandwidth used by telephony, Internet applications and end-users have widely disparate bandwidth, latency and loss requirements. For example, the failure of an Internet T3 link (45 MB) may impact one large weather simulation at a supercomputer center, or several thousand web-surfing dial-up users. In our analysis, we make no effort to quantify the significance of Internet outages based on the number of users affected. Instead, we focus on the number of individual link or interface failures, and the number of unreachable network destinations.

In general, 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. No resources are explicitly reserved for each datagram or IP data flow. Instead, the end-to-end quality of Internet performance depends on the impact of loss, queuing

delay and network congestion on each of the flow's individual datagram packets. So, for example, although the initial "call setup" of an Internet telephony application may succeed, all subsequent voice datagrams in the connection may be lost due to network congestion. The relationship between loss, latency and end-to-end performance remains an area of active research.

In addition, the explosive growth in demand for Internet facilities and features has resulted in a significantly more rapid Internet software and hardware evolutionary testing and development cycle than traditional amongst PSTN equipment suppliers. For example, telephony switches typically undergo development cycles on the order of several years or even decades. In contrast, some Internet backbone routers and switches have development cycles lasting six months or less. Internet vendors regularly market backbone equipment featuring new software algorithms even before these protocols have advanced into official standards [13, 24]. The technological demands associated with the Internet's growth are so severe that Internet providers often depend on these newly released products or software features to sustain their network's continued expansion. The abbreviated development cycle has led to a trade-off between reliability and time-to-market. As a result, the reliability of the Internet infrastructure has arguably suffered.

The rapid growth of IP backbones has also led to a decline in the relative level of experience and degree of coordination amongst Internet backbone operators. A number of significant recent Internet outages have stemmed from human error. Other outages have originated, or been exacerbated by lack of coordination between the backbone engineering staff of different Internet providers. In the PSTN network, a comparatively small number of telecommunication companies interact via well-defined, standardized channels using uniform management, measurement and operational procedures. The significantly more diverse and less uniform Internet does not enjoy the same degree of coordination. Specifically, the Internet lacks central administration and coordination. Unlike traditional PSTN standards bodies whose formal membership requirements are defined by international treaty [29], the only requirement for participation in the three yearly Internet standards meetings is showing up [13].

We briefly describe some recent Internet outages which directly, or indirectly, impacted a majority of Internet backbone paths. Although several major incidents stemmed from underlying PSTN failures, we focus below on faults specific to the Internet. We provide the following summaries as anecdotal evidence of the sources of major Internet failures.

- April 25, 1997 — A misconfigured router maintained by a small Virginia service provider injected an incorrect routing map into the global Internet. This map indicated that the Virginia company's network provided optimal connectivity to all Internet destinations. Internet providers that accepted this map automatically diverted all of their traffic to the Virginia provider. The resulting network congestion, instability, and overload of Internet router table memory effectively shut down most of the major Internet backbones for up to two hours. Incorrect published contact information for operations staff, and lack of procedures for inter-provider coordination exacerbated the problem [2].

- August 14, 1998 – A misconfigured critical Internet database server incorrectly referred all queries for Internet machine names ending in ".net" to the wrong secondary database server. As a result, a majority of connections to ".net" Internet web servers and other end stations failed for a period of several hours [26].

- November 8, 1998 – A malformed routing control message stemming from a software fault

triggered an interoperability problem between core Internet backbone routers manufactured by different vendors. This problem lead to a persistent, pathological oscillation and failure in the communication between most Internet core backbone routers. As a result, Internet end-users experienced wide-spread loss of network connectivity, and increased packet loss and latency. The majority of backbone providers resolved the outage within several hours after adding filters which removed the malformed control message [25].

Overall, both Internet and telephony outages stem from a wide range of sources, including faults in the underlying telecommunication switching system, and the higher level software and hardware components. Like Pradhan [30], we are interested in estimating the reliability of Internet backbone paths at specified probability and duration thresholds such as the mean number of events per year, and the mean time spent in events. 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 provides further background on Internet routing and related work. Section 3 describes the infrastructure used in our characterization of backbone failures and the analysis of both inter and intra-domain path stability. Section 4 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 Background

The Internet is divided into a large number of distinct regions of administrative control, commonly called *Autonomous Systems* (AS). An autonomous system (also called a routing domain) typically consists of a network service provider or a large organizational unit, such as a college campus or a corporate network. In turn, each AS interconnects a number subnetworks, such as remote corporate offices or customer networks. Autonomous systems usually have distinct routing policies and connect to one or more remote autonomous systems at neutral private or public *exchange points*.

The routers in the Internet are responsible for receiving and forwarding packets through this interconnected maze of subnetworks and autonomous systems. Each router makes routing decisions based on its knowledge of the topology, the conditions on the network, and complex routing policies as specified by network administrators within their domain. In order to make such dynamic decisions, routers exchange path and topology information using special purpose routing protocols. We often distinguish between two classes of routing protocols: An *inter-domain (or exterior) routing protocol* is used to exchange information between peer routers in different autonomous systems. An *intra-domain (or interior) routing protocol*, in contrast, is used to pass information between routers within an autonomous system.

Figure 1 provides a simple example to illustrate the concepts behind the routing architecture of the Internet. In the diagram, two Internet service providers, or autonomous systems, interconnect at both a public exchange point and via private peering. The routers R1 and R8, and R2 and R7 communicate via an inter-domain routing protocol. The internal routers (R4,R5,R6 and R10,R11,R9) use an intra-domain routing protocol to exchange internal routes.

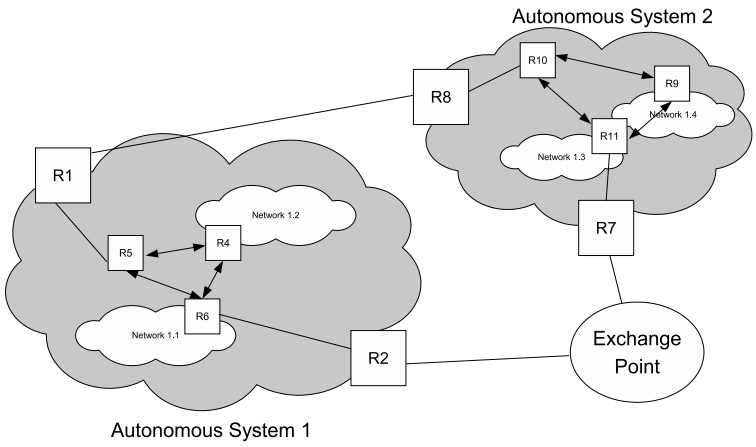

Figure 1: An Illustration of Internet Routing Architecture.

The global Internet encompasses a conglomeration of thousands of interconnected wide-area backbones, each under separate realms of administrative control. Even though backbones may utilize a different underlying technology, such ATM or Frame Relay, all Internet backbones interconnect and interoperate at some level using a standard set of IP based protocols and services. The Internet Protocol *(IP)* provides a common substrate for the encapsulation and forwarding of network datagrams. All Internet services, including reliable transport *(TCP)*, datagram *(UDP)*, and multicast protocols, use IP packets as their foundation. Internet hosts segment application level streams into one or more independently routed IP datagrams. At the edge of every Internet backbone, routers forward these individual IP datagrams to the appropriate next-hop router in adjacent networks. Internet routers build these next-hop routing tables based topological information conveyed in routing control messages exchanged with other routers.

Ten to twelve large Internet service providers dominate the Internet. These national and international providers, often referred to as *tier one* providers, account for majority of routes and bandwidth that comprise the public Internet. Approximately four to six thousand smaller regional

networks, or *tier two* providers peer with the tier one providers at one or more private or public *Internet eXchange Points* (IXPs). At the end of the NSFNet in 1995, the National Science Foundation established five Network Access Points (NAPs) in the continental U.S. These large public exchange points are often considered the *core* of the Internet where providers *peer*, or exchange both routing information and traffic.

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 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. Aggregation provides a critical layer of abstraction that limits the level of globally visible policy and reachability information in the Internet core. Some classes of customer addresses resist easy aggregation, including customers with historic address assignments outside their provider's aggregate, and *multi-homed* customers. A multi-homed customer obtains connectivity from one more providers for redundancy or load-sharing purposes.

The most common inter-domain (exterior) routing protocol used by autonomous systems in the Internet is the Border Gateway Protocol (*BGP*) [11]. BGP is an *incremental* protocol that sends update information only upon changes in network topology or routing policy. In contrast to many interior protocols that build their own reliability on top of a datagram service, BGP uses TCP as its underlying transport mechanism. Routes information exchanged in BGP includes a number of associated attributes, including the address of the next-hop router and a record of the inter-domain path the route has followed through different providers. We refer to this path record of next-hops as the route's *ASPath*. Since BGP routers within a domain synchronize using internal protocols, BGP information collected from any border router should reflect the routing behavior of the each autonomous system pending local router policies, and local hardware or software failures.

Internally within an autonomous system, routers use a variety of intra-domain (interior) protocols to distribute local routing information, including Open Shortest Path First (OSPF), ISIS, and IGRP [12]. In this paper, we focus on the OSPF protocol used by the provider in our case study. Unlike exterior gateway protocols, which only exchange routing updates between adjacent neighbors, every router participating an OSPF backbone periodically floods information to all other OSPF routers throughout the network. The flooded OSPF information describes a router's local interface and link configuration, as well as information received from adjacent OSPF routers. OSPF router uses knowledge of network topology learned in the received updates to build an "open shortest path first tree" to every network and interface in the backbone. The shortest path tree then forms the basis for the router's forwarding, or routing table. In our intra-domain analysis, we focus on the OSPF router links state advertisement (RLSA). Every OSPF router periodically generates a RLSA update to describe the current state of all the router's interfaces. In addition to periodic announcements, OSPF routers also generate RLSA upon any change in the status of an interface.

After a policy change or network failure affects the availability of a path to a set of prefix destinations, the routers topologically closest to the failure will detect the fault, withdraw the route and make a new local decision on the preferred alternative route, if any, to the set of destinations. These routers will then propagate the new topological information to each router within the autonomous system. The network's border routers will in turn propagate the updated information to each

external peer router, pending local policy decisions. Routing policies on an autonomous system's border routers may result in different update information being transmitted to each external peer.

At the border of backbones, most providers use BGP as the inter-domain protocol to redistribute topological information to other backbone providers. The interaction between internal and external gateway protocols varies based on network topology and backbone provider policy. In the case where a customer network is *single-homed*, or only has a single path to the Internet core, providers may choose to statically route the customer. In this configuration, the route to the customer network always will remain static, or constant, in the BGP inter-domain information. At the other extreme, providers may choose to inject all OSPF information directly into BGP. In this configuration, BGP will re-announce all changes to the OSPF shortest path tree affecting the customer network. As an intermediate solution, most providers aggregate OSPF information at the backbone boundary. Multiple OSPF routes will fall under a larger, aggregate prefix. The backbone order router will maintain a path to an aggregate super-net prefix as long as a path to one or more of the component prefixes is available. This effectively limits the visibility of instability stemming from unstable customer circuits or routers to the scope of a single routing domain.

Routing *instability*, informally defined as the rapid change of network reachability and topology information, has a number of origins including router configuration errors, transient physical and data link problems, and software bugs. In both this and our earlier studies [18, 19], we analyze routing control messages exchanged between routers at the major US exchange points.

High levels of network instability can lead to packet loss, increased network latency and time to convergence. At the extreme, high levels of routing instability have led to the loss of internal connectivity in wide-area, national networks. Experience with the NSFNet and wide-area backbones has demonstrated that a router which fails under heavy routing instability can instigate a "*route flap storm*." In this mode of pathological oscillation, overloaded routers are marked as unreachable by BGP peers as they fail to maintain the required interval of Keep-Alive transmissions. As routers are marked as unreachable, peer routers will choose alternative paths for destinations previously reachable though the "down" router and will transmit updates reflecting the change in topology to each of their peers. In turn, after recovering from transient CPU problems, the "down" router will attempt to re-initiate a BGP peering session with each of its peer routers, generating large state dump transmissions. This increased load will cause yet more routers to fail and initiate a storm that begins affecting ever larger sections of the Internet. Several route flap storms in the past year have caused extended outages for several million network customers. The latest generation of routers from several vendors provide a mechanism in which BGP traffic is given a higher priority and Keep-Alive messages persist even under heavy instability.

In our earlier study [18], we found that a majority routing instability (99 percent) consisted of *pathological* updates which did not reflect actual network topological or policy changes. The majority of these pathologies stemmed from specific vendor hardware and software bugs. Since the publication of our findings, the majority of these faulty implementations have been resolved. In this paper, we turn our attention to analysis of "legitimate" faults that reflect actual link or network failures.

In this section, we provided an overview of the Internet architecture and routing protocols primarily to set a context for discussing our experimental results. A number of books, including [11, 29] provide excellent references for additional information. In the next section, we describe our experimental data collection infrastructure.

3 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 [23] and IPMA [16] projects for the collection, analysis and post-processing of our data.

3.1 Inter-domain Data Collection

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 several 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. During the ten month period of our RouteViews study from January 1997 to November 1998, we collected over nine gigabytes of information on the default-free routing tables of three providers.

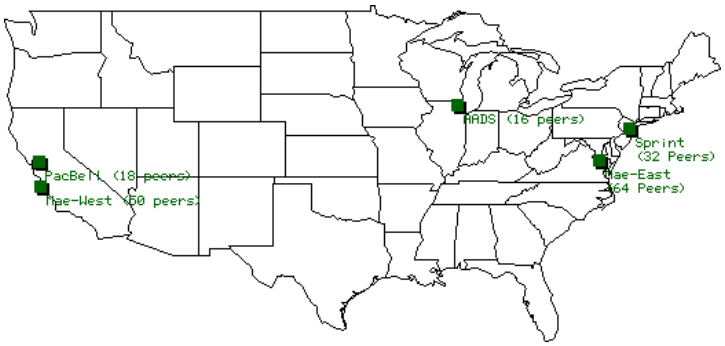

Figure 2: Map of major U.S. Internet exchange points.

In addition, we corroborated our RouteViews data with routing information recorded from the experimental instrumentation of key portions of the Internet infrastructure. Over the course of three years, we logged all BGP routing messages exchanged with probe machines located at five of the major U.S. network exchange points: Mae-East, Sprint, AADS, PacBell and Mae-West. At these geographically diverse exchange points, network service provider routers peer, or provide routing information about their customer routes. Based on the size of the ISP and quality of aggregation, the number of customer routes announced by each provider ranged from a few dozen to several thousand. Figure 2 shows the location of each exchange point, and the number of service providers peering with our probe machine at each exchange.

3.2 Intra-domain Data Collection

We base our analysis of intra-domain failures on a case study of a medium size regional network. This provider's backbone, shown in Figure 3, connects educational and commercial customers in

132 cities via high speed serial lines and frame-relay links at speeds up to OC3. 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.

Figure 3: Network Map of the Regional Provider in Our Case Study.

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 [28], 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. We also use the internal backbone data for confirmation of several previously postulated sources of failure in wide-area networks.

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 the router interfaces throughout the network. The CNMS monitors interfaces through the periodic (every ten minutes) transmission and receipt of "ping" packets. An interface that does not respond to a ping, is marked as "down" and the interface state transition is logged to disk. Similarly, when a down interface responds to a ping, the CNMS records the interface as "up." 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. Usually trouble tickets involve CNMS alerts lasting more than a few minutes (e.g. long enough to get an operator's attention) or a prolonged degradation in the quality of service (e.g. increased latency) to one or more customer sites.

As a means of validating our CNMS-based data on internal failures, we also analyzed six months of router OSPF messages (RLSAs) flooded throughout the regional network between March 1997 and November 1998. We also used this OSPF data for analysis of frequency components of network faults in Section 4.4.

4 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 examine the failure logs and routing data from our case-study of a medium size regional provider. We first examine the frequency and duration of backbone failures and then categorize the sources of the observed failures. Finally, we discuss the relationship between the frequency of intra-domain failures with the behavior of inter-domain routing changes.

4.1 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 3, 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. As of November 1998, the Internet default-free routing contained approximately 55,000 route entries. Due to local policies and network topology, the exact number of routes in the default-free tables varied slightly for each provider over the course of our study.

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 backbone prefix.

In the taxonomy below, we distinguishes between five 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 re-routing of traffic to a given prefix destination after a network failure. Recall from Section 2 that the ASPath represents the routing path of the prefix through different interconnected autonomous systems.

Policy Fluctuation: A route is implicitly withdrawn and replaced with by an alternative route with differing attributes, but the same next-hop and ASPath.

Pathological Routing: The repeated transmission of BGP withdrawals for a prefix that is currently unreachable, or the duplicate announcement of routes with all the same path attributes. This is pathological behavior as BGP routers should only transmit BGP updates that reflect legitimate policy or forwarding changes.

Since both Pathological Routing and Policy Fluctuation do not reflect "legitimate" network failures, we omit both of these classes of routing events from our discussion in the remaining sections. Both Pathological Routing and Policy Fluctuation were addressed in [18, 19].

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. As described in Section 2, default-free routes represent large, aggregated sections of the Internet. Most providers will maintain a route to an aggregate, supernet address as long as a path to one or more the component prefixes is available. This aggregation effectively limits the interdomain visibility of most faults stemming from unstable customer circuits or routers.

Lacking internal knowledge of the policies and design of remote autonomous systems, 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. For example, we may not know if ISP1's loss of a route to network 1.1 in Figure 1 represents a circuit cut, or the the decision of a customer in the 1.1 netblock to renumber their IP addresses.

In an effort to focus our analysis on actual failures, we limit our dataset to only those prefixes present in each ISP's routing table for more than an aggregate 60 percent (170 days) of our nine month study. This availability requirement filtered an average of 11,000 routes (20 percent) from our default-free routing table analysis. The removal of these short-lived routes provides a more conservative estimate of network failures.

In addition, we modified our analysis to limit any bias in our data stemming from failures in our local network infrastructure. The BGP protocol specification requires that all routes from an adjacent router be withdrawn upon the loss of the remote peering session. Since our RouteViews BGP connections to ISP1,ISP2 and ISP3 transit several commodity networks, we do not include the loss of a RouteViews TCP peering session as a source of failure for routing table entries. This modification may filter some "legitimate" failures in the ISP1, ISP2, and ISP3 backbones, and artificially biases our failure data towards longer mean-time to failures.

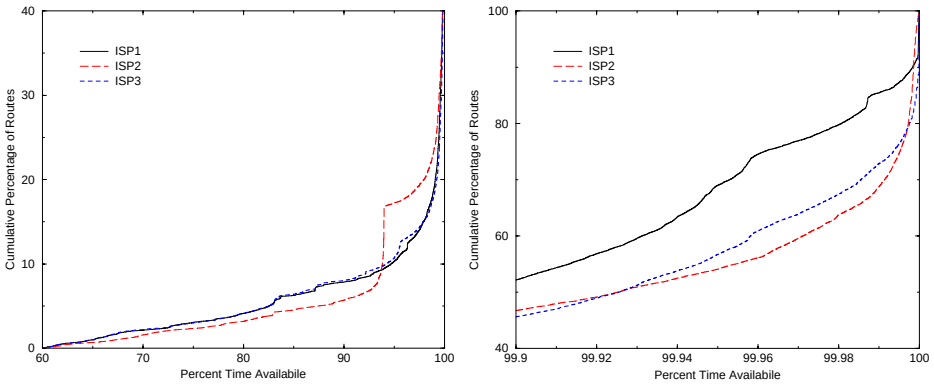

Figure 4: Cumulative distribution of the route availability of three service providers.

As a final modification to our data, we applied a fifteen minute filter window to all BGP route transitions. Specifically, we count multiple failures during this filter period as a single failure. This filter limits the bias of high frequency pathological BGP described in our earlier analysis [18, 19]. Moreover, ongoing discussions with providers suggests that fifteen minutes is approximately the time required for Internet routing to reach convergence [24].

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 2, provider's regularly aggregate multiple more specific network addresses into a single supernet advertisement.

The graphs in Figure 4 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. As described earlier, both graphs only include routes available for more than 60 percent of the time during our study. Both graphs in Figure 4 represent the same data, but Figure 4(b) provides an expanded view of route availability above 99.9 percent.

A recent study study [17] found that the PSTN averaged an availability rate better than 99.999 percent during a one year period. From the graph in Figure 4(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 dramatic 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 4(a). The step in the curve for ISP3 at 95 percent availability represents a multi-hour loss of inter-provider connectivity due the November 8 incident described in Section 1. ISP1 exhibits significant less availability above 99.9 than ISP2 and ISP3 as evinced by the higher curve in Figure 4(b).

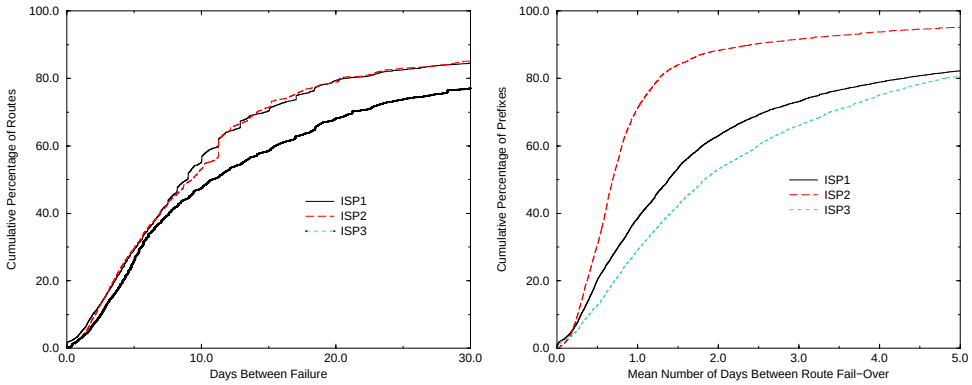

(a) Failure (b) Fail-Over

Figure 5: Cumulative distribution of the mean-time to failure and mean-time to fail-over for default-free routes from three ISPs.

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 5 show the cumulative distribution of the mean number of days between route failures (a), and route fail-over (b) for routes from ISP1, ISP2 and ISP3. The horizontal axes represent the number of ISP routes that exhibit a specific mean-time to failure/fail-over or less; the vertical axes show the cumulative proportion of the ISP's routing table entries for all such events. Examining the graph in Figure 5(a), 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 described in Section 2, a growing number of Internet sites are multi-homed, or possess redundant connectivity to the Internet through multiple providers. In addition, 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. We describe this dynamic rerouting around a fault as fail-over. As noted earlier, changes in the routing path of a destination may actually represent routing policy changes.

Since not all Internet routes enjoy redundant connectivity, we focus our analysis on fail-over by modifying the vertical axis in Figure 5(b) 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 Paxon's 1994 study [27] which found 2/3's of Internet paths persisted for either days

or weeks.

The graph in Figure 6(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 6(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 6(a), but at a slower rate in the Figure 6(b) MTTR graph. The lower average mean-time to failure, but slower mean-time to repair contributes to ISP1's overall lower availability in Figure 4.

Overall, analysis of our MTTR data agrees with our qualitative findings in Section 4.2 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 [27] that most Internet outages are short-lived – lasting on the order seconds or minutes.

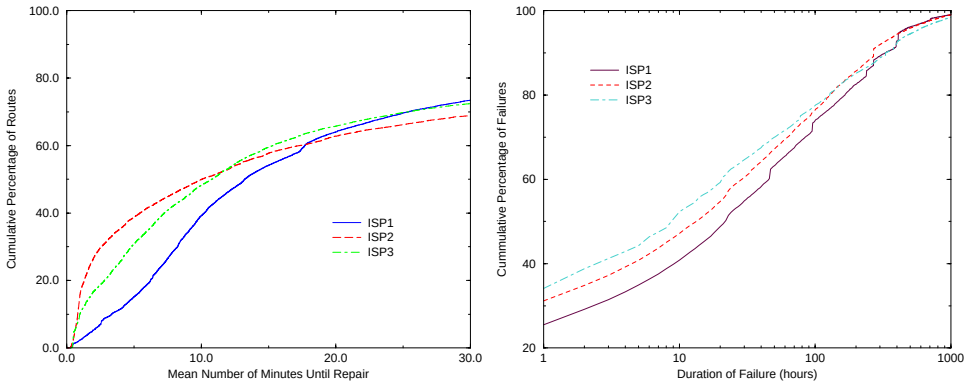

(a) MTTR (b) Duration

Figure 6: Cumulative distribution of MTTR and failure duration for routes from three providers.

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 6(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 6(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 6(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 6(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 [18] that a small fraction routes are responsible for the majority of network instability.

In this section, we described our analysis of the failures in the inter-domain paths between providers. In the next subsection, we look inside the "black box," and explore failures that occur inside of a provider's backbone.

4.2 Analysis of Intra-Domain Network Stability

Having examined the stability of inter-domain paths, we now focus on intra-domain failures. As described in Section 2, intra-domain routing serves as the basis for much of the information exchanged in inter-domain routing. Analysis of the faults associated with an intra-domain network also provides insight into failures in other areas of the Internet.

We provide a case-study of the intra-domain stability of a medium-size, regional Internet provider. This network, described in Section 3, connects educational and commercial customers in 132 cities via high speed serial lines and frame-relay links at speeds up to OC3. Internally, this network uses OSPF as its intra-domain routing protocol. Externally, this network BGP peers with multiple providers for Internet connectivity.

The graph in figure 7(a) 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 7(a), 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 7(a) 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 our central data collection host.

The graph in Figure 7(b) shows the cumulative mean-time to repair for the two different categories

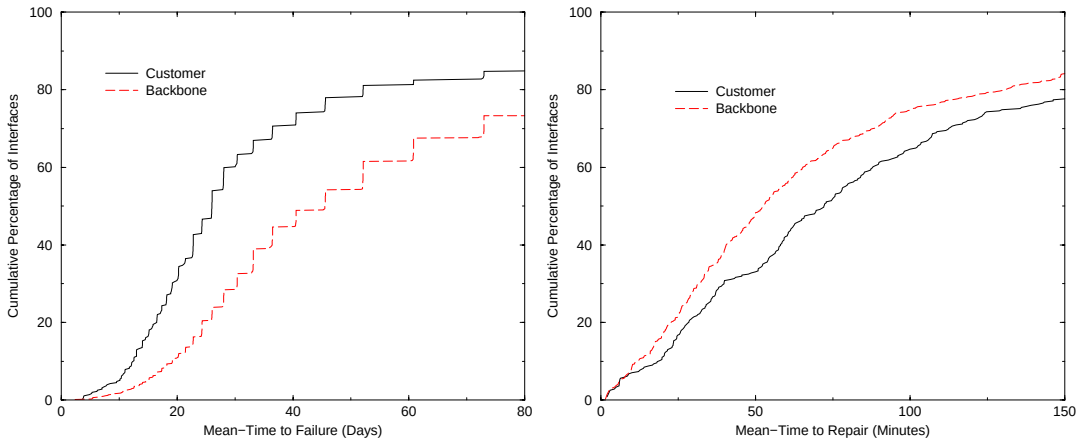

(a) MTTF (b) MTTR

Figure 7: Cumulative distribution of the mean-time to failure and mean-time to repair for backbone interface in a regional Internet provider.

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 under 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.

4.3 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 sections. As discussed in Section 3, we base our characterization of network failures on the operational trouble logs of a regional ISP.

Each trouble ticket provides a number of fields, including the trouble found, duration of ticket, source of outage, and type of equipment involved. The regional ISP we studied manages a large and diverse number of devices, ranging from dial-in modems and web servers to backbone routers. In this paper, we limit our discussion to outages associated with routers.

Table 1 shows a breakdown of all the outages recorded during our one-year case study (November 1997 to November 1998). As diagnosing and categorizing 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. 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.

Outage Category	Number of Occurrences	Percentage
Maintenance	272	16.2
Power Outage	273	16.0
Fiber Cut/Circuit/Carrier Problem	261	15.3
Unreachable	215	12.6
Hardware problem	154	9.0
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

Table 1: Category and number of recorded outages Internet in a regional Internet provider between November 1997 and November 1998.

From Table 1, 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 Table 1 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.

Table 2 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 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 Table 2 with the source of outages in Table 1. 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

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	625.50
flpop	20	99.90	4.45	506.70
grpop	49	99.67	11.80	1733.18
ironmt	9	99.82	18.33	955.11
jackson	19	99.82	9.26	926.00
lssu	3	99.69	68.33	1635.33
ltupop	36	99.81	10.00	1014.97
michnet1	17	99.96	3.76	210.65
michnet5	142	99.82	10.23	964.87
msu	49	99.87	8.55	686.80
mtu	15	99.71	15.93	1538.67
muskpop	43	99.70	12.77	1572.19
nmu	12	99.85	24.75	788.08
oakland	44	99.82	14.57	932.89
oakland3	8	99.90	10.88	520.38
saginaw	24	99.96	4.63	213.33
tcity	20	99.68	11.40	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

Table 2: Availability of Router Interfaces during a one year case study (November 1997 - November 1998) of a regional network.

last under two hours. Carrier problems tend to be harder to estimate as the length of time down is independent of the regional provider.

4.4 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. We were particularly interested in the high frequency component of routing instability in our analysis. Other work has been able to capture the lower frequencies through both routing table snapshots [10] and end-to-end techniques [27]. 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 [18], 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 [22].

In this section, we repeat our analysis of [18] in identifying frequency components in the OSPF data collected during our case study. A more rigorous approach to identifying temporal trends in the routing updates was undertaken using time series analysis. Specifically, we analyzed time-stamped logs of all received BGP and OSPF updates using spectrum analysis. The BGP data covered August through September of 1997, while our OSPF data included November 1997 to November 1998. We detrended both datasets in a manner similar to the treatment of Beverage's wheat prices

by Bloomfield in [3]. The rate of routing updates is modeled as $x_t = T_t I_t$, where T_t is the trend at time t and I_t is an irregular or oscillating term. Since all three terms are strictly positive, we conclude that $log x_t = log T_t + log I_t$. T_t can be assumed as some value of x near time t, and I_t some dimensionless quantity close to 1; hence $log I_t$ oscillates about 0. This avoids adding frequency biases that can be introduced due to linear filtering.

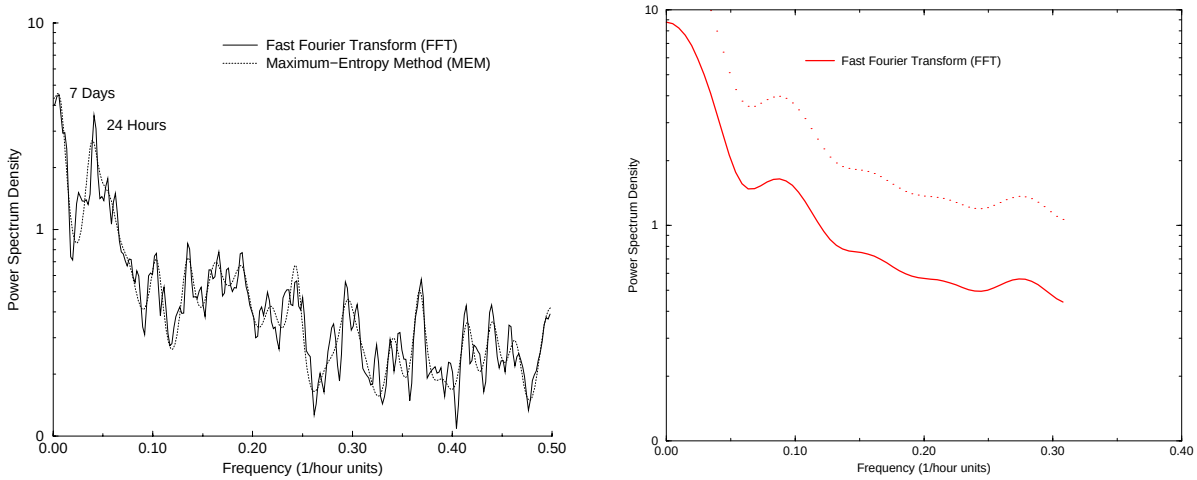

(a) BGP

(b) OSPF

Figure 8: Results from time series analysis of the inter-domain routing updates measured at the Mae-East exchange point during August and September 1996, and the intra-domain routing updates measured inside a regional provider during October and November 1998 using hourly aggregates.

Figure 8 shows a correlogram of both datasets generated by a traditional fast Fourier transform (FFT) of the autocorrelation function of the data. The graph of BGP data in Figure 8(a) shows significant frequencies at seven days, and 24 hours. The seven day frequencies correspond to the low amount of instability measured on the weekends in comparison to the relatively high amount during the week. The 24 hour periodicity in inter-domain routing instability coincides with the observations in [18] that during the early morning hours, the Internet is fairly stable, in sharp contrast to spikes encountered during North American business hours.

In marked contrast to the BGP data, the correlogram of OSPF routing information in Figure 8(b) does not exhibit any significant frequency components. The absence of intra-domain 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 [18, 20] that significant levels BGP instability stem from congestion collapse. As described in Section 2, BGP uses TCP as its underlying transport. 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. Autonomous systems use an internal form of BGP, *IBGP*, to fully connect and synchronize their border routers. Through a mesh of TCP connections traversing the interior of the autonomous system, border routers reach consensus on routes exported and polices applied to neighboring ASes. The probability that an IBGP session's datagram is lost at one of these congested interior routers increases during high network utilization making this result broader. The latest generation of routers from several vendors provide a mechanism in which BGP traffic is given a higher priority and Keep-Alive messages persist even under congestion.

5 Conclusion

In the early days of the DARPANet[1], few, if any, mission-critical public or commercial services placed significant dependency on the network. Outages were often frequent, prolonged and generally went unnoticed. Although the imminent "death of the Internet" has yet to materialize [21], our analysis confirms the widely held belief [21] that the Internet exhibits significantly less availability and reliability than the telephony network

Still, the Internet has proven remarkably robust. Underlying advances and upgrades in Internet hardware and software infrastructure have forestalled the most serious problems of bandwidth shortages and a periodic lack of router switching capacity. In today's Internet, commercial and mission critical applications are increasingly migrating towards using the now ubiquitous network as a communication medium. It is important to understand and characterize routing instability for protocol design and system architecture evolution.

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 groups, including [15, 24] 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 [20] 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.

By directly measuring the network availability of several Internet Service Providers, this paper identified several important trends in the stability and failure characteristics of the Internet. This work in conjunction with several other research efforts has begun to examine inter-domain routing

[1] The DARPANet was the precursor to the NSFNet (1987-1996) and today's commercial Internet

through experimental measurements. These research efforts help characterize the effect of added topological complexity in the Internet since the end of of the NSFNet backbone. Further studies are crucial for gaining insight into routing behavior and network performance so that a rational growth of the Internet can be sustained.

Acknowledgments

We wish to thank Sean Donelan, Sue Hares, Susan R. Harris, Gerald R. Malan, Allyn Romanow, and Mark Turner for their comments and helpful insights.

References

[1] Alliance for Telecommunications Industry Solutions, "A Technical Report on Reliability and Survivability Aspects of Interactions Between the Internet and Public Telecommunications Network", T1A1.2/98-001R6.

[2] R. Barrett, S. Haar, R. Whitestone, "Routing Snafu Causes Internet Outage," Interactive Week, April 25, 1997.

[3] Bloomfield P."Fourier Analysis of Time Series: An Introduction." John Wiley & Sons, New York. 1976.

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

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

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

[7] S.Dawson, F. Jahanian, and T. Mitton, "Experiments on Six Commercial TCP Implementations Using a Software Fault Injection Tool (tech report version)," Software Practice and Experience, vol. 27, no. 12, pp. 1385-1410, December 1997.

[8] Dettinger, M.D., Ghil, M., Strong, C.M., Weibel, W., and Yiou, P. "Software Expedites Singular-Spectrum Analysis of Noisy Time Series," in Eos, Trans. American Geophysical Union, v. 76(2), p. 12, 14, 21.

[9] S. Floyd, and V. Jacobson, "The Synchronization of Periodic Routing Messages," *IEEE/ACM Transactions on Networking*, V.2 N.2, p. 122-136, April 1994.

[10] 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.

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

[12] C. Huitema, "Routing In the Internet," Prentice Hall, New Jersey, 1995.

[13] G, Malkin, "The Tao of the IETF", RFC-1718.

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

[15] IOPS home page, http://www.iops.org

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

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

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

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

[20] 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.

[21] B. Metcalf, "Predicting the Internet's Catastrophic Collapse and Ghost Sites Galore in 1996," *InfoWorld*, December 4, 1995.

[22] S. Feldman, "State of Mae-East," presented at "NANOG8," Ann Arbor, Michigan, October 1996.

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

[24] North American Network Operators Group (NANOG) home page, http://www.nanog.org/

[25] North American Network Operators Group (NANOG) mailing list, http://www.merit.edu/mail.archives/html/nanog/msg03039.html.

[26] North American Network Operators Group (NANOG) mailing list, http://www.merit.edu/mail.archives/html/nanog/msg00569.html.

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

[28] 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.

[29] D. Piscitello, A. Chapin, "Open Systems Networking," Addison-Wesley, Reading, MA, 1993.

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

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