Internet Routing Instability *

Craig Labovitz, G. Robert Malan, and Farnam Jahanian
University of Michigan
Department of Electrical Engineering and Computer Science
1301 Beal Ave.
Ann Arbor, Michigan 48109-2122
{labovit, rmalan, farnam}@eecs.umich.edu

Abstract

This paper examines the network inter-domain routing information exchanged between backbone service providers at the major U.S. public Internet exchange points. Internet routing instability, or the rapid fluctuation of network reachability information, is an important problem currently facing the Internet engineering community. 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 lead to the loss of internal connectivity in wide-area, national networks. In this paper, we describe several unexpected trends in routing instability, and examine a number of anomalies and pathologies observed in the exchange of inter-domain routing information. The analysis in this paper is based on data collected from BGP routing messages generated by border routers at five of the Internet core's public exchange points during a nine month period. We show that the volume of these routing updates is several orders of magnitude more than expected and that the majority of this routing information is redundant, or pathological. Furthermore, our analysis reveals several unexpected trends and ill-behaved systematic properties in Internet routing. We finally posit a number of explanations for these anomalies and evaluate their potential impact on the Internet infrastructure.

1 Introduction

Since the end of the NSFNet backbone in April of 1995, the Internet has seen explosive growth in both size and topological complexity. This growth has placed severe strain on the commercial Internet infrastructure. Regular network performance degradations stemming from bandwidth shortages and a lack of router switching capacity, have lead the popular press to decry the imminent death of the Internet [13]. *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. Instability, also referred to as "route flaps", significantly contributes to poor end-to-end network performance and degrades the overall efficiency of the Internet infrastructure. All of these sources of network instability result in a large number of routing updates that are passed to the core Internet exchange point routers. Network instability can spread from router to router and propagate throughout the network. At the extreme, route flaps have led to the transient loss of connectivity for large portions of the Internet. Overall, instability has three primary effects: increased packet loss, delays in the time for network convergence, and additional resource overheard (memory, CPU, etc.) within the Internet infrastructure.

The Internet is comprised of a large number of interconnected regional and national backbones. The large public exchange points are often considered the "core" of the Internet, where backbone service providers *peer*, or exchange traffic and routing information with one another. 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.

The Internet is divided into a large number of different regions of administrative control commonly called *autonomous systems*. These autonomous systems (AS) usually have distinct routing policies and connect to one or more remote autonomous systems at private or public *exchange points*. Autonomous systems are traditionally composed of network service providers or large organizational units like college campuses and corporate networks. At the boundary of each autonomous system, peer border routers exchange reachability information to destination IP address blocks [2], or *prefixes*, for both transit networks, and networks originating in that routing domain. Most autonomous systems exchange routing information through the Border Gateway Protocol (*BGP*) [12].

Unlike interior gateway protocols, such as IGRP and OSPF, that periodically flood an intra-domain network with all known routing table entries, BGP is an *incremental* protocol that sends update information only upon changes in network topology or routing policy. Moreover, BGP uses TCP as its underlying transport mechanism in contrast to many interior protocols that build their own reliability on top of a datagram service. As a path vector routing protocol, BGP limits the distribution of a router's reachability information to its *peer*, or neighbor routers. A *path* is a se-

quence of intermediate autonomous systems between source and destination routers that form a directed route for packets to travel. Router configuration files allow the stipulation of *routing policies* that may specify the filtering of specific routes, or the modification of path attributes sent to neighbor routers. Routers may be configured to make policy decisions based on both the announcement of routes from peers and their accompanying attributes. These attributes, such as *Multi Exit Descriptor* (MED), may serve as hints to help routers chose between alternate paths to a given destination.

Backbone border routers at public exchange points commonly have thirty or more external, or *inter-domain*, peers, as well as a large number of *intra-domain* peering sessions with internal backbone routers. After each router makes a new local decision on the best route to a destination, it will send that *route*, or path information along with accompanying distance metrics and path attributes, to each of its peers. As this reachability information travels through the network, each router along the path appends its unique AS number to a list in the BGP message. This list is the route's *ASPATH*. An ASPATH in conjunction with a prefix provide a specific handle for a one-way transit route through the network.

Routing information shared between peers in BGP has two forms: announcements and withdrawals. A route *announcement* indicates a router has either learned of a new network attachment or has made a policy decision to prefer another route to a network destination. Route *withdrawals* are sent when a router makes a new local decision that a network is no longer reachable. We distinguish between *explicit* and *implicit* withdrawls. Explicit withdrawls are those associated with a withdrawl message; whereas an implicit withdrawl occurs when an existing route is replaced by the announcement of a new route to the destination prefix without an intervening withdrawl message. A BGP *update* may contain multiple route announcements and withdrawals. In an optimal, stable wide-area network, routers only should generate routing updates for relatively infrequent policy changes and the addition of new physical networks.

In this paper, we measured the BGP updates generated by service provider backbone routers at the major U.S. public exchange points. Our experimental instrumentation of these exchanges points has provided significant data about the internal routing behavior of the core Internet. This data reflects the stability of inter-domain Internet routing, or changes in topology or policy between autonomous systems. Intra-domain routing instability is not explicitly measured, and is only indirectly observed through BGP information exchanged with a domain's peer. We distinguish between three types of inter-domain routing updates: *forwarding instability* may reflect legitimate topological changes and affects the paths on which data will be forwarded between autonomous systems; *routing policy fluctuation* reflects changes in routing policy information that may not affect forwarding paths between autonomous systems; and *pathological* updates are *redundant* BGP information that reflect neither routing nor forwarding instability. We define *instability* as an instance of either forwarding instability or policy fluctuation. Although some of the preliminary results of our study have been reported at recent NANOG, IETF, and IEPG meetings, this paper is the first detailed written report of our findings. The major results of our work include:

- The number of BGP updates exchanged per day in the Internet core is one or more orders of magnitude larger than expected.

- Routing information is dominated by pathological, or redundant updates, which may not reflect changes in routing policy or topology.

- Instability and redundant updates exhibit a specific periodicity of 30 and 60 seconds.

- Instability and redundant updates show a surprising correlation to network usage and exhibit corresponding daily and weekly cyclic trends.

- Instability is not dominated by a small set of autonomous systems or routes.

- Instability and redundant updates exhibit both strong high and low frequency components. Much of the high frequency instability is pathological.

- Discounting policy fluctuation and pathological behavior, there remains a significant level of Internet forwarding instability.

- This work has led to specific architectural and protocol implementation changes in commercial Internet routers through our collaboration with vendors.

The remainder of this paper is organized as follows: Section 2 describes the infrastructure used to collect the routing stability data analyzed in this paper. Section 3 provides further background on Internet routing and related work. Section 4 describes a number of anomalies and pathologies observed in BGP routing information. It defines a taxonomy for discussing the different categories of BGP update information, and posits a number of plausible explanations for the anomalous routing behavior. Section 5 describes key trends and characteristics of forwarding instability. Finally, the paper concludes with a discussion on the possible impact of different categories of instability on the performance of the Internet infrastructure.

2 Methodology

Our analysis in this paper is based on data collected from the experimental instrumentation of key portions of the Internet infrastructure. Over the course of nine months, we logged BGP routing messages exchanged with the Routing Arbiter project's route servers at five of the major U.S. network exchange points: AADS, Mae-East, Mae-West, PacBell, and Sprint. At these geographically diverse exchange points, network service providers peer by exchanging both traffic and routing information. The largest public exchange, Mae-East located near Washington D.C., currently hosts over 60 service providers, including ANS, BBN, MCI, Sprint, and UUNet. Figure 1 shows the location of each exchange point, and the number of service providers peering with the route servers at each exchange.

Although the route servers do not forward network traffic, they do peer with the majority (over 90 percent) of the service providers at each exchange point. The route servers provide aggregate route server BGP information to a number of client peers. Unlike the specialized routing hardware used by most service providers, the route servers are Unix-based systems which provide a unique platform for exchange point statistics collection and monitoring.

The Routing Arbiter project has amassed 12 gigabytes of compressed data since January of 1996. In January 1997, the operational phase of the Routing Arbiter project ended.
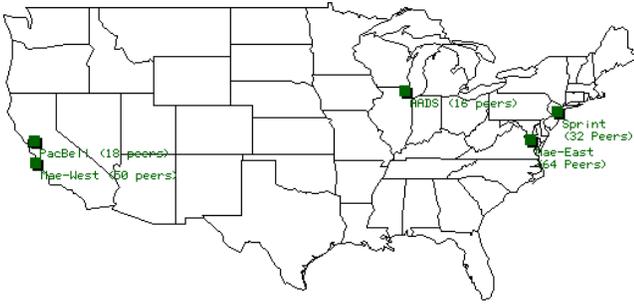
2

Figure 1: Map of major U.S. Internet exchange points.

Data collection and analysis has continued under the auspices of the Internet Performance Measurement and Analysis (IPMA) project [8]. We use several tools from the Multithreaded Routing Toolkit (MRT) toolkit [9] to decode and analyze the BGP packet logs from the route server peering sessions. Although we analyze data from all of the major exchange points, we simplify the discussion in much of this paper by concentrating on the logs of the largest exchange, Mae-East. We analyze the BGP data in an attempt to characterize and understand both the origins and operational impact of routing instability. For the purposes of data verification, we have also analyzed sample BGP backbone logs from a number of large service providers [1].

Increasingly, major Internet service providers (ISP) are utilizing private peering points for the exchange of inter-domain traffic. However, this role was not significant during the data collection period represented by the analysis in this work. A greater level of cooperation with the major ISPs will be needed in the future for continued measurement of Internet routing instability.

3 Background

The fluctuation of network topology can have a direct impact on end-to-end performance. A network that has not yet reached convergence may drop packets, or deliver packets out of order. In addition, through analysis of our data and ongoing discussions with router vendors, we have found that a significant number of the core Internet routers today are based on a *route caching* architecture [11]. In this architecture, routers maintain a routing table cache of destination and next-hop lookups. As long as the router's interface card finds a cache entry for an incoming packet's destination addresses, the packet is switched on a "fast-path" independently of the router's CPU. Under sustained levels of routing instability, the cache undergoes frequent updates and the probability of a packet encountering a cache miss increases. A large number of cache misses results in increased load on the CPU, increased switching latency and the loss of packets. A number of researchers are currently studying the effects of loss and out-of-order delivery on TCP and UDP-based applications [23]. A number of vendors have developed a new generation of routers that do not require caching and are able to maintain the full routing table in memory on the forwarding hardware. Initial empirical observations suggest these routers do not exhibit the same pathological loss under heavy routing update load [11].

Internet routers may experience severe CPU load and

memory problems at heavy levels of routing instability. Many of the commonly deployed Internet routers are based on the older Motorola 68000 series processor. Under stable network conditions, these low-end processors are sufficient for most of the router's computational needs since the bulk of the router's activity happens directly on the forwarding hardware, leaving the processor to handle the processing of BGP and interior gateway protocol (IGP) messages. But heavy instability places larger demands on a router's CPU and may frequently lead to problems in memory consumption and queuing delay of packet processing. Frequently, the delays in processing are so severe that routers delay routing Keep-Alive packets and are subsequently flagged as down, or unreachable by other routers. We have deterministically reproduced this effect under laboratory conditions with only moderate levels of route fluctuation. These experiments are corroborated by the experience of router vendors and ISP backbone engineers.

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 choose alternative paths for destinations previously reachable through 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 (including Cisco Systems and Ascend Communications) provide a mechanism in which BGP traffic is given a higher priority and Keep-Alive messages persist even under heavy instability.

Instability is not unique to the Internet. Rather, instability is characteristic of any dynamically adaptive routing system. Routing instability has a number of possible origins, including problems with leased lines, router failures, high levels of congestion and software configuration errors. After one or more of these problems 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.

The ASPATH attribute present in each BGP announcement allows routers to detect, and prevent *forwarding loops*. We define a forwarding loop as a steady-state cyclic transmission of user data between a set of peers. As described earlier, upon receipt of an update every BGP router performs loop verification by testing if its own autonomous system number already exists in the ASPATH of an incoming update. Until recently, many backbone engineers believed that the ASPATH mechanism in BGP was sufficient to ensure

network convergence. A recent study, however, has shown that under certain unconstrained routing policies, BGP may not converge and will sustain persistent route oscillations [21].

A number of solutions have been proposed to address the problem of routing instability, including the deployment of route dampening algorithms and the increased use of route aggregation [18, 19, 2]. *Aggregation*, or *supernetting*, combines a number of smaller IP prefixes into a single, less specific route announcement. Aggregation is a powerful tool to combat instability because it can reduce the overall number of networks visible in the core Internet. Aggregation also hides, or abstracts, information about individual components of a service provider's networks at the edges of the backbone. Aggregation is successful when there is cooperation between service providers and well-planned network addressing. Unfortunately, the increasingly competitive Internet is sometimes lacking both.

Compounding the problem, a rapidly increasing number of end-sites are choosing to obtain redundant connectivity to the Internet via multiple service providers. This redundant connectivity, or *multi-homing*, requires that each core Internet router maintain a more specific, or longer, prefix in addition to any less specific aggregate address block prefixes covering the multi-homed site.

Our study shows that more than 25 percent of prefixes are currently multi-homed and non-aggregatable. Further, we find that the prevalence of multi-homing exhibits a relatively steep linear rate of growth. This result is consistent with some of the recent findings of Govindan and Reddy [6].

Route servers provide an additional tool to help backbone operators cope with the high levels of Internet routing instability. Each router at an exchange point normally must exchange routing information with every other peer router. This requires $O(N^2)$ bilateral peering sessions, where N is the number of peers. Although route servers do not help limit the flood of instability information, they do help offload computationally complex peering from individual routers onto a centralized route server. This server maintains peering sessions with each exchange point router and performs routing table policy computations on behalf of each client peer. The route server transmits a summary of post-policy routing table changes to each client peer. Each peer router then needs only to maintain a single peering session with the route server, reducing the number of peering sessions to $O(N)$.

A number of vendors have also implemented route dampening [22] algorithms in their routers. These algorithms "hold-down", or refuse to believe, updates about routes that exceed certain parameters of instability, such as exceeding a certain number of updates in an hour. A router will not process additional updates for a dampened route until a preset period of time has experienced. Route dampening algorithms, however, are not a panacea. Dampening algorithms can introduce artificial connectivity problems, as "legitimate" announcements about a new network may be delayed due to earlier dampened instability. A number of ISPs have implemented a more draconian version of enforcing stability by either filtering all route announcements longer than a given prefix length or refusing to peer with small service providers.

Overall, our research has shown that the Internet continues to exhibit high levels of routing instability despite the increased emphasis on aggregation and the aggressive deployment of route dampening technology. Further, recent studies have shown that the Internet topology is growing increasingly less hierarchical with the rapid addition of new exchange points and peering relationships [6]. As the topological complexity grows, the quality of Internet address aggregation will likely decrease, and the potential for instability will increase as the number of globally visible routes expands. Since commercial and mission critical applications are increasingly migrating towards using the Internet as a communication medium, it is important to understand and characterize routing instability for protocol design and system architecture evolution.

The behavior and dynamics of Internet routing stability have gone virtually without formal study, with the exception of Govindan and Reddy [6], Paxson [17] and Chinoy [3]. Chinoy measured the instability of the NSFNet backbone in 1993. Unlike the current commercial Internet, the now decommissioned NSFNet had a relatively simple topology and homogeneous routing technology. Chinoy's analysis did not focus on any of the pathological behaviors or trends we describe in this paper [3].

Paxson studied routing stability from the standpoint of end-to-end performance [17]. We approach the analysis from a complimentary direction − by analyzing the internal routing information that will give rise to end-to-end paths. The analysis of this paper is based on data collected at Internet routing exchange points. Govidian examined similar data, but focused primarily on gross topological characterizations, such as the growth and topological rate of change of the Internet [6].

4 Analysis of Pathological Routing Information

In this section, we first discuss the expected behavior of a well-behaved inter-domain routing system. We then describe the *observed* behavior of Internet routing, and define a taxonomy for discussing the different classifications of routing information. We will demonstrate that much of the behavior of inter-domain routing is pathological and suggests widespread, systematic problems in portions of the Internet infrastructure. We distinguish between three classes of routing information: forwarding instability, policy fluctuation, and pathologic (or redundant) updates. In this section we focus on the characterization of pathological routing information. In Section 5, we will discuss long-term trends and temporal behavior of both forwarding instability and policy fluctuation.

Although the default-free Internet routing tables currently contain approximately 45,000 prefixes [8], our study has shown that routers in the Internet core currently exchange between three and six million routing prefix updates each day. On average, this accounts for 125 updates *per network* on the Internet *every day*. More significantly, we have found that the flow of routing update information tends to be extremely bursty. At times, core Internet routers receive bursts of updates at a rates exceeding several hundred prefix announcements a second. Our data shows that on at least one occasion, the total number of updates exchanged at the Internet core has exceeded 30 million per day[2]. This aggregate rate of instability can place a substantial load on recipient routers as each route may be matched against a potentially extensive list of policy filters and operators. The current high level of Internet instability is a significant problem for all but the most high-end of commercial routers. And even high-end routers may experience increasing levels

---

[2]Our data collection infrastructure failed for the day after recording 30 million updates in a six hour period. The number of updates that day may actually have been much higher.

of packet loss, delay, and time to reach convergence as the rate of instability increases.

In this paper, we analyze sequences of BGP updates for each (prefix, peer) tuple over the duration of our nine month study. As we describe later, the majority of BGP updates from a peer for a given prefix exhibit a high locality of reference, usually occurring within several minutes of each other. In these sequences of updates for a given (prefix, peer) tuple, we identify five types of successive events:

WADiff: A route is explicitly withdrawn as it becomes unreachable and it is later replaced with an alternative route to the same destination. The alternative route differs in its ASPATH or nexthop attribute information. This is a type of forwarding instability.

AADiff: A route is implicitly withdrawn and replaced by an alternative route as the original route becomes unreachable, or a preferred alternative path becomes available. This is a type of forwarding instability.

WADup: A route is explicitly withdrawn and then reannounced as reachable. This may reflect transient topological (link or router) failure, or it may represent a pathological oscillation. This is generated by either forwarding instability or pathological behavior.

AADup: A route is implicitly withdrawn and replaced with a duplicate of the original route. We define a *duplicate route* as a subsequent route announcement that does not differ in the nexthop or ASPATH attribute information. This may reflect pathological behavior as a router should only send a BGP update for a change in topology or policy. Since our initial study only examined the attributes reflective of inter-domain forwarding path (ASPATH and nexthop), this may also reflect policy fluctuation.

WWDup: The repeated transmission of BGP withdrawals for a prefix that is currently unreachable. This is pathological behavior.

4.1 Gross Observations

In the remainder of the paper, we will refer to AADiff, WADiff and WADup as *instability*. We will refer to WWDup as *pathological instability*. AADup may represent either pathological instability or policy fluctuation. A BGP update may contain additional attributes (MED, communities, localpref, etc.), but only changes in the (Prefix, NextHop, ASPATH) tuple will reflect inter-domain topological changes, or forwarding instability. Successive prefix advertisements with differences in other attributes may reflect routing policy changes. For example, a network may announce a route with a new BGP community. The new community represents a policy change, but may not directly reflect a change in the inter-domain forwarding path of user data.

In principle, the introduction of classless inter-domain routing (CIDR) [19] has allowed backbone operators to group a large number of customer network IP addresses into one or more large "supernet" route advertisements at their autonomous system's boundaries. A high level of aggregation will result in a small number of globally visible prefixes, and a greater stability in prefixes that are announced. In general, an autonomous system will maintain a path to an aggregate supernet 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 autonomous system.

Unfortunately, portions of the Internet address space are not well-aggregated and contain considerably more routes than theoretically necessary. Although aggregation of a single site, or campus-level network is relatively straightforward, aggregation at a larger scale, including across multiple backbone providers, is considerably more difficult and requires close cooperation between service providers.

Perhaps the largest factor contributing to poor aggregation is the increasing trend towards multi-homing of customer end-sites [6]. Since the multi-homed customer prefixes require global visibility, it is problematic for these addresses to be aggregated into larger supernets. In addition, the lack of hierarchical allocation of the early, pre-CIDR IP address space exacerbates the current poor level of aggregation. Prior to the introduction of RFC-1338, most customer sites obtained address space directly from the Internic instead of from their provider's CIDR block. Similarly, the technical difficulties and associated reluctance of customer networks to renumber IP addresses when selecting a new service provider contribute to the number of unaggregated addresses.

The suboptimal aggregation of Internet address space has resulted in large number of globally visible addresses. More significantly, many of these globally visible prefixes are reachable via one or more paths. We would expect Internet instability to be proportional to the total number of available paths to all of the globally visible network addresses or aggregates. Analysis of our experimentally collected BGP data has revealed significantly more BGP updates than we originally anticipated. The Internet "default-free" routing tables currently contain approximately 45,000 prefixes with 1,500 unique ASPATHs interconnecting 1,300 different autonomous systems [8]. As shown later in this paper, instability is well-distributed over destination prefixes, peer routers, and origin autonomous system space. In other words, no single prefix or path dominates the routing statistics or contributes a disproportionate amount of BGP updates. Thus, we would expect that instability should be proportional to the 1,500 paths and 45,000 prefixes, or substantially less than the three to six million updates per day we currently observe.

The majority of these millions of unexpected updates, however, may not reflect legitimate changes in network topology. Instead, our study has shown that the majority of inter-domain routing information consists of pathological updates. Specific examples of these pathologies include: repeated, duplicate withdrawal announcements (WWDup), oscillating reachability announcements (WADup), and duplicate path announcements (AADup). Figure 2 shows the relative distribution of each class of instability over a seven month period. For the clarity and simplification of the following discussions, we have excluded WWDup from Figure 2 so as not to obscure the salient features of the other data. The breakdown of instability categories shows that both the AADup and WADup classifications consistently dominate other categories of routing instability. The relative magnitude of AADup updates was unexpected. Closer analysis has shown that the AADup category is dominated by policy changes that do not directly affect forwarding instability and will be the topic of future work. Only a small portion of the BGP updates (AADiff, WADiff) each day may directly reflect possible exogenous network events, such as router failures and leased line disconnectivity. In Section 6, we discuss the impact of the pathological updates on In-

ternet infrastructure. In general, the repeated transmission of these pathological updates is a suboptimal use of critical Internet infrastructure resources.

Network	Announce	Withdraw	Unique
Provider A	1127	23276	4344
Provider B	0	36776	8424
Provider C	32	10	12
Provider D	63	171	28
Provider E	1350	1351	8
Provider F	11	86417	12435
Provider G	2	61780	10659
Provider H	21197	77931	14030
Provider I	259	2479023	14112
Provider J	2335	1363	853

Table 1: Partial list of update totals per ISP on February 1, 1997 at AADS. This data is representative of daily routing update totals. These totals should not be interpreted as performance of particular backbone provider. Data may be more reflective of a provider's customers and the relative quality of address aggregation.

Analysis of nine months of BGP traffic indicates that the majority of BGP updates consist entirely of pathological, duplicate withdrawals (WWDup). Most of these WWDup withdrawals are transmitted by routers belonging to autonomous systems that never previously announced reachability for the withdrawn prefixes. On average, we observe between 500,000 to 6 million pathological withdrawals per day being exchanged at the Mae-East exchange point. As Table 1 illustrates, many of the exchange point routers withdraw an order of magnitude more routes then they announce during a given day. For example, Table 1 shows that ISP-I announced 259 prefixes, but transmitted over 2.4 *million* withdrawals for just 14,112 different prefixes.

The 2.4 million updates illustrates an important property of inter-domain routing – the disproportionate effect that a single service provider can have on the global routing mesh. Our analysis of the data shows that all *pathological routing incidents* were caused by small service providers. We define a pathological routing incident as a time when the aggregate level of routing instability seen at an exchange point exceeds the normal level of instability by one or more orders of magnitude. Further interaction with these providers has revealed several types of problems including misconfigured routers, and faulty new hardware/software in their infrastructure.

Our data also indicates that not all service providers exhibit this pathological behavior. Empirical observations show that there is a strong causal relationship between the manufacturer of a router used by an ISP and the level of pathological BGP behavior exhibited by the ISP. For example, in a particular case, we observed that before a large service provider's transition to a backbone infrastructure based on particular router, the service provider exhibited well-behaved routing. Immediately following the transition, the service provider began demonstrating pathological behavior similar to behaviors described previously.

Our analysis of the data also indicates that routing updates have a regular, specific periodicity. We have found that most of these updates demonstrate a periodicity of either 30 or 60 seconds, as discussed below. We define the *persistence* of instability and pathologies as the duration of time routing information fluctuates before it stabilizes. Our data indicate that the persistence of most pathological BGP behaviors are under five minutes. This short-lived pathological behavior suggests some type of delay in convergence between inter-domain BGP routers, or multiple IGP/EGP

routing protocols operating within an autonomous system.

4.2 Possible Origins of Routing Pathologies

Our analysis indicates that a small portion of the extraneous, pathological withdrawals may be attributable to a specific router vendor's implementation decisions. In particular, one Internet router vendor has made a time-space trade-off implementation decision in their routers not to maintain state on the information advertised to the router's BGP peers. Upon receipt of any topology change, these routers will transmit announcements or withdrawals to all BGP peers regardless of whether they had previously sent the peer an announcement for the route. Withdrawals are sent for every explicitly and implicitly withdrawn prefix. We will subsequently refer to this implementation as *stateless BGP*. At each public exchange point, this stateless BGP implementation may contribute an additional $O(N * U)$ updates for each legitimate change in topology, where N is the number of peer routers and U is the number of updates. It is important to note that the stateless BGP implementation is compliant with the current IETF BGP standard [12]. Several products from other router vendors do maintain knowledge of the information transmitted to BGP peers and will only transmit updates when topology changes affect a route between the local and peer routers. After the initial presentation of our results [10], the vendor responsible for the stateless BGP implementation updated their router operating software to maintain partial state on BGP advertisements. Several ISPs have now begun deploying the updated software on their backbone routers. Preliminary results after deployment of this new software indicate that it limits distribution of WWDup updates. As we describe below, although the software update may be effective in masking WWDup behavior, it does not explain the origins of the oscillating WWDup behavior.

Overall, our study indicates that the stateless BGP implementation by itself contributes an insignificant number of additional updates to the global routing mesh. Specifically, the stateless BGP implementation does not account for the oscillating behavior of WWDup, and AADup updates. In the case of a single-homed customer and a number of stateless peer routers, every legitimate announce-withdrawal sequence should result in at most $O(N)$ updates at the exchange point, where N is the number of peers. Instead, empirical evidence suggests that each legitimate withdrawal may induce some type of short-lived pathological network oscillation. We have observed that the persistence of these updates is between one and five minutes.

In general, Internet routing instability remains poorly understood and there is no consensus among the research and engineering communities on the characterization or significance of many of the behaviors we observed. Researchers and the members of the North American Network Operators Group (NANOG) have suggested a number of plausible explanations for the periodic behavior, including: CSU timer problems, misconfigured interaction of IGP/BGP protocols, router vendor software bugs, timer problems, and self-synchronization.

Most Internet leased lines (T1, T3) use a type of broadband modem referred to as a Channel Service Units (CSU). Misconfigured CSUs may have clocks which derive from different sources. The drift between two clock sources can cause the line to oscillate between periods of normal service and corrupted data. Unlike telephone customers, router interface cards are sensitive to millisecond loss of line carrier
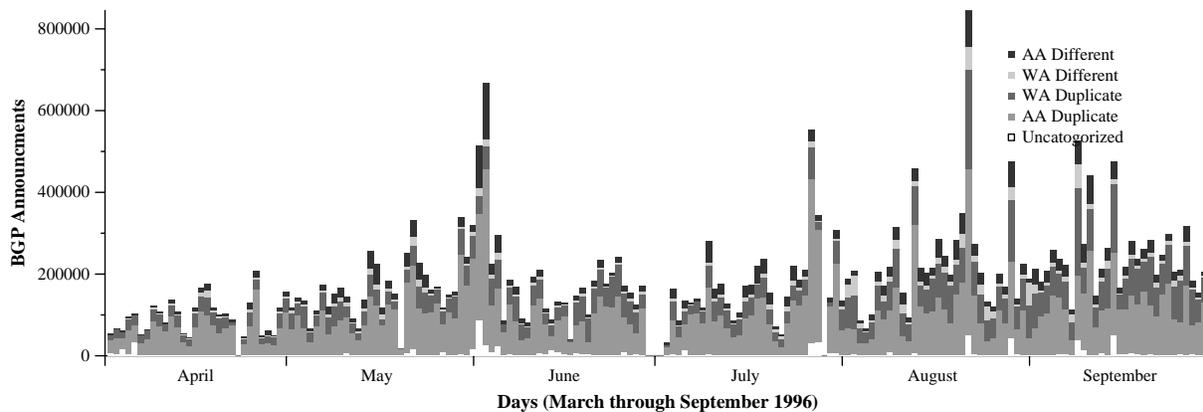
Figure 2: Breakdown of Mae-East routing updates from April through September 1996.

and will flag the link as down. If these CSU problems are widespread, the resulting link oscillation may contribute a significant number of the periodic BGP route withdrawals and announcements we describe.

Another possible explanation involves a popular router vendor's inclusion of an unjittered 30 second interval timer on BGP's update processing. Most BGP implementations use a small, jittered timer to coalesce multiple outbound routing updates into a single BGP update message in order to reduce protocol processing overhead on the receiving peer [11]. The combination of this timer and a stateless BGP implementation may introduce some unintended side-effects. Specifically, we examine the sequence of an announcement for a prefix with ASPATH A1, followed by an announcement (and subsequent implicit withdrawal for A1) for the prefix with ASPATH A2, followed by a re-announcement of the prefix with ASPATH A1. If the sequence A1,A2,A1 occurs within the expiration of the timer interval, the routing software may flag the route as changed and transmit a duplicate route announcement at the end of the interval. A similar sequence of events for the availability of a route, W,A,W, could account for WWDup behavior of some routers. Overall, the 30 second interval timer may be acting as an artificial route dampening mechanism, and as such, the WWDup and AADup behavior may be masking real instability. We will discuss the implication and effects of redundant BGP updates and pathological behavior more in Section 5.

Unjittered timers in a router may also lead to *self synchronization*. In [5], Floyd and Jacobson describe a means by which an initially unsynchronized system of apparently independent routers may inadvertently become synchronized. In the Internet, the unjittered BGP interval timer used on a large number of inter-domain border routers may introduce a weak coupling between those routers through the periodic transmission of the BGP updates. Our analysis suggests that these Internet routers will fulfill the requirements of the Periodic Message model [5] and may undergo abrupt synchronization. This synchronization would result in a large number of BGP routers transmitting updates simultaneously. Floyd and Jacobson describe self-synchronization behavior with Decnet DNA protocol, the Cisco IGRP protocol, and the RIP1 protocol on the NSFNet backbone. The simultaneous transmission of updates has the potential to overwhelm the processing capacity of recipient routers and lead to periodic link or router failures. We have discussed the possibility of self-synchronization with router vendors and are exploring the validity of this conjecture.

Another plausible explanation for the source of the periodic routing instability may be the improper configuration of the interaction between interior gateway protocols and BGP. The injection of routes from IGP protocols, such as OSPF, into BGP, and vice versa, requires a complex, and often mishandled, filtering of prefixes. Since the conversion between protocols is lossy, path information (e.g., ASPATH) is not preserved across protocols and routers will not be able to detect an inter-protocol routing update oscillation. This type of interaction is highly suspect as most IGP protocols utilize internal timers based on some multiple of 30 seconds. We are working closely with router vendors and backbone providers on an ongoing analysis of these interactions.

As described earlier, Varadhan et al. [21] show that unconstrained routing policies can lead to persistent route oscillations. Only the severely restrictive shortest-path route selection algorithm is provably safe. Since the end of the NSFNet, routing policies have been growing in size and complexity. As the number of peering arrangements and the topological complexity of the Internet continue to grow, the potential for developing persistent route oscillation increases. We note, however, that there have been no known reports to date of persistent route oscillation occurring in operational networks. The evaluation and characterization of potentially dangerous unconstrained policies remains an open issue currently being investigated by several research groups.

5 Analysis of Instability

In the previous section we explored characteristics of pathological routing behavior. In this section, we focus on the trends and characteristics of both forwarding instability and route policy fluctuation. The remainder of this discussion presents routing statistics collected at the Mae-East exchange point. It is important to note that these results are representative of other exchange points, including PacBell and Sprint.

5.1 Instability Density

Ignoring attribute changes and pathological traffic (AADup and WWDup) we examined the remaining BGP updates for any overall patterns and trends. Figure 3 represents Internet routing instability for a seven month period. This instability is measured as the sum of AADiff, WADiff, and WADup updates seen during the day for seven months. Each day is

represented by a vertical slice of small squares, each of which represent a ten minute aggregate of instability updates. The black squares represent a level of instability above a certain threshold; the light-gray squares a level below; and the white squares represent times for which data is not available. Additionally, the horizontal axis has a raised indentation that represents weekends. The raw data were detrended using a least-square regression – routing instability increased linearly during the seven month period. Moreover, because we were looking for gross trends, the magnitude of the difference between minimal and maximal instability was reduced by examining the logarithm of this detrended data. Figure 3 represents the modified data. The threshold was chosen as a point above the mean of the modified data, and as such represents a significant level of raw updates that varies depending on the date. The values for the threshold correspond to a raw update rate from 345 updates per 10 minute aggregate in April to 770 updates in October.

Figure 4: Representative week of raw forwarding instability updates (August 3 through 9, 1996) aggregated at ten minute intervals.
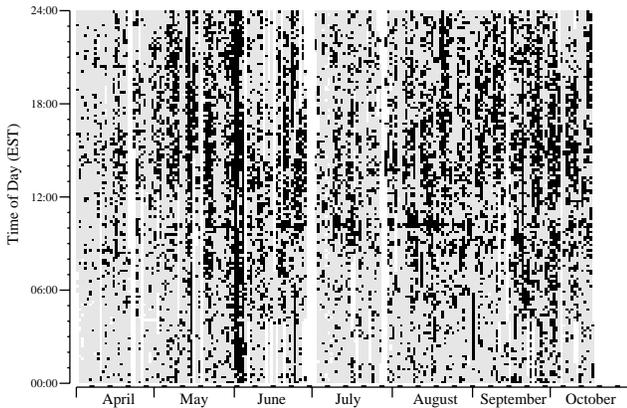

Figure 3: Internet forwarding instability density measured at the Mae-East exchange point during 1996.

Figure 3 shows several interesting phenomena. The bottom of the graph represents midnight EST for each given day. Notice that during the hours of midnight EST (9:00pm PST) to 6:00am EST there are significantly fewer updates than during the rest of the day; the updates appear to be heaviest during North American network usage hours. In particular, from noon to midnight are the densest hours. The second major trend is represented by vertical stripes of less instability (light gray) that correspond to weekends. Perhaps the most striking visual pattern that emerges from the graph are the bold vertical lines at the end of May and beginning of June. These represent the state of the Internet during a major ISP's infrastructure upgrade. Some networks experienced especially high levels of congestion, disconnectivity, and latency during this period. Another interesting pattern is the horizontal line of dense updates at approximately 10:00am (7:00am PST). This line represents large spikes of raw updates that are consistently measured. A plausible explanation for this localized density is that this time may correspond to backbone maintenance windows. Finally, notice that the updates measured during June, July and early August from about 5:00pm to midnight are sparser than those times in May and late August and September. This may represent summer vacation at most of the educational hosts in the Internet, and reflects a pattern closer to the usage of business.

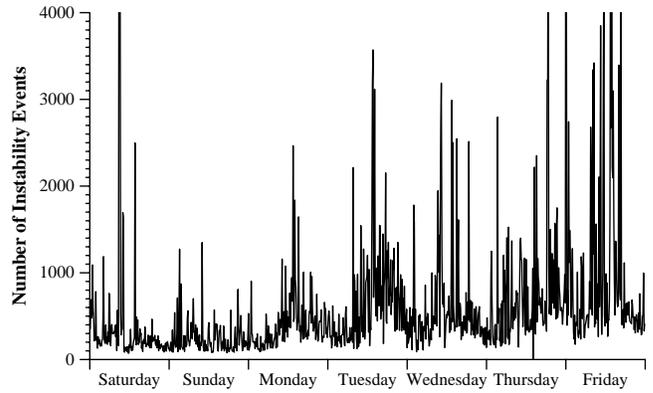The week of routing updates represented in figure 4 provides a representative display of the general trends over a week. From the data there appears to be a bell-shaped curve of raw updates that peaks during the afternoon. Similarly, there is relatively little instability during the weekend. The exception is Saturday's spike. Saturdays often have high amounts of temporally localized instability. We have no immediate explanation for this occurrence.

A more rigorous approach to identifying temporal trends in the the routing updates was undertaken using time series analysis. Specifically, the modified data represented in figure 3 were analyzed using spectrum analysis. The data from August through September were used due to their completeness. Again, these detrended data were ideal for harmonic analysis having been filtered in a manner similar to the treatment of Beverage's wheat prices by Bloomfield in [1]. 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.
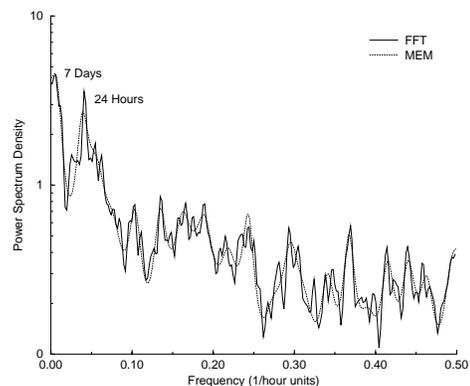

Figure 5: Results from time series analysis of the Internet forwarding instability updates measured at the Mae-East exchange point during August and September 1996 using hourly aggregates.

Figure 5 shows a correlogram of the data generated by two techniques: a traditional fast Fourier transform (FFT) of the autocorrelation function of the data; and maximum-

entropy (MEM) spectral estimation. These two approaches differ in their estimation methods, and provide a mechanism for validation of results. They both find significant frequencies at seven days, and 24 hours. These confirm the visual trends identified in figures 3 and 4.

It is somewhat surprising that the measured routing instability corresponds so closely to the trends seen in Internet bandwidth usage [15] and packet loss. As to the causality of these phenomena, we can only offer suppositions. With a high level of packet loss and a significant rate of BGP updates, keep-alive messages can become delayed long enough to drop BGP connections between peering routers. The specific levels of update load and congestion necessary to sever these connections vary depending on the routing technology in place. Once a BGP connection is severed, all of the peer's routes are withdrawn. An alternate explanation is that this cycle is due to Internet engineering activity that occurs within a business day. However, the data seem to indicate that a significant level of instability remains until late evening, correlating more with Internet usage than engineering maintenance hours. While the relationship between network usage and routing instability may seem intuitively obvious to some, a more rigorous justification is problematic due to the size and heterogeneity of the Internet. We are continuing to investigate this relationship in our current work [8].

5.2 Fine-grained Instability Statistics

Having examined aggregate instability statistics, we now analyze the data at a finer granularity: autonomous system and route contributions. To simplify the following presentation, we focus on a single month of instability, August 1996, measured at the Mae-East exchange point. This month was chosen since it typifies the results seen at the other exchange points across our measurements. Specifically, we show that:

- **No** single autonomous system consistently dominates the instability statistics.

- There is **not** a correlation between the size of an AS (measured at the public exchange point as the number of routes which it announces to non-customer and non-transit peers) and its proportion of the instability statistics.

- A small set of paths or prefixes do **not** dominate the instability statistics; instability is evenly distributed across routes.

The graphs in figure 6 break down the routing updates seen during August measured in each of the route server's peers. Three update categories (AADiff, WADiff, and WADup) are shown where points represent the normalized number of updates announced by a peer on a specific day. That is, there is a point for every peer for every day in August. The horizontal axes show the proportion of the Internet's default-free routing table for which the peer is responsible on a specific day; the vertical axes signify the proportion of that day's route updates that the peer generated. The diagonal represents the break-even points: where a peer generates a proportion of announcements equal to its responsibility for routes in the routing table. If routing updates were equally distributed across all routes, we would expect to see autonomous systems generating them at a rate equal to their share of the routing table. Generally, we do not see that: few days cluster about the line which indicates that there is

not a correlation between the size of an AS, and its share of the update statistics.

The Internet routing tables are dominated by six to eight ISPs. These ISPs represent the clusters of points highlighted in figure 6a. Over the course of the month, their share of the default-free routing tables did not change significantly. Over the course of our analysis no single ISP consistently contributes disproportionately to the measured instability in all three categories. The exception, shown in the figures, is *ISP-E* which during August was going through an infrastructure transition. While it is not characteristic of *ISP-E*'s behavior for every month, it was characteristic of our analysis that at least one of the major ISPs was going through an infrastructure change at any given point in time. Some autonomous systems always represent a somewhat larger share of instability, but this may be explained by a large number of factors. For example, *ISP-A* provides connectivity to a large number of international networks; *ISP-B* is a relatively new ISP that has a much younger customer base and has been able to provide address space from under its own set of aggregated CIDR blocks, perhaps hiding internal instability through better aggregation. Additional factors that can skew ISP behavior include: customer behavior, routing policies, and quality of aggregation.

We now focus on the instability on a per-route basis. Specifically, we look at the instability measured at the Mae-East exchange point during August for (prefix, AS-peer) pairs, or *Prefix+AS*. A Prefix+AS represents a set of routes that an AS announces for a given destination. It is more specific than a prefix since the same prefix could be reached through several ASes; and more general than a route which uniquely specifies the ASPATH. By aggregating routing updates on Prefix+AS pairs, we can pinpoint several routing update phenomena: updates that oscillate over several routes for a given prefix; AS contribution for given prefix; and prefix behavior.

Figure 7 shows the cumulative distribution of Prefix+AS instability for the four BGP announcement categories. In all four graphs, the horizontal axes represent the number of Prefix+AS pairs that exhibited a specific number of BGP instability events; the vertical axes show the cumulative proportion of all such events. The graphs contain lines that represent daily cumulative distributions for August 1996. Examining these graphs, one can see that from 80 to 100 percent of the daily instability is contributed by Prefix+AS pairs announced less than fifty times. For example, figure 7a shows that depending on the day, from 20 to 90 percent (median of approximately 75%) of the AADiff events are contributed by routes that changed ten times or less. Together, these graphs show that no single route consistently dominates the instability measured at the exchange point. However, there are days where a single Prefix+AS pair contributes substantially, such as August 11, a day where several prefix+AS pairs contributed about 40% of the daily aggregate AADiffs, graphically displayed as the lowest curve in figure 7a. Specifically, in this example, *ISP-A* announced seven routes each between 630 and 650 times. These same seven routes had an equal amount of AADups that day and also account for the low curve in figure 7c. Moreover, there are zero withdrawals on these seven prefixes.

When comparing the four types of routing updates in figure 7, one can see that WADiff climbs to a plateau of about 95% faster than the other three categories. WADiff also has the fewest number of Prefix+AS pairs that dominate their days. In fact, there are very few days where a Prefix+AS has more than 100 WADiff events. Similarly, there are very few

(a) AS AADiff Contribution

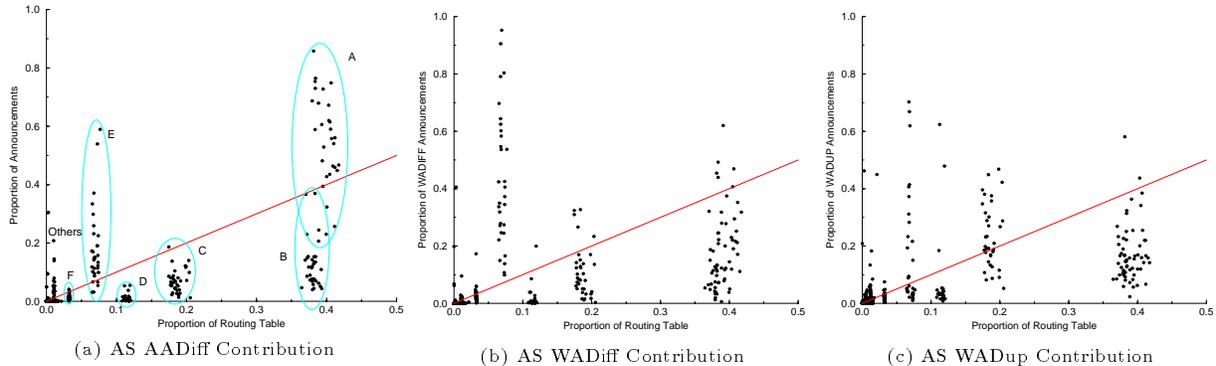(b) AS WADiff Contribution

(c) AS WADup Contribution

Figure 6: AS contribution to routing updates measured at the Mae-East exchange point during August 1996. These graphs measure the relative level of routing updates generated by backbone providers. This data does not represent relative performance of ISPs, and may be more reflective of customer instability and address allocation policies.

(a) AADiff

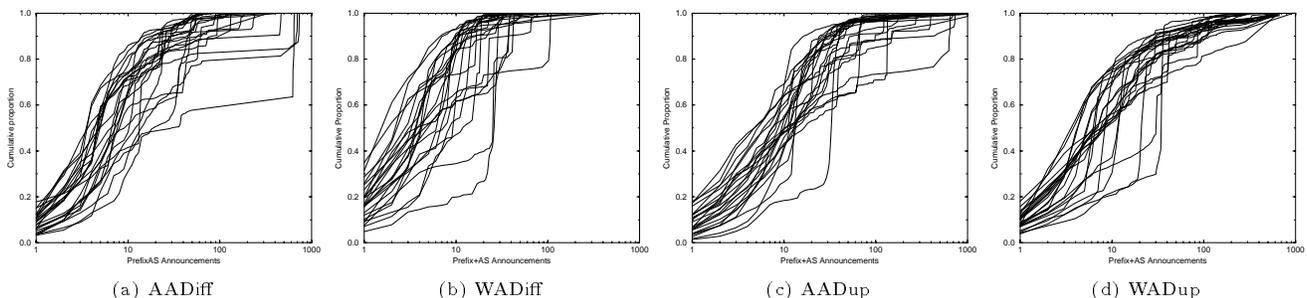(b) WADiff

(c) AADup

(d) WADup

Figure 7: Cumulative distribution of Prefix+AS routing updates measured at the Mae-East exchange point during August 1996. Each line in a graph represents the update distribution for a single day.

days where a Prefix+AS sees more than 200 AADiff events. Taken together, this information is comforting since these categories perhaps best represent actual topological instability. In contrast, the categories than may represent redundant instability information, AADup and WADup, both have a significant number of days where from 5% to 10% of their events come from Prefix+AS pairs that occur 200 times or more. An investigation of instability aggregated on prefix alone generated results similar to those shown in this section and have been omitted.

5.3 Temporal Properties of Instability Statistics

We next turn our attention to the temporal properties of Internet routing instability. Section 5.1 described the aggregate temporal behavior and identified the weekly and daily frequencies. Here we investigate the frequency distributions for instability events at the Prefix+AS level. Again our analysis looks at the statistics from August 1996 measured at the Mae-East exchange point. 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 [6] and end-to-end techniques [17]. Our measurement apparatus allowed a unique opportunity

to examine the high frequency components. Our results are shown in figure 8. The graphs in figure 8 represents a histogram distribution for each of the four instability categories. The graphs' horizontal axes mark the histogram bins in a log-time scale that ranges from one second (1s) to one day (24h); the vertical axes show the proportion of updates contained in the histogram bins. The data shown in these graphs take the form of a modified box plot: the black dot represents the median proportion for all the days for each event bin; the vertical line below the dot contains the first quartile of daily proportions for the bin; and the line above the dot represents the fourth quartile.

As illustrated figure 8, the predominant frequencies in each of the graphs are captured by the thirty second and one minute bins. The fact that these frequencies account for half of the measured statistics was surprising. Normally one would expect a 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. The thirty second periodicity suggests some wide-spread, systematic influence in the origin, or on the flow of instability information. There are several possible causes for this periodicity including routing software timers, self synchronization, and routing loops. The presence of these frequencies in the more *legitimate* instability categories, such as WADiff and AADiff almost certainly represents some pathology which may be caused by CSU handshaking timeouts on leased lines or a flaw in the

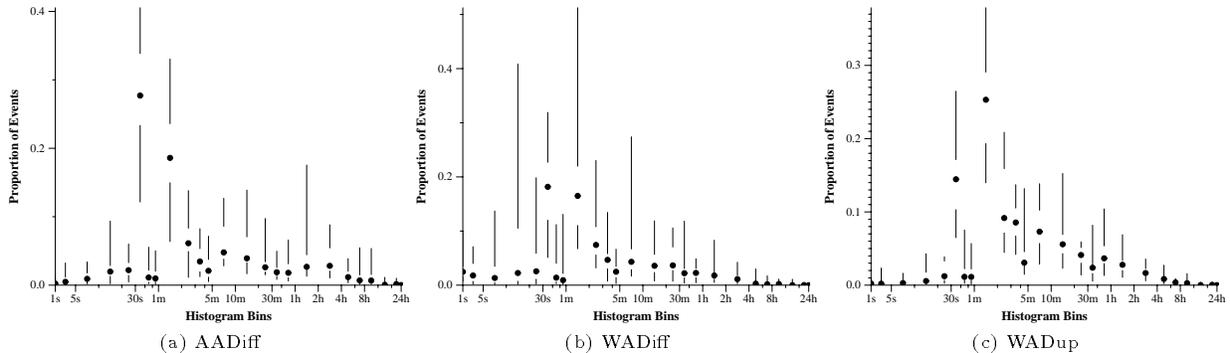(a) AADiff   (b) WADiff   (c) WADup

Figure 8: Histogram distribution of update inter-arrival time distances for Prefix+AS instability measured at the Mae-East exchange point during August 1996.

routing protocols.

6 Impact of Routing Instability and Conclusion

As we described earlier, forwarding instability can have a significant deleterious impact on the Internet infrastructure. Instability that reflects real topological changes can lead to increased packet loss, delay in time for network convergence, and additional memory/CPU overhead on routers. In the current Internet, network operators routinely report backbone outages and other significant network problems directly related to the occurrence of route flaps [14].

Our analysis in this paper demonstrated that the majority (99 percent) of routing information is pathological and may not reflect real network topological changes. We defined a taxonomy for discussing routing information and suggested a number of plausible explanations that may account for some of the anomalous behaviors. Router vendors and ISPs are currently proceeding with the deployment of updated routing software to correct some of the potential problems we described.

Since pathological, or redundant, routing information does not affect a router's forwarding tables or cache, the overall impact of this phenomena may be relatively benign and may not substantially impact a router's performance. Most of the pathological updates will be quickly discarded by routers and will not undergo policy evaluation. More importantly, these pathological updates will not trigger router cache churn and the resultant cache misses and subsequent packet loss.

A number of network operators, however, believe that the the sheer volume of pathological updates may still be problematic [16]. Even pathological updates require some minimal router resources, including CPU, buffers and the expense of marshaling pathological prefix data into both inbound and outbound packets. Informal experiments with several popular routers suggest that sufficiently high rates of pathological updates (e.g. 300 updates per second) are enough to crash a widely deployed, high-end model of commercial Internet router. We define crash as a state in which the router is completely unresponsive and does not respond to future routing protocol messages, or console interrupts. Other studies have reported high CPU consumption and loss of peering sessions at moderate rates of routing instability. Although our analysis of the impact of redundant information on Internet performance is still ongoing, we believe pathological updates are a suboptimal use of Internet

resources.

Our analysis of the data showed that instability is well distributed across both autonomous systems and prefix space. More succinctly, no single service provider or set of network destinations appears to be at fault. We described a strong correlation between the version and manufacturer of a router used by an ISP and the level of pathological behavior exhibited by that ISP. As noted earlier, router vendors responded to our finding, and developed software updates to limit several pathologies. Updated software is now actively being deployed by backbone operators. Preliminary results indicate that it will be successful in limiting the flow of some pathologies, particularly those involving WWDup updates.

We also showed that instability and redundant information exhibit strong temporal properties. We describe a strong correlation between the level of routing activity and network usage. The magnitude of routing information exhibits the same significant weekly, daily and holiday cycles as network usage and congestion. Although the relation between instability and congestion may seem intuitive, a formal explanation for this relationship is more difficult.

Instability and redundant routing information also exhibit a strong periodicity. Specifically, we described 30 and 60 second periodicity in both instability and redundant BGP information. We offered a number of plausible explanations for this phenomena, including: self-synchronization, misconfiguration of IGP/BGP interactions, router software problems, and CSU link oscillation. The origins of this periodic phenomena, however, remain an open question.

If we ignore the impact of redundant updates and other pathological behaviors, Figure 9 shows that most (80 percent) of Internet routes exhibit a relatively high level of stability. Only between 3 and 10 percent of routes exhibit one or more WADiff per day, and between 5 and 20 percent exhibit one or more AADiff each day. This conforms with empirical observations by most end-users that the Internet usually seems to work. Our data also agrees with Paxson's findings that only a very small fraction of routes exhibit some type of topological instability each day[17].

One of our difficulties in evaluating the impact of instability on Internet performance is that we have not yet fully been able to characterize and understand the significance of the different classes of routing information. Figure 9 shows that between 35 and 100 percent (50 percent median) of prefix+AS tuples are involved in at least one category of routing update (policy fluctuation, forwarding instability, pathological information) each day. Specifically, we do not
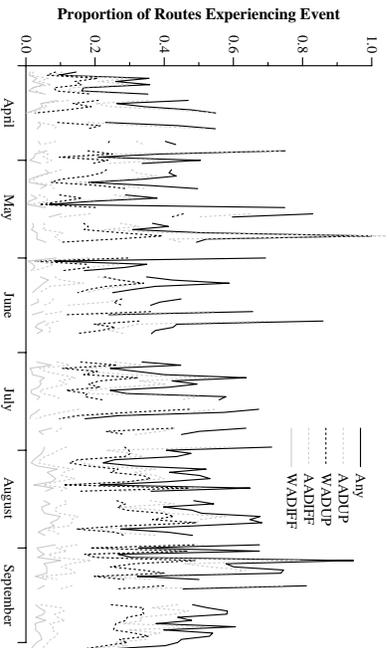
11

Figure 9: Proportion of Internet Routes affected by routing updates (1996). Days shown have at least 80 percent of the date's data collected.

know what percentage of redundant updates may actually be reflective of "legitimate" changes in forwarding information. As we described earlier, some of our analysis suggests that a portion of the AADup and WWDup behaviors may originate in the interaction between forwarding instability and the 30 second interval timer on some routers. If this is the case, then some portion of pathological behavior may reflect legitimate topological changes.

By directly measuring the BGP information shared by Internet Service Providers at several major exchange points, this paper identified several important trends and anomalies in inter-domain routing behavior. This work in conjunction with several other research efforts has begun to examine inter-domain routing 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

References

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

[2] H.-W. Braun, P. S. Ford and Y. Rekhter, "CIDR and the Evolution of the Internet," SDSC Report GA-A21364, in Proceedings of INET'93, Republished in ConneXions Sep 1993 (InterOp93 version).

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

[4] D. Estrin, Y. Rekhter, and S. Hotz, "A Scalable Inter-domain Routing Architecture," in Proceedings of the ACM SIGCOMM '92, pp. 40-52, Baltimore, MD, August, 1992.

[5] 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.

[6] 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.

[7] J. Honig, D. Katz, M. Mathis, Y. Rekhter, and J. Yu, "Application of the Border Gateway Protocol in the Internet," RFC-1164, June 1990.

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

[9] C. Labovitz, "Multithreaded Routing Toolkit," Merit Technical Report, 1996.

[10] C. Labovitz, NANOG presentation, Washington, D.C., May 1996.

[11] Dave O'Leary, Cisco Systems, Inc. Private communication, January 1997.

[12] K. Lougheed and Y. Rekhter, "A Border Gateway Protocol (BGP)," RFC-1163 June 1990.

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

[14] Merit Joint Technical Staff mail archives, http://www.merit.edu/mjts/msg00078.html.

[15] MFS Communications Mae-East Statistics Page, http://www.mfst.com/MAE/east.stats.html.

[16] North American Network Operators Group, http://www.nanog.org.

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

[18] Y. Rekhter, "Scalable Support for Multi-homed Multi-Provider Connectivity," NANOG, Ann Arbor, MI. October 1996.

[19] Y. Rekhter and C. Topolcic, "Exchanging Routing Information Across Provider Boundaries in the CIDR Environment," RFC-1520. September 1993.

[20] Routing Arbiter web pages, http://www.ra.net.

[21] K. Varadhan, R. Govindan, and D. Estrin, "Persistent Route Oscillations in Inter-Domain Routing," USC/ISI, Available at the Routing Arbiter project's homepage at USC/ISI.

[22] C. Villamizar, R. Chandra, and R. Govindan, "draft-ietf-idr-route-dampen-00-preview", Internet Engineering Task Force Draft, July 21, 1995.

[23] C. Villamizar, "TCP Response Under Loss Conditions", NANOG Presentation, San Francisco, February 1997.