

Joint Analysis of Network Incidents and Intradomain Routing Changes

Amelie Medem, Renata Teixeira, Nick Feamster and Mickael Meulle
UPMC Paris Universit s and CNRS, Georgia Tech, Orange Labs R&D

Abstract—This paper studies how intradomain routing instability relates to events in network trouble tickets for two networks: a VPN provider and the Internet2 backbone network. Our goal in performing this *joint analysis* of routing and trouble tickets is to better understand the likely underlying causes of intradomain routing instability. We develop a method to correlate trouble tickets with instability events and find that, although unplanned events last longer than scheduled maintenance, there is no single underlying cause for most instability, and that these causes differ across networks. In comparison to a similar study from Labovitz *et al.* from ten years ago, we find that, while certain causes of instability such as maintenance and circuit problems remain significant, power issues have become much less prevalent, and software-related problems have become more common.

I. INTRODUCTION

Networks must recover quickly from faults that result from software bugs, hardware failures, misconfigurations, and other causes. Various work has attempted to improve network uptime by developing methods to quickly route around failures when faults occur [17], to prevent or help debug router misconfigurations [5], [11], and to prevent software bugs from occurring [3]. Ultimately, designing protocols to most effectively reduce downtime requires a deeper understanding of the *causes* of faults that induce routing changes. The goal of this paper is to try to better understand the causes of these routing changes in two different networks.

Determining the underlying causes of intradomain routing changes is challenging because these messages themselves contain no hints as to the underlying causes of the changes. The early work by Labovitz *et al.* [10] explains routing failures in one large regional Internet backbone; this paper extends that study by performing a similar analysis on more recent data and for two types of networks. Other previous studies have used routing-message logs [2], [7], [10], [12], [16], [18] or end-to-end measurements [4], [15] to characterize the duration and ultimate effect of network incidents, but they do *not* attempt to explain the cause of these instabilities. Feamster *et al.* [5] and Huang *et al.* [7] use trouble tickets logs to understand the causes of routing changes, but their work largely focused on real-time diagnosis and detection; they did not characterize the nature and causes of the observed instabilities. Both researchers and operators still need a better understanding of the underlying causes of these changes.

This paper attempts to identify the underlying causes of faults that induce routing changes within two networks—a large Virtual Private Network (VPN) backbone; and Internet2 (formerly Abilene), the US research backbone. To identify the

underlying causes of these routing changes, we correlate IS-IS routing messages from these networks with corresponding trouble tickets (i.e., text documents that document incidents affecting network uptime). *Our analysis assumes that intradomain routing changes that correlate with a trouble ticket both in time and in space (i.e., location in the network) can be explained by that trouble ticket.* Correlating these two data sources provides a better understanding of (1) how a particular incident that induces a trouble ticket can appear in IS-IS routing; (2) the possible underlying causes of any observed IS-IS routing changes.

Despite the availability of both routing data and trouble-ticket data, *correlating* these two data sets is challenging. Ideally, to perform such correlation, one would like to have events in these two data sets coincide if the two events were related; unfortunately, the nature of an “event” in each data set differs. In IS-IS routing, an event consists of repeated “link flaps” or routing changes; each change is short in duration, but a single event may comprise many routing messages. The challenge is aggregating these routing messages into a single routing event. In contrast, trouble-ticket data consists of “tickets” that humans manually enter into a ticketing system. Each ticket has a start and end time; unfortunately, due to the manual, human nature of how these tickets are created and resolved, the start and end time on these tickets may not be accurate. Thus, the challenge with correlating these two timeseries boils down to correlating a noisy timeseries with accurate timing but no information about causes (i.e., the IS-IS data) with a coarse timeseries with inaccurate timing but rich causal information (i.e., the trouble tickets). To solve this problem, we develop a correlation algorithm, which we describe in more detail in Section IV.

Our analysis (Section V) reveals the following findings:

- The primary cause of routing instability in Internet2 and the VPN provider are quite different. The largest fraction of Internet2’s routing changes result from scheduled maintenance, whereas the majority of the VPN provider instabilities are unplanned.
- Unplanned events last longer than scheduled maintenance.
- Hardware and circuit problems cause more downtime than software-related problems, even though software problems cause approximately one-third of routing instabilities for both networks.
- In comparison to the study of Labovitz *et al.* [10] in 1999, power-related problems are less prevalent, but software-

TABLE I
AN EXAMPLE TROUBLE TICKET.

Report time: Wed, 17 Jun 2005 20:10:09 GMT Subject: Internet2 IP Network Core Node ATLA Maintenance Affected: Core Node ATLA and its Connectors, Peers and Participants Start time: Sunday, June 21, 2005, 3:00 AM (0300) UTC End time: Sunday, June 21, 2005, 4:00 AM (0400) UTC Description: During the above time, Internet2 engineers will be performing diagnostics on line cards on the router ATLA. This may cause intermittent disruptions on any given network connection. The router itself will not lose connectivity; as line cards are temporarily taken off line, connectivity will be disrupted for a short amount time until the line card is placed back in service. Ticket Identity: 10014

related problems are more prevalent.

II. DATA

This section presents our datasets. We analyze data from two operational IP networks: a VPN provider and Internet2. The VPN provider and Internet2 are highly different in their topology and in the network services they provide. Internet2 interconnects research labs and universities in the United States. The VPN provider is the backbone of a large provider of MPLS VPN service. Internet2 has only 11 backbone routers, whereas the VPN provider counts hundreds of routers all over the country. The VPN provider interconnects with two other networks: one that is the international backbone of the VPN service, and the other that provides commodity Internet connectivity to its customers. Internet2 interconnects with hundreds of peers and customers. We use trouble tickets and IS-IS messages for Sep. 19, 2006 to Apr. 16, 2007 for the VPN provider and Aug. 8, 2004 to Apr. 6, 2007 for Internet2 [1].

Trouble tickets. Network administrators use trouble tickets to track troubleshooting and maintenance activities. Trouble tickets can give useful information on network incidents. In Table I, we present one example of trouble ticket. From this ticket, we learn that there will be a scheduled incident on the interface cards of the router ATLA. Although, the information stored in a ticket may vary across systems, trouble tickets always contain the three pieces of information we need for our study: a free-text part that describes the incident or activity; the affected equipment; and the timing information (when the ticket was opened and closed).

Different networks use different trouble-ticketing systems. Internet2 uses an email system available on the Internet to record trouble tickets [1], whereas the VPN provider uses a private system that stores trouble tickets in a database. In Internet2, each email contains a finished or ongoing description of a ticket. Multiple emails may refer to the same trouble ticket identifier, so we select the latest email with a given identifier to represent the trouble ticket. The latest email often contains the updated description of the problem. From a total of 2,574 emails, we extract 1,613 unique trouble tickets. The database from the VPN provider contains more details about the problem than the email descriptions from Internet2. The database we study contains seven months of trouble tickets of the backbone network operations team. For confidentiality reasons, we cannot disclose the exact numbers of trouble

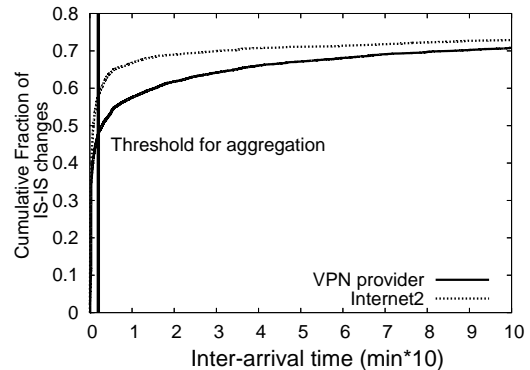


Fig. 1. Time between IS-IS changes per link/prefix.

tickets for the VPN provider, so our analysis only shows percentages for this network.

IS-IS messages. Both the VPN provider and Internet2 use IS-IS as their intradomain routing protocol and deploy monitors to capture IS-IS messages. IS-IS domain topology changes always trigger the flooding of many IS-IS messages all over the network. The IS-IS flooding process ensures that every message reaches every router, including the monitor. The VPN provider collects IS-IS messages from one switch over an IP link, whereas Internet2 deploys one monitor per router. The monitor timestamps every message it receives.

We focus on *IS-IS updates*, which are messages that report changes to the topology of the IS-IS domain. The main IS-IS changes are the following: link down, link up, prefix down, prefix up, metric and overload bit changes. Other IS-IS topology changes also exist (e.g. changes on router names) but we do not consider them in our analysis because they did not appear in our data. During the period of our study, we find 81,487 topology changes in Internet2 and thousands of changes in the VPN provider. For privacy reasons, we do not disclose the exact number of the VPN provider.

III. PREPROCESSING THE DATA

Jointly analyzing IS-IS routing events and trouble tickets is challenging because the notion of an “event” in each dataset is different: the IS-IS data has accurate timestamps but is noisy; the trouble-ticket data, on the other hand, is more coarse, but has inaccurate timestamps. In this section, we describe how we process the trouble tickets and IS-IS data to derive events for each timeseries. We parse trouble tickets to extract a concise and structured description of the network incidents. Then, we describe how we aggregate multiple IS-IS changes to create a single event.

A. Parsing Trouble Tickets

We process the content of trouble ticket fields to extract information that we need for matching tickets with IS-IS messages, plus the root causes of network incidents.

We first extract the root cause for each ticket using the TroubleMiner tool [13]. TroubleMiner uses machine learning techniques to process the free-text description of tickets and extract a concise description of the incident. The type *T.type* of a ticket *T* corresponds to the most relevant words in its

description. For each ticket, T , we set $T.open$ and $T.close$ with the time when T was opened and closed, respectively. The $T.affected$ reflects the names of routers affected by the incident. We process each word in this field and in the ticket description field to extract the names of the routers affected by the trouble ticket.

We focus on *intradomain* routing events, so we ignore the trouble tickets that do not correspond to incidents on internal backbone routers. For Internet2, we ignore tickets reporting incidents that occur on neighboring networks (71.3%) and tickets reporting BGP issues (1.7%). In Internet2, few unplanned tickets happen to share the same affected routers, with an inter-arrival time of less than a day. We manually check these tickets and find that most of them relate to the same incident. Therefore, we merge them into a single event. In the VPN provider, we ignore tickets that correspond to routers that do not run IS-IS (14.4%). In the VPN provider, we also coalesce multiple tickets into a single event because the operators of this network often open one trouble ticket for each router that is affected by a single network failure. We only coalesce trouble tickets that affect the same routers with overlapping time periods and the same description.

B. From IS-IS messages to IS-IS events

This section describes our method for *aggregating* multiple IS-IS messages into a single event. A single network incident can trigger many IS-IS messages. For example, if a link fails both end-points of the link will generate messages at nearly the same time reflecting the same event. We aggregate IS-IS messages into IS-IS events in two steps. First, we group IS-IS messages that happen close in time for a given link or prefix. Second, we aggregate concurrent per-link or per-prefix events that happen close in time on different links or prefixes to create IS-IS events.

For each link, we aggregate IS-IS messages in two sequential steps. First, we **group pairs of related down and up messages**. A “link down” followed by a “link up” typically represents a failure and its recovery, so we consider these two messages as part of the same change. Second, we **aggregate IS-IS changes that occur close in time**. A link can go down and up repeatedly over some time interval; such repeated “flapping” may produce many IS-IS messages with the same underlying cause. The challenge is to determine the time interval to distinguish link flapping for a single event from multiple events.

We choose this threshold experimentally by analyzing the distribution of inter-arrival times of IS-IS changes presented in Figure 1. We crop the x-axis at 100 minutes. This figure shows that many IS-IS changes happen close to the previous change—50% of IS-IS changes in the VPN network follow the previous change by less than 100 seconds; this fraction is even higher (60%) for Internet2. In contrast, the inter-arrival time for some IS-IS changes lasts for days or weeks. We pick the knee of the curves in Figure 1 and select a threshold of 100 seconds (marked by the vertical line). If two IS-IS changes

TABLE II
CARDINALITY OF TROUBLE TICKETS. INCIDENTS IN TICKETS ARE DIFFERENT IN THE VPN PROVIDER AND INTERNET2.

Trouble tickets	VPN provider	Internet2
	100%	435(100%)
Planned	30.1%	324 (74.5%)
<i>Router Software</i>	10.0%	74 (17.0%)
<i>Router Hardware</i>	5%	4 (1%)
<i>Router/Card decommissioned</i>	10%	0 (0%)
<i>Backbone Circuit/Fiber</i>	0%	237 (54.5%)
<i>Other maintenance</i>	4.5%	9 (2.0%)
Unplanned	70%	111 (25.5%)
<i>Router Software bugs</i>	25.2%	8 (1.8%)
<i>Router hardware crash</i>	16.8%	7 (1.6%)
<i>Backbone Circuit/Fiber failures</i>	4.5%	87 (20%)
<i>Power issues</i>	2.8%	3 (0.7%)
<i>Switch issues</i>	2.8%	2 (0.5%)
<i>Configuration error</i>	1.7%	4 (0.9%)
<i>Overload</i>	1.7%	0 (0%)
<i>Other unplanned</i>	14.5%	0 (0%)

happen within 100 seconds, we consider them as part of the same event.

We **aggregate concurrent per-link and per-prefix events that occur close in time**. A single IS-IS event can affect two different links with a given delay. These concurrent events may happen for many reasons. For example, when an interface card crashes, all its links will fail almost at the same time. In their work, Markopoulou *et al.* [12] show that concurrent failures that affect different links in less than 12 seconds reflect the same event. In our work, we use the same threshold to merge IS-IS events on different links or prefixes into the same IS-IS event. At the end of this process, almost 25% and 27.7% of IS-IS events in the VPN provider and Internet2 involve different links or prefixes. The start time of an IS-IS event I , $I.start$, is the minimum of the timestamps for all IS-IS messages in an aggregate; the end time, $I.end$, is the maximum timestamp in the group. The event’s location $I.location$ is a set that contains the identities and names of the affected routers.

IV. JOINT ANALYSIS OF TICKETS AND IS-IS EVENTS

We aim to determine whether an IS-IS event coincides with a trouble ticket. To do so, we need to get a better understanding of the occurrence of network incidents in the VPN provider and Internet2, and a thorough knowledge of network operational practices in both networks. First, we perform a preliminary statistics on trouble tickets alone. Then, we present our methodology to jointly examine tickets and IS-IS events.

A. Preliminary statistics on trouble tickets

We carry out the statistical analysis of different types of trouble tickets, the occurrence of tickets per router and their duration. Our analysis exhibits the following findings.

Different networks show different proportion of incidents. Table II illustrates the types of network incidents that appear in trouble tickets for the VPN provider and Internet2. Different types of networks exhibit the same type of incidents

but in different proportion. Therefore, we may expect not to see the same trend for the causes of routing changes in both networks. About 70% of the incidents in the VPN provider are unplanned; in contrast, roughly the same fraction of incidents in the Internet2 network are due to scheduled incidents such as planned maintenance. In our study, we make a distinction between planned maintenance and emergency maintenance. Network operators scheduled planned maintenance activities and set their time period in advance. Emergency maintenance is an urgent repair of a critical network condition. We analyze emergency maintenance as unexpected incidents. The VPN provider has more failures on routers internally, whereas Internet2 has more failures on circuits or connection. The widely differing incidents might be explained by the fact that these two networks have different goals and connectivity requirements: The VPN provider is a commercial backbone with a large number of customer enterprises and nearly a hundred times more routers than Internet2.

On average, a single router is affected by unexpected failures with an inter-arrival time of more than one day. In the VPN provider, 85% of unplanned trouble tickets only concern a single router during the whole period of analysis. In contrast, in Internet2, each router appears in at least one trouble ticket (On average in 15 tickets). For routers that appear in many trouble tickets, we find that the inter-arrival time between these tickets are always more than a day (with a maximum of 31 days for the VPN provider and 49 days for Internet2).

Trouble tickets last longer in Internet2 than in the VPN provider. This observation is true for planned tickets. Indeed, 80% of the VPN provider’s planned tickets last less than 30 minutes. In Internet2, we count only 12% of planned tickets in the same duration windows. Internet2 counts a lot of backbone circuit related planned tickets (54.5%). Maintenance work on backbone circuits takes more time compared to other planned incidents in the VPN provider. On the contrary, the difference in durations of unplanned tickets in both networks is smaller, which indicates that the VPN provider and Internet2 may take similar time to react to such types of events.

B. Methodology for joint analysis

We use two criterions to jointly analyze trouble tickets and IS-IS events:

- 1) *Location*: Trouble tickets explicitly contain the *location* of network incidents in their description. Given a trouble ticket T , we use the location of the incident, stored in $T.affected$, to correlate it with IS-IS events. We only compare a ticket T with an IS-IS event I if the event impact its affected routers. Therefore, a trouble ticket T and an IS-IS event I coincide in location, if $T.affected$ overlaps with $I.location$. If this condition holds, we then check whether T and I coincide in time.
- 2) *Time*: We compare the times between I and T using two cases: (1) The IS-IS event starts during the trouble ticket ($I.start \geq T.open$ and $I.start \leq T.close$) and (2) The IS-IS event starts before the time period indicated in the

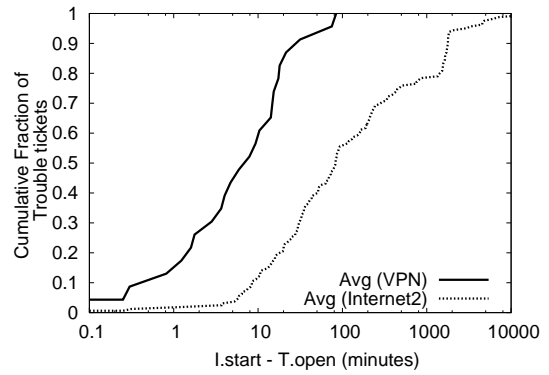


Fig. 2. Average differences between the start time of IS-IS events and the open time of planned tickets (I starts during T).

trouble ticket ($I.start < T.open$). We do not consider the case $I.start > T.close$, where the IS-IS event starts after the closed time of the ticket. Network operators close trouble tickets when they have finished with the resolution of an incident. IS-IS events that start after trouble tickets are likely not to reflect the incidents.

C. Challenges in the correlation

1) *Inexact timestamps in trouble tickets*: Jointly analyzing trouble tickets and IS-IS events in time is challenging because operators open trouble tickets manually. Thus, the time of a ticket can be minutes or hours from that of the corresponding IS-IS events. Operators may timestamp the ticket with the time of a failure as reported in syslog or in some monitoring system. Alternatively, they may enter the planned time for a maintenance activity or the time when a customer reported a failure or an operator detected it. In the former case, the timing of the ticket should be close to that of IS-IS messages (within a few minutes); whereas in the latter we should expect larger delays. To deal with possible inaccuracies in tickets’ time, we adopt a period of observation of *one day* before the trouble ticket. We suppose that one day is a reasonable time for network administrators to have noticed the consequences of a problem inside their network.

2) *Network incidents invisible in IS-IS*: During our period of observation, we find that not all trouble tickets are visible as IS-IS events on their affected routers. Table III presents the cardinality of trouble tickets when they coincide with IS-IS events one day before and during their time period. We see that about 29.6% ($100 - 70.4$) of trouble tickets from the VPN provider and 32% ($100 - 68$) of tickets from the Internet2 did not coincide with any IS-IS events in one day. This finding indicates that studying routing instability alone will not capture all failures. Incidents corresponding to trouble tickets may not be visible in IS-IS routing for a number of reasons: (1) incidents can be masked by lower layers: In the VPN provider, of the previous 35%, 24% affect IP sub-layers interfaces with no IP addresses (e.g., SONET); (2) incidents can affect IP interfaces with no IS-IS; (3) They can happen on backup network equipment; (4) planned incidents can be mitigated with techniques that facilitate “hitless” maintenance,

TABLE III

CARDINALITY OF TROUBLE TICKETS WHEN THEY COINCIDE WITH IS-IS EVENTS. WE CONSIDER IS-IS EVENTS ONE DAY BEFORE AND DURING TICKET PERIOD. EACH PERCENTAGE IS RELATIVE TO THE TOTAL OF TICKETS WITHIN THE SAME CLASS.

Trouble tickets	The VPN provider			Internet2		
	Total	Planned	Unplanned	Total	Planned	Unplanned
Total T	70.4%	33.6%	66.4%	296(68%)	207 (70%)	89(30%)
$I.start \leq T.start$	51.4%	30.5%	62%	78(37.67%)	110(34.1%)	43(40.5%)
$I.start \in [T.start, T.end]$	63.5%	64%	63.4%	256(59.8%)	164(79.2%)	69(65.0%)

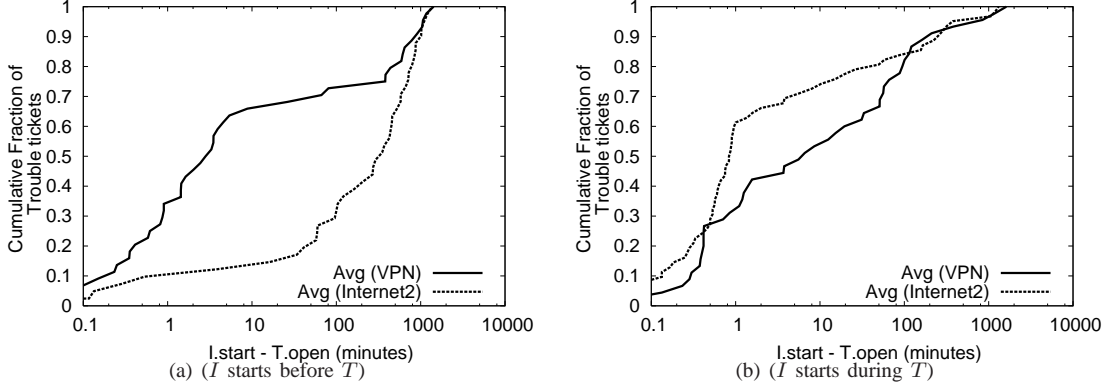


Fig. 3. Differences between the start time of IS-IS events and the open time of unplanned trouble tickets.

(e.g., Cisco Online Insertion); and (5) incidents may not be severe enough to affect routing.

D. Insights from joint analysis

This section examines planned maintenance and unexpected incidents separately because both the VPN provider and Internet2 management these types of events differently.

IS-IS events around planned maintenance activities. If $I.start \geq T.open$ and $I.start \leq T.close$, then we consider that they are correlated. We observe that in both networks, the vast majority of IS-IS events related to planned tickets happen after the ticket is opened. From Table III, 64% and 79.2% of the VPN provider's and Internet2's planned trouble tickets had overlapping IS-IS events. IS-IS events will often happen during the time of a ticket in case of planned maintenance. IS-IS events within planned tickets reflect the maintenance work performed by network administrators. Network operators usually open a ticket before planned maintenance, perform the planned task, and then close the ticket. Knowing this, we expect to see the corresponding IS-IS events within the period of planned trouble tickets.

IS-IS events start closer to planned tickets in the VPN provider than in Internet2. Figure 2 plots the cumulative distribution function (CDF) of the difference $I.start - T.open$ when I starts during T . Each trouble ticket T can coincide with many IS-IS events I (for example, a router crash affects many links), therefore we plot the average of all the differences $I.start - T.open$. On average, for 60% of the VPN provider's planned tickets, an IS-IS event starts within 10 minutes. In Internet2, this happens for only 10% of planned tickets. This difference may be explained by the fact that network administrators in Internet2 compared to the VPN provider,

after opening the ticket, take more time to start doing the maintenance work. Planned trouble tickets last longer in Internet2 than in the VPN provider. Of IS-IS events that start within planned tickets, 96.2% and 97.7% end within the ticket period in the VPN provider and Internet2, respectively. For the remaining IS-IS events, we find that the vast majority last few minutes after the ticket is closed by network administrators.

If $I.start \leq T.start$, we consider that they coincide if I starts few minutes, immediately, before T . We do not expect IS-IS events to correlate before planned maintenance. However, the time in planned ticket is often manual and this case may happen. For example, an administrator may set the time of a maintenance at 3:00 AM and effectively starts the work at 2:58 AM.

IS-IS events around unexpected incidents. We find that 63.4% of the VPN provider's unplanned trouble tickets and 65% of Internet2's unplanned tickets had overlapping IS-IS events. IS-IS events within unplanned ticket may reflect the effect of the unexpected incident on the network or may also reflect the troubleshooting procedure. In the first case, the timestamp of trouble tickets can come from syslog or some other automatic alarm system.

Figures 3 plot the CDF of the minimum, the maximum and the average of differences $I.start - T.open$ when I starts before T and I starts during T . The different distributions for the VPN provider and Internet2 in Figure 3(a) and 3(b) reflect the different operational practices of these networks (and indicate that we should configure thresholds differently for each network). The IS-IS time of unplanned events is often close to the ticket time (when the ticket timestamp comes from some automatic logging), but tickets corresponding to unplanned events are more often preceded by an IS-IS event.

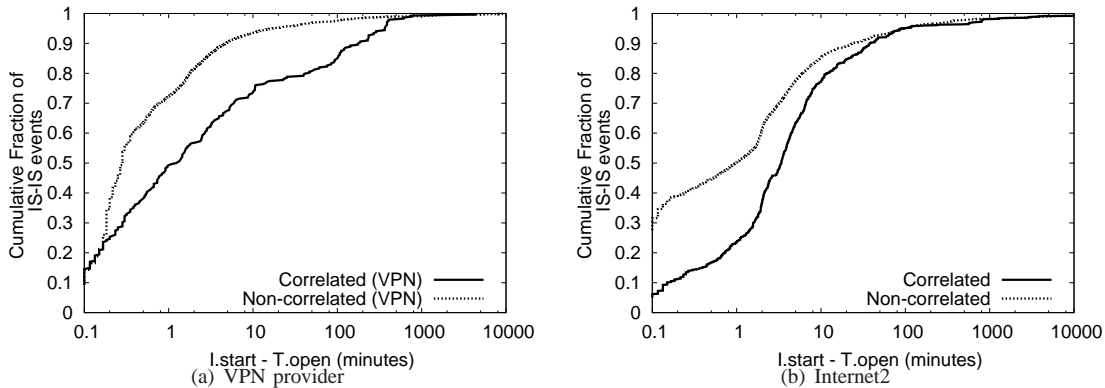


Fig. 4. Duration of correlated and non-correlated IS-IS events.

This case might happen when a ticket is opened to perform an emergency repair or to troubleshoot some unknown event.

In the VPN provider, we say that I is correlated with an unplanned trouble ticket T if $I.start - T.open \leq 10$ minutes. In Internet2, the threshold for correlation is 60 minutes. We pick 10 minutes as threshold for the VPN provider (compared to 60 minutes for Internet2) because in average, 65% of IS-IS events in one day start before the ticket is opened (compared to 13% for Internet2).

When the IS-IS event ends *after* the trouble ticket (i.e., $I.end \geq T.close$), sometimes the IS-IS event continues well beyond the end of the trouble ticket; in these cases, it is unlikely that the IS-IS event is related to the ticket. However, some IS-IS events may legitimately finish after the ticket is closed, because the affected link or prefix takes time to become fully stable. We analyze the distribution of $I.end - T.close$ (more details in our technical report [14]). We conclude that if an IS-IS event continues for more than 5 minutes after the end of the trouble ticket the event is not correlated with the ticket.

E. Comparing correlated and non-correlated IS-IS events

Correlated IS-IS events last longer than non-correlated IS-IS events and most of non-correlated IS-IS events are short. Figures 4 plot the cumulative distribution of the duration of IS-IS events that correlate and do not correlate with trouble tickets. In the VPN provider, almost 60% of correlated events last less than three minutes, while the percentage is 90% for non-correlated IS-IS events. We find similar trend for Internet2 (23% against 51% last less than one minute).

These findings are not surprising because routing is designed to automatically mask many types of failures. A few of non-correlated IS-IS events are long: 3% of events in the VPN provider network and 7% of events in the Internet2 network are longer than one hour. These longer events may correspond to incidents that did not affect the communications in the network or incidents that were not reported by customers or operators. Our analysis focuses on the IS-IS events that correlate with trouble tickets.

Correlated IS-IS events affect more links or prefixes than non-correlated IS-IS events. In the VPN provider, we find that

almost 50% of correlated events affect one link, against around 75% for non-correlated events. These results contrast with those of Internet2, where most of correlated and non-correlated IS-IS events affect only a single link. In internet2, about 75% of non-correlated IS-IS events affect a single prefix, whereas the percentage is only around 10% for the correlated one. This result highly depends on the nature of network incidents. The VPN provider's trouble tickets mostly concern router failures, whereas most of Internet2's trouble tickets relate to circuit problems. We find that, in average, circuit problems only disrupt one link at a time, but affect many more prefixes.

V. MAJOR CAUSES OF ROUTING CHANGES

We correlate trouble tickets and IS-IS events from two operational networks: the VPN provider and the Internet2 backbone. We apply the results of our joint analysis from Section IV to study how each of these events contributes to overall periods of instability.

Table IV summarizes some causes of IS-IS routing changes in both the VPN provider network and in Internet2. The table also indicates how frequently each of these events occurred within the time period of the analysis (column 2), and their contribution to the overall routing instability, as determined by correlating each ticket with the corresponding IS-IS event (column 3). The effect on IS-IS in terms of the duration (column 4), and the average number of IS-IS updates generated per event (column 5).

Hardware and circuit problems cause more instability than software. Most unplanned events in the VPN provider result from router problems, and specifically hardware card crashes. Although software bugs are a significant fraction of problems overall, they contribute less to periods of downtime than the hardware problems. For example, although hardware problems were responsible for only 35% of all trouble tickets, they were responsible for about 41.39% of all periods of routing instability. Overall, router problems were the predominant contributor to both overall number of trouble tickets and periods of instability, whereas in the Internet2 network, the major contributors to instability were related to circuits. *Both* networks experienced problems caused by crashes of router line cards: In the VPN provider, this accounted for 34% of all

TABLE IV
THE UNDERLYING CAUSES OF IS-IS EVENTS IN TWO NETWORKS.

Internet2				
	Trouble tickets		Effect on IS-IS	
	Percent.	Down-time (%)	Duration Avg. (min)	Updates Avg.
Planned	196 (71.5%)	68.6	374.2	20
<i>Router</i>	78 (28.9%)	2.20	28.9	13.6
<i>Software</i>	74 (27.4%)	2.14	30.1	13.2
<i>Upgrade</i>	72 (26.7%)	2.08	30.7	13.2
<i>Configuration</i>	2 (0.7%)	0.06	7.2	14.0
<i>Hardware</i>	3 (1.1%)	0.04	9.8	22.6
<i>Card replaced</i>	2 (0.7%)	0.02	8.1	28.0
<i>Equipment replaced</i>	1 (0.3%)	0.02	13.3	12.0
<i>Other Router</i>	1 (0.3%)	0.02	6.7	20.0
<i>Circuit / Fiber</i>	115 (42.6%)	66.40	628.2	24.5
<i>Hardware replaced</i>	23 (8.5%)	1.13	53.3	23.1
<i>Relocation / rollback</i>	47 (17.4%)	49.67	1107.1	25.7
<i>Power module</i>	2 (0.7%)	0.10	59.0	8.0
<i>Other Circuit</i>	43 (16%)	15.50	176.5	24.0
Unplanned	89 (28.5%)	31.4	420.4	137.0
<i>Router</i>	19 (7.0%)	0.10	4.3	11.8
<i>Software</i>	8 (3.0%)	0.05	3.6	8.8
<i>Software bug</i>	4 (1.5%)	0.01	4.0	10.0
<i>Routing problems</i>	4 (1.5%)	0.04	2.7	9.0
<i>Hardware</i>	7 (2.6%)	0.02	4.1	16.2
<i>Card crash</i>	3 (1.1%)	0.01	1.8	22.0
<i>Other Hardware</i>	4 (1.5%)	0.01	6.0	12.0
<i>Other Router</i>	4 (1.5%)	0.03	6.9	8.0
<i>Circuit / Fiber</i>	56 (20.8%)	31.16	564.7	174.8
<i>Card failures</i>	14 (5.2%)	3.40	262.4	428.0
<i>Other Circuit</i>	42 (15.6%)	27.76	658.7	96.3
<i>Power issue</i>	2 (0.7%)	0.15	4.3	10.0

The VPN provider				
	Trouble tickets		Effect on IS-IS	
	Percent.	Down-time (%)	Duration Avg. (min)	Updates Avg.
Planned	25%	0.37	2.81	3.8
<i>Router</i>	25%	0.37	2.81	3.8
<i>Software</i>	6%	0.027	0.85	3.8
<i>Upgrade</i>	6%	0.027	0.85	3.8
<i>Hardware</i>	18%	0.32	3.5	3.8
<i>Card Replaced</i>	5%	0.05	1.9	4.8
<i>Router decommissioned</i>	13%	0.27	4.0	3.5
<i>Other Router</i>	1%	0.01	3.46	4.0
<i>Circuit / Fiber</i>	0%	—	—	—
Unplanned	75%	99.63	252.3	49.9
<i>Router</i>	69%	91.21	251.37	53.0
<i>Software</i>	26%	24.25	174.84	17.4
<i>Software bug</i>	21%	6.11	54.5	16.3
<i>Routing problems</i>	5%	18.14	680	22
<i>Hardware</i>	35%	41.39	221.6	84.9
<i>Card crash</i>	34%	41.36	228.01	87.2
<i>Other Hardware (Adaptor)</i>	1%	0.03	6	6
<i>Other Router</i>	8%	25.55	684.1	25.8
<i>Circuit / Fiber</i>	3%	7.58	473.5	24.6
<i>Power issue</i>	3%	0.83	52.0	4.0

tickets and about 41.36% of all instability periods; in Internet2, this accounted for 5.2% of all trouble tickets and 3.4% of all periods of instability.

Maintenance and circuit problems remain significant; software problems have become more prevalent. We compared our findings to Labovitz *et al.*, who performed a similar study of a regional provider network in 1999 [10]. That study found that maintenance was the most significant contributor to outages; we also find that maintenance is a predominant contributor to the number of events and overall periods of instability. However, the previous study also found that power outages were responsible for 16% of outages—in contrast, we see that in both networks we studied, power issues were responsible for no more than 3% of all incidents, and no more than 1% of all periods of instability. Software issues have become more predominant: ten years ago, Labovitz *et al.* found that software was the least significant contributor to outages, responsible for only 1.3% of outages; in contrast, software issues are responsible for closer to 30% of incidents in both of the networks we studied.

When they do occur, software problems can be a significant contributor to periods of instability, particularly for the VPN provider. Although more instances that caused IS-IS instability were caused by hardware, both networks also experienced significant problems resulting from either planned or unplanned software problems: In Internet2, for example, software upgrades were responsible for about 26% of all trouble tickets, and about 2% of all periods of instability; in total, software was responsible for about 30% of all events. On average, software events can last considerably longer than hardware events: for example, software upgrades in Internet2

last 30 minutes on average, and unplanned software problems in the VPN provider last about 174 minutes on average. In the VPN provider network, software problems were responsible for about 26% of all trouble tickets, and about 24% of all downtime—unlike in the Internet2 network, most of these problems were unplanned.

Unplanned events last longer than planned events. We define the IS-IS duration of a network incident in tickets as the time difference between the last and the first correlated IS-IS events. We use this duration to evaluate the IS-IS instability triggered by an incident. Figure 5(a) shows the CDF of the IS-IS duration of planned and unplanned trouble tickets for both networks. This figure shows that the IS-IS duration for planned activities is relatively short in the VPN provider. The average duration of tickets in IS-IS is 2.81 minutes for planned and 252.3 minutes for unplanned events. More than half (60%) of the planned tickets lasted less than one minute, whereas only 15% of unplanned tickets lasted less than one minute.

Hardware and circuit problems tend to generate more updates per event. A greater number of IS-IS updates for an incident means more network instability but also a footprint that can be easily detected by network operations centers' troubleshooting tools. Figure 5(b) shows the cumulative distribution function of the IS-IS updates generated by planned and unplanned trouble tickets in the VPN provider and Internet2. For both networks, our algorithm correlates from one to about three thousands IS-IS updates with any incident. As expected, maintenance activities trigger a lower number of updates compared to unexpected incidents, on average. We do not provide the plots for the number of links and the number of prefixes affected in IS-IS events correlated trouble tickets due

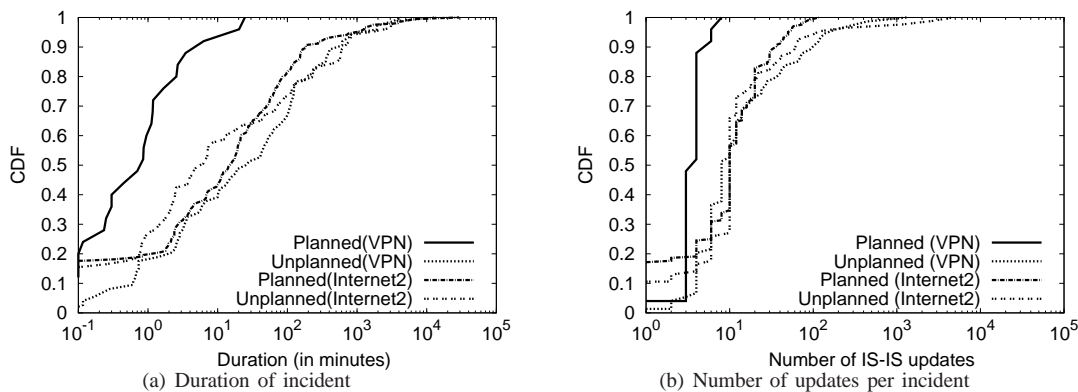


Fig. 5. Duration and number of IS-IS updates for planned and unplanned trouble tickets.

to a lack of space. Router incidents also tend to affect more than two links both in the VPN provider and in Internet2: in the VPN provider, we observed that these events can affect about ten links of a single router, more than any other cause of IS-IS events. On average, software incidents affect twice as many links as hardware incidents, likely because software incidents usually affect the entire router, while hardware incidents are more isolated.

VI. RELATED WORK

Researchers have mainly studied network events using routing messages [5]–[8], [10]–[12] or traceroutes [4], [15]. These studies improved our understanding of network events and their impact on routing and end-to-end paths, but they cannot systematically explain the causes of network events, which is our goal. Our work searches for the causes of intradomain routing changes inside trouble tickets logs. Some studies already use trouble tickets logs [5], [7], [9], [10], but with different goal. Most similar to ours is the early work by Labovitz et al. [10] which studied trouble tickets logs from one ISP to explain routing instabilities. Our study complements theirs. First, we use trouble tickets logs from two different networks, which is important to show the variability of events according to the network. Second, we find the causes of intradomain routing instabilities by correlating these trouble tickets with intradomain routing updates using our correlation algorithm.

VII. CONCLUSION

This paper studied the underlying causes of intradomain routing instability by correlating IS-IS intradomain routing changes with trouble tickets. Our work has revealed some interesting findings, particularly in comparison to the Labovitz et al. study of a regional provider ten years ago [10]. Similar to this study, we found that planned maintenance is a significant contributor to instability overall; however, power outages have become much less significant contributors to downtime, while software problems have become considerably more prevalent (although they still contribute to instability and downtime considerably less than hardware problems). By studying different types of networks, we also discovered that different types of

networks may have drastically different causes for downtime or network instability.

REFERENCES

- [1] Abilene, <https://listserv.indiana.edu/archives/internet2-ops-l.html>, Accessed: September 2007.
- [2] C. Boutremans, G. Iannaccone, and C. Diot. Impact of Link Failures on VoIP Performance. In *Proc. of NOSSDAV workshop*, May 2002.
- [3] M. Caesar and J. Rexford. Building bug-tolerant routers with virtualization. In *ACM SIGCOMM Workshop on Programmable Routers for the Extensible Services of Tomorrow (PRESTO)*, 2008.
- [4] M. Dahlin, B. Chandra, L. Gao, and A. Nayate. End-to-End WAN Service Availability. *IEEE/ACM Trans. Networking*, 2003.
- [5] N. Feamster and H. Balakrishnan. Detecting BGP Configuration Faults with Static Analysis. In *Proc. USENIX Symposium on NSDI*, May 2005.
- [6] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs. Locating Internet Routing Instabilities. In *Proc. ACM SIGCOMM*, Sep 2004.
- [7] Y. Huang, N. Feamster, A. Lakhina, and J. Xu. Diagnosing network disruptions with network-wide analysis. In *Proc. ACM SIGMETRICS*, Jun 2007.
- [8] G. Iannaccone, C. nee Chuah, R. Mortier, S. Bhattacharyya, and C. Diot. Analysis of Link Failures in an IP Backbone. In *Proc. Internet Measurement Workshop*, Nov 2002.
- [9] S. Kandula, R. Mahajan, P. Verkaik, S. Agarwal, J. Padhye, and P. Bahl. Detailed diagnosis in enterprise networks. *Proc. ACM SIGCOMM*, 2009.
- [10] G. Labovitz, A. Ahuja, and F. Jahanian. Experimental study of the Internet Stability and Backbone failures. In *Proc. International Symposium on FTCS*, Jun 1999.
- [11] R. Mahajan, D. Wetherall, and T. Anderson. Understanding BGP Misconfiguration. In *Proc. ACM SIGCOMM*, Oct 2002.
- [12] A. Markopoulou, G. Iannaccone, S. Bhattacharyya, C. Chuah, and C. Diot. Characterization of Failures in an IP Backbone. In *Proc. IEEE INFOCOM*, 2004.
- [13] A. Medem, M.-I. Akodjenou, and R. Teixeira. Troubleminder: Mining network trouble tickets. In *Proc. 1st IFIP/IEEE international workshop on Management of the Future Internet (Manfi 2009)*, 2009.
- [14] A. Medem, R. Teixeira, N. Feamster, and M. Meulle. Analyzing the causes of intra-domain routing changes. Technical report, UPMC Paris Universit as and CNRS, 2009.
- [15] V. Paxson. End-to-End Internet Packet Dynamics. *IEEE/ACM Trans. Networking*, 1999.
- [16] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang. BGP routing stability of popular destinations. In *Proc. Internet Measurement Workshop*, Nov 2002.
- [17] J.-P. Vasseur, M. Pickavet, and P. Demeester. *Network Recovery: Protection and Restoration of Optical, SONET-SDH, IP, and MPLS*. Morgan Kaufmann Publishers Inc., 2004.
- [18] F. Wang, Z. M. Mao, J. Wang, L. Gao, and R. Bush. A Measurement Study on the Impact of Routing Events on End-to-End Internet Path Performance. *Proc. ACM SIGCOMM*, 2006.