Troubleshooting on Intra-Domain Routing Instability

Zhang Shu
National Institute of Information and Communications Technology
4-2-1 Nukui-kitamachi, Koganei, Tokyo
Japan 184-8795
zhang@koganei.wide.ad.jp

Youki Kadobayashi
Nara Institute of Science and Technology
8916-5 Takayama, Ikoma, Nara
Japan 630-0101
youki-k@is.aist-nara.ac.jp

ABSTRACT

Routing instability is a problem directly affecting the reliability of the Internet. While a great deal of effort has been committed to inter-domain routing instability, studies on intra-domain routing have been quite limited. Most network operators still do not have sufficient knowledge on this problem and often complain that: (i) They do not know to what extent the intra-domain routing instability can occur on their networks because this is difficult to detect, and (ii) the causes of this instability are difficult to find. In this paper, we first present the results of some passive measurements we did on intra-domain routing instability. We show the statistical results of OSPF routing information (for both IPv4 and IPv6) we collected on the WIDE Internet and APAN Tokyo-XP network. Through the statistics, we demonstrate how seriously routing instability can occur on a service network. We then propose an approach to help network operators isolate the causes of this. We emphasize the importance of gathering useful data for troubleshooting in event-driven fashion and propose using SNMP or telnet for this. We then explain what kind of data should be collected for the purposes of troubleshooting and how to use this data to isolate the problem.

Categories and Subject Descriptors

C.2.2 [**Computer Communication Networks**]: Network Protocols—*Routing protocols*; C.2.3 [**Computer Communication Networks**]: Network Operations—*Network monitoring*

General Terms

Performance, Management, Measurement

Keywords

intra-domain, routing, instability, monitoring, OSPF, cause

SIGCOMM'04 Workshops, Aug. 30-Sept. 3, 2004, Portland, Oregon, USA.

1. INTRODUCTION

The Internet has greatly extended its influence during the last decade. Nowadays, people use it to routinely obtain information from the WWW, send e-mails, exchange files, and so on. It has become an indispensable tool in their daily lives. However, the current Internet has a lot of room for improvement. Comparing the Internet with the traditional Public Switched Telephone Network (PSTN), although it excels at its ability to provide a variety of services, it lacks reliability. End users often find that sometimes they have extremely poor performance while using the Internet. There are many factors leading to this poor performance, including the link bandwidth, the efficiency of the application software, and the robustness of the protocol. In this paper, we focus on routing instability because it directly affects the efficiency of IP packet forwarding.

Routing instability is also called a route flap and generally it means changes in routes that network operators did not expect (i.e., changes caused by some unknown reason other than network maintenance). Normally, routers use their routing tables (or forwarding tables) to decide the next hop to which they should forward a packet. When the route for a network is mistakenly determined based on information that does not reflect the correct topology, packets for that network will be sent in the wrong direction or simply dropped. What is worse, when the next hop mistakenly determined during route calculation happens to be the router from which the packet was received, and which has not yet recalculated the routing table, a routing loop occurs. The packet is repeatedly transferred between these two routers until its TTL becomes zero or the route for this packet converges on these two routers. When the number of packets involved in the same routing loop reaches a certain level, link bandwidth is significantly wasted and the load on routers is increased, thus making it more likely that other oscillations in routing information will occur.

While much research has been done on inter-domain routing instability, the study on intra-domain routing has been quite limited. Most network operators still have an insufficient understanding of this problem and often complain that:

1. They do not know to what extent intra-domain routing instability can occur on their networks because this is difficult to detect.
2. The causes of this instability are difficult to detect.

In this paper, we first present the results of case studies on intra-domain routing to show how frequently routing

instability can occur on a network that is used daily. We then propose an approach to help network operators isolate the causes of intra-domain routing instability by collecting data that will be useful for troubleshooting in an event-driven fashion. Although we focus our discussion mainly on OSPFv2, the same approach also applies for other link-state routing protocols such as OSPFv3 and IS-IS.

Thus far, there has been some research on the issue of intra-domain routing instability. One study was conducted in Michigan, U.S.A. [5] and they presented the results of an OSPFv2 measurement on a regional network. Analysis of how routing messages could be lost on a congested network has been reported [3].

We have organized this paper as follows: Section 2 presents the statistical results of routing data that we collected from the WIDE [1] Internet and APAN Tokyo-XP network. In Section 3, we describe the method we propose to gather data useful for troubleshooting based on event-driven fashion and explain how to use this data to determine what causes instability. Finally, we conclude our work and introduce future directions in Section 4.

2. MEASUREMENTS ON INTRA-DOMAIN ROUTING INSTABILITY

In this section, we present the methodology and the results of two routing instability measurements conducted on the WIDE Internet and the APAN Tokyo-XP network.

2.1 Methodology

Basically our approach can be divided into two steps: data collection and data analysis.

We chose the WIDE Internet, a national academic network in Japan, as our main target for measurement. The WIDE Internet consists of more than 100 routers (in the whole routing domain) and connects hundreds of organizations. Because it uses OSPF as its main routing protocol, we used OSPF routing messages for our analysis. We collected raw data from a network being utilized by many organizations but not from an experimental network because we thought results obtained under a real workload would be more convincing.

The method we used to collect data was quite simple. We connected a FreeBSD PC to an Ethernet segment which was configured as a part of the OSPF backbone area. As OSPF uses multicast to propagate most of its routing information, we could record all the OSPF messages flooded on this segment with tcpdump, a tool widely used to collect messages over a shared link.

Currently there are five kinds of messages defined in the OSPFv2 specification [2]. Of all these messages, the Link State Update (LSU) packet is used to carry the routing information. In particular, four kinds of Link State Advertisements (LSAs) used in the calculation of intra-domain routes are defined in the LSU packet. They are Router-LSAs, Network-LSAs, Network-Summary-LSAs and ASBR-Summary-LSAs. By analyzing the collected LSAs, we can calculate to what extent each LSA changes in its content. Although we analyzed all these four kinds of LSAs, our discussion are focused on Router-LSAs because they include the most fundamental routing information.

We also collected and analyzed routing messages of OSPFv3 [7] of the WIDE Internet.

The data collection began in August 2000 and is still being conducted. The results we present here are based on data from August 2000 to January 2004. The collected data of this period amounts approximately to 32.8 GB at an average of 27.0 MB per day.

We analyzed all the data with *ospfanaly*, a tool we wrote in C language. *Ospfanaly* uses *libpcap* to read data recorded by *tcpdump* and outputs the changes in each LSA. We also wrote Perl scripts for use in data processing.

2.2 Measurements of OSPF on WIDE Internet

Here we present the statistical results of the OSPF LSAs and summarize some of their oscillation patterns.

2.2.1 Statistical Results of Router-LSAs

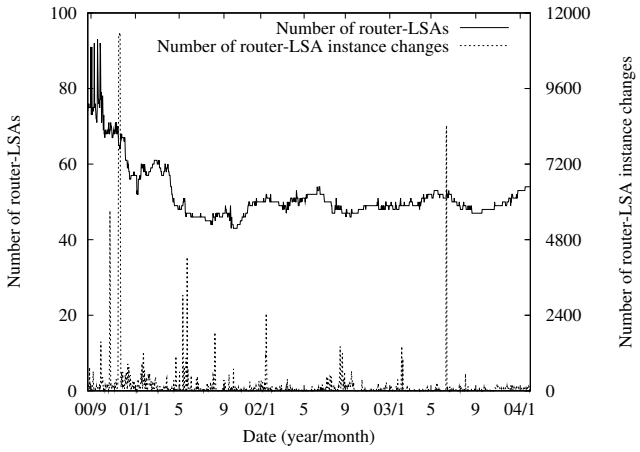

Figure 1: Number of Router-LSAs and their changes in backbone area

The number of Router-LSAs that appeared in the backbone area each day ranges from 43 to 94. In general, because there should not be many changes in the network topology, we had thought that these routers would not originate many changing Router-LSAs. However, we found just the opposite in our investigations. Figure 1 plots the total number of Router-LSAs and changes in them each day. We can see that, although the total number of changes is not that great on most days, there were periods when many changes occurred and sometimes the changes reach 11,000 times in a single day.

Figure 2 is a enlarged version of Fig. 1, where we set 1,000 as the upper bound. From this we can see that although the number of changes tends to be less than 1,000 times per day, there are still significant parts where more than 200 changes occurred.

If these changes had been originated by most of the backbone area routers, we could consider them as normal topology changes due to network maintenance. However, after our analysis, we found that this was not the case. The changes tend to be originated by only a few routers. For example, on November 20, 2000, a total of 11,380 changes occurred, but 98.6% of these changes were due to two Router-LSAs. On April 22, 2001, 1093 of 1097 changes were caused by only one Router-LSA.

From Fig. 2 we can also see that the total number of

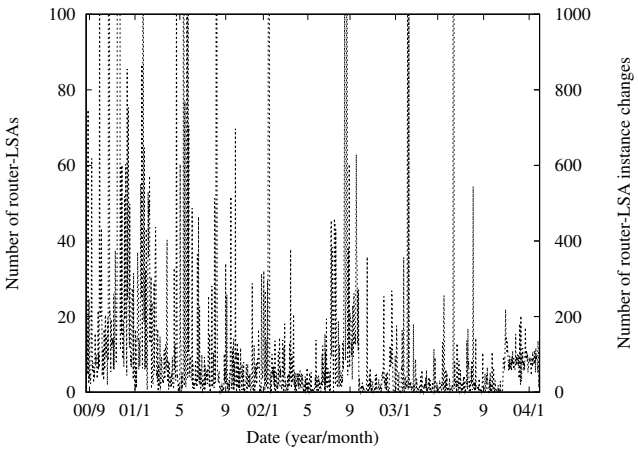

Figure 2: Number of changes in Router-LSAs in backbone area (limited version)

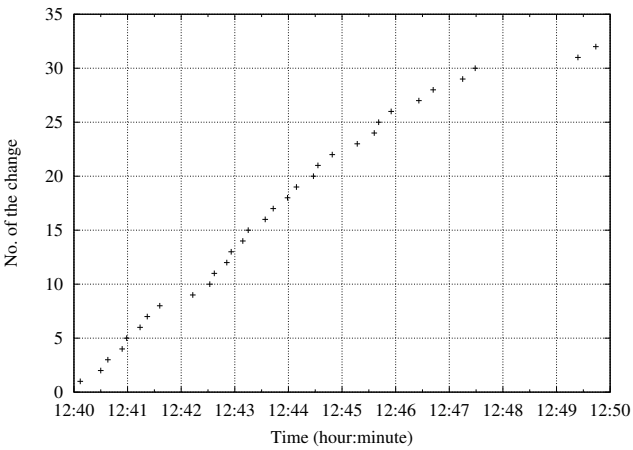

Figure 3: Typical changes in broadcast and NBMA network interface

changes in the Router-LSA is decreasing. This is mainly because in the implementation of modern routers, routing packets are often given higher priority for processing than other packets to avoid dropping routing packets. Another reason for this is lower network congestion due to the increased network bandwidth in recent years.

2.2.2 *Oscillation Patterns*

We classify the observed LSA changes into following categories based on their characteristics.

1. Changes in broadcast and NBMA network interfaces

   When a router on a broadcast or NBMA network finds that there are other OSPF routers on the same link, it describes this network as a transit network (Type 2) in its Router-LSA. Otherwise, this network is treated as a stub network (Type 3).

   Typical changes to a broadcast interface during a 10-minute period are plotted in Fig. 3. During this period, the Router-LSA changes a total of 32 times, or 3-4 times per minute.

   Figure 4 has part of the output generated by *ospfanaly* for the 10-minute period. '+' indicates the addition of a link compared with the last instance of an LSA and '-' indicates the opposite. We can see that most of the oscillations consist of repeated declarations of up/down events for a transit network.

2. Changes in point-to-point network interfaces

   Point-to-point connections are often used on the current Internet to connect distant places. When a router detects a link-down or does not receive its peer's hello packet within a certain time (RouterDeadInterval) over a point-to-point interface, it originates a new LSA, in which the point-to-point link is removed, and floods this new LSA to tell other routers (in the same area) that the link is no longer available. This kind of change is different from that in broadcast or NBMA networks in that the router on the other side will also detect the

```
12:42:32 +link=203.178.137.64 type=3
         -link=203.178.137.77 type=2
12:42:37 +link=203.178.137.77 type=2
         -link=203.178.137.64 type=3
12:42:51 +link=203.178.137.64 type=3
         -link=203.178.137.77 type=2
12:42:56 +link=203.178.137.77 type=2
         -link=203.178.137.64 type=3
12:43:09 +link=203.178.137.64 type=3
         -link=203.178.137.77 type=2
```

Figure 4: Changes in broadcast and NBMA network interface

   failure and originate a new LSA. Thus, when a point-to-point connection fails, we see two LSA changes originated by the two endpoints at about the same time. We also frequently observed this kind of change in our investigations.

   Figure 5 plots typical changes in point-to-point interfaces during a 30-minute period. Figure 6 is part of the output of *ospfanaly* for that period.

3. Changes in metric

   We occasionally found rapid metric changes in Router-LSAs. For example, on one day in April 2001, a layer-3 switch located in Otemachi (in central Tokyo) repeatedly did the following:

   (a) Declared interface-down for one of its interfaces

   (b) Declared the recovery of the down interface, but with a metric of 0

   (c) Declared its metric change from 0 to 1

   This cycle occurred about 560 times throughout a 37-hour period. The reasons for this phenomenon are still unknown.

4. Miscellaneous changes

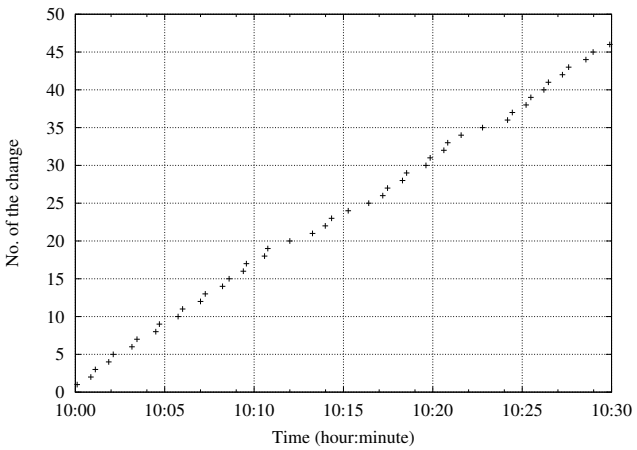

Figure 5: Typical changes in point-to-point network interface

```
10:01:52 -link=203.178.136.22 type=1
10:02:07 +link=203.178.136.22 type=1
10:03:10 -link=203.178.136.22 type=1
10:03:27 +link=203.178.136.22 type=1
10:04:30 -link=203.178.136.22 type=1
10:04:42 +link=203.178.136.22 type=1
10:05:45 -link=203.178.136.22 type=1
10:06:00 +link=203.178.136.22 type=1
10:07:00 -link=203.178.136.22 type=1
```

Figure 6: Changes in point-to-point network interface

We can occasionally see rapid changes in a specific Router-LSA. On one day in early May 2001, a Router-LSA kept changing every 5-6 seconds, with completely different contents each time. This lasted for hours until we found it was caused by a misconfiguration of two routers' using the same router-ID.

2.2.3 OSPFv3 Measurement

Figure 7 shows the statistical results for OSPFv3 Router-LSAs on the WIDE Internet. As IPv6 routing on the WIDE Internet is still in the experimental phase, OSPFv3 shows higher instability compared with what we have observed in OSPFv2.

2.3 Measurements on APAN Tokyo-XP

The Asia-Pacific Advanced Network (APAN) [6] is an international academic network that connects many different countries (or regions) around the Asia-Pacific area. It consists of several exchange points in different countries and usually each of the exchange points is an independent Autonomous System (AS). Among all of the exchange points, the Tokyo-XP is the largest one.

Because the main purpose of the Tokyo-XP network is to provide transit services among academic organizations in different countries, the network itself is relatively small in scale, with no more than ten routers in its OSPF backbone area. However, through our 8-month study, we found we could also observe intra-domain routing instability on such a small-scale network.

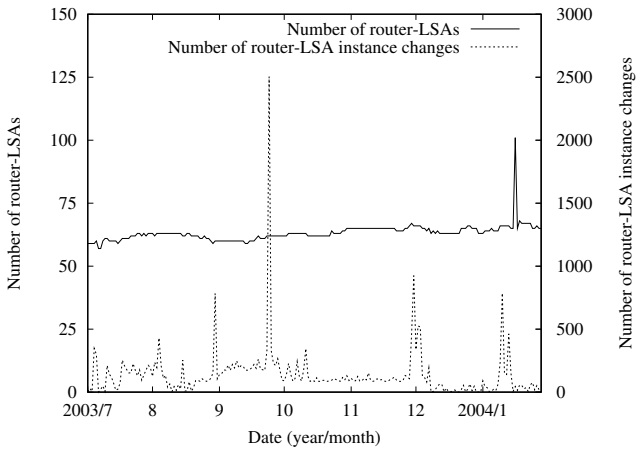

Figure 7: Number of total OSPFv3 Router-LSAs and their changes on WIDE Internet

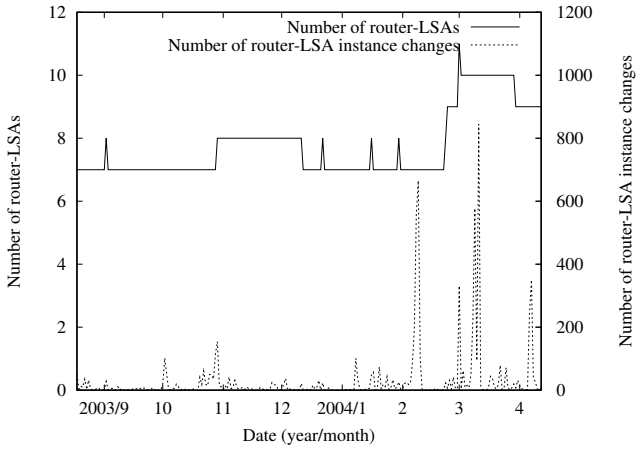

Figure 8: Number of total Router-LSAs and their changes on APAN Tokyo-XP network

Figure 8 has the statistical results for Router-LSAs observed in the APAN Tokyo-XP network. We can see these are like the ones for the WIDE Internet, i.e., although most of the time the routing is relatively stable, the number of changes in Router-LSAs can be quite high at times.

2.4 Causes of Observed Oscillations

During the measurement, we identified the following causes that accounted for part of the observed oscillations.

- Network congestion

  For many times we found that when unusual oscillations were detected, reports of DDoS attacks arrived at almost the same time. Therefore, network congestion due to DDoS attacks may have been part of the causes of the observed oscillations because it has been proved that the network congestion can cause OSPF adjacency to be lost [3].

- Layer-2 failure

  Layer-2 failure is the second serious factor that accounted for the instability. This kind of failure includes trouble with switches and link failures. Because layer-2 switches are becoming more complicated, instability due to trouble with them is increasing gradually.

- Misconfiguration

  The third serious factor affecting LSA stability is misconfiguration due to network operators. We referred to an example in Section 2.2.2 of misconfiguration when the same router-ID is used on two routers.

- Software bugs

  Bugs in routing software is another factor. During our measurements, we were told that a Cisco router using a specific IOS had mistakenly flushed the Router-LSA of another router whose refresh timer was longer than 35 minutes.

3. A SYSTEM TO HELP TROUBLESHOOTING ON INTRA-DOMAIN ROUTING INSTABILITY

As we showed in last section, although we found many causes for the observed oscillations, many remain where the causes are still unknown. The main reason we could not identify all causes was due to the lack of data for troubleshooting. As routing instability usually occurs intermittently, it is difficult to trace all unexpected oscillations timely. For many times, when we got the report of a routing problem and started the troubleshooting, the problem had gone. In this section, we propose a system for troubleshooting routing instability problems which can automatically collect data that informs us of the network's status when routing problems occurred so that we can exactly know what has been happening on the network. The collected data can be used for later troubleshooting even when the network becomes normal.

3.1 Features of the Proposed System

The proposed system aids network operators in the three steps of troubleshooting: instability detection, data collection for troubleshooting and cause isolation.

1. Instability detection

   Currently most network operators depend on reports from users (sometimes operators themselves) to detect routing instability. In the proposed system, we monitor all of the intra-domain routing information flooded on the whole network so that we can detect all routing changes in real time. When detecting any routing information change, it generates a "route changed" event.

2. Data collection for troubleshooting

   Usually the next thing operators have to do after the detection of instability is to interact with all kinds of network equipments, routers or switches, to get the current status of the network. We automate this step in our system. We automatically start the data collection process based on the "route changed" event generated by the monitoring agent. This makes it possible to quickly collect all valuable data right after the instability is detected.

3. Cause isolation

   In this step, traditionally operators try to identify the problem based on their experience as well as data they collected from the interaction with network equipments. Our system also automates this process to further relieve the operators.

3.2 Instability Detection

Instability detection involves finding routing information changes in an OSPF LSA. It can be achieved by monitoring all OSPF LSU packets flooded in an OSPF routing domain and comparing the contents of different instances of a same LSA. When finding any change in the content, it generates a "route changed" event.

There are two methods to monitor the LSU packets:

1. Make the detection agent act as an OSPF router and monitor all of the received OSPF messages.

2. Monitor all OSPF messages on an shared segment in completely passive fashion by packet capture technology, e.g., libpcap.

The former method increases other routers' load but works when there is not any OSPF traffic on a network segment.

3.3 Data Collection

Here we explain how to collect data that is useful for troubleshooting. As we described in Section 3.1, when receiving a "route changed" event, the proposed system automatically starts to collect data from sources considered to hold useful data for troubleshooting.

3.3.1 Data Sources

The following data sources are considered to hold useful data.

1. Routers advertising changed LSAs

   Usually the router advertising a changed LSA is the one most likely to know why it originated the LSA. First, we try to query this router for useful data.

2. Routers described by changed LSAs

   Routers described by changed LSAs can also hold useful data. For example, when changes in a type-2 (point-to-point) link in a Router-LSA are detected, we should query the router on the other end as well as the advertising router.

3. Layer-2 network equipment directly connecting two sources above

   Sometimes layer-2 network equipment, such as an Ethernet switch, can provide important data on troubleshooting, e.g., the link status and the amount of traffic of router's interface.

3.3.2 Methods

We use two methods to obtain data useful for troubleshooting: SNMP and telnet.

Currently most commercial routers support SNMP, which defines a set of objects that can be used to trace a routers' status. In our system, we use the following objects to obtain troubleshooting data.

- OSPF object

  The OSPF object defined in [8] is used to describe all kinds of variables for OSPF. We can obtain information on OSPF interfaces and the data source's neighbors by querying this object.

- Interface object

  The interface object is used to show the status of a router's network interfaces. Entries of ifAdminStatus and ifOperStatus are defined in this object to indicate the desired state and the current operational state of a interface. By comparing the value of these two entries at different times, we know whether the state of a network interface has changed frequently between the up/down states. Entries such as ifInOctets, ifInUcastPkts, ifOutOctets, ifOutUcastPkts are also defined in the interface object to indicate inbound or outbound amount of traffic and packet number. By comparing values of these entries, we can estimate the forwarding rate for a router, thus determining whether or not network congestion has occurred.

- Object of CPU usage

  Sometimes high CPU utilization of a router can cause the router to drop important routing messages such as OSPF Hello packets. By obtaining the CPU usage of a data source when a routing change is detected, we can know whether or not the change has been caused by the data source's high load. The object that indicates CPU usage depends on the vendors of network equipments. For example, CpmCPU is used on Cisco products.

- RMON MIB

  The Remote Monitoring MIB (RMON MIB) was developed by the IETF in order to monitor remote network equipments. Recently, more and more network equipment vendors began to implement it on their products. By querying the etherStats and etherHistory object defined in this MIB, we can get more detailed information on the traffic and even historical data.

For network equipment that does not support or enable SNMP, we use telnet as an alternative method of obtaining troubleshooting data. Telnet is widely supported on both commercial and UNIX-based routers to achieve remote network operation. By executing commands through it, we can obtain information on the data sources. For example, if the data source is a Cisco product, we can use commands such as "show ip ospf", "show interface" or "show processes cpu", to obtain the router's OSPF information, CPU load and interfaces' status. If the data source is UNIX-based, we can use "uptime" and "netstat -i" instead.

Generally a router is attached to more than one interface because its basic function is to forward packets from one network to another. In such cases, we should try to query the data source with all of its IP addresses until we get the appropriate data. This is because when routing becomes unstable, we cannot ensure IP reachability by using only one of the target's addresses. We need to try all of its addresses to maximize our possibilities of obtaining the right data.

3.4 Cause Isolation

Traditionally, it has been the network operators' responsibility to identify what has caused the routing problem and this has usually been conducted manually. In the system we propose, the isolation process is automated. The system will analyze the data collected for troubleshooting to find what problem occurred. For example, when a change in the network interface status is found, the router is identified as the problem. If the amount of traffic on a network interface is over the limit, it indicates that network congestion has occurred. If an OSPF Router-LSA changes frequently with completely different contents, it indicates the possibility of misconfiguration of the router-ID.

4. CONCLUSION AND FUTURE WORK

In this paper, we first presented the results of some passive measurements we did on OSPF LSA oscillations by analyzing OSPF routing messages collected on the WIDE Internet and the APAN Tokyo-XP network. We showed that although most network operators hardly notice them, frequent and persistent routing changes can occur on a network. We also presented some patterns for OSPF LSA oscillations. We listed some causes of the observed oscillations. These were network congestion, layer-2 problems, operators' misconfiguration and bugs in routing software. However, there were still many oscillations that we could not explain because we lacked an effective monitoring system that could provide the data for troubleshooting.

We then proposed a system of helping network operators identify routing problems by obtaining useful data for troubleshooting in an event-driven fashion. We showed what kind of data may be useful for the purposes of troubleshooting and explained how to obtain this through either SNMP or telnet.

We intend to fully implement the proposed system to confirm its validity.

5. REFERENCES

[1] WIDE (Widely Integrated Distributed Environment) Project, http://www.wide.ad.jp

[2] J. Moy: *OSPF Version 2,* RFC 2328, April 1998

[3] Aman Shaikh, Anujan Varma, Lampros Kalampoukas and Rohit Dube: *Routing Stability in Congested Networks: Experimentation and Analysis* SIGCOMM '00, Stockholm, Sweden

[4] Zebra routing daemon: http://www.zebra.org

[5] David Watson, Farnam Jahanian and Craig Labovitz: *Experiences With Monitoring OSPF on a Regional Service Provider Network,* ICDCS2003, May 2003, Providence, Rhode Island, US

[6] APAN (Asia-Pacific Advanced Netowrk), http://www.apan.net

[7] R. Coltun, D. Ferguson, J. Moy.: *OSPF for IPv6,* RFC 2740, December 1999

[8] F. Baker, R. Coltun.: *OSPF Version 2 Management Information Base,* RFC 1850, November 1995