Rtanaly: A System to Detect and Measure IGP Routing Changes

SHU ZHANG[†] and KATSUSHI KOBAYASHI[†]

Routing changes of the interior gateway protocol (IGP), especially unexpected ones, can significantly affect the connectivity of a network. Although such changes can occur quite frequently in a network, most operators have hardly noticed them because of a lack of effective tools. In this paper, we introduce Rtanaly, a system to (i) detect IGP routing changes in real-time and instantly alert operators of the detected changes, (ii) quantify routing changes over the long term to provide operators with a general view on the routing stability of a network, (iii) estimate the impact of routing changes, and (iv) help operators troubleshoot in response to unexpected changes. Rtanaly has the following features: (i) it supports all three widely deployed IGPs - OSPFv2, OSPFv3, and IS-IS, (ii) it uses a completely passive approach, (iii) it visually displays the measurement results, and (iv) it is accessible through the web. We present the results of measurements that we have performed with Rtanaly as well as some observed pathological behavior to show its effectiveness. We have released the first version of Rtanaly as free software and its distribution is based on a BSD-style license.

1. Introduction

The Internet plays an increasingly important role in our society. People not only frequently obtain information that they use in their daily lives through the Internet, but also use it for critical tasks such as investment, banking, and shopping transactions, which require higher levels of reliability. However, the infrastructure of the Internet is still fragile and reachability can be seriously degraded by small failures. In this paper, we focus on interior gateway protocol (IGP) routing changes which can directly affect the performance of an IP service network.

We use the term “IGP routing changes” to refer to changes that can be observed in the link state update messages of an IGP in an operational network. In a routing domain, all routers calculate intra-domain routes based on link state update messages, so changes in these messages directly affect the connectivity of a network. IGP routing changes can be divided into two categories. The first includes changes due to network maintenance and traditionally these changes are considered unavoidable[☆]. The second type of IGP routing change includes those arising from unknown causes, most of which are network failures. In our work, we mainly deal with such unexpected changes.

The results of measurements we have performed and of other studies^{1)~4)} have shown that although network operators are rarely

aware of them, unexpected IGP routing changes can occur in a service network, and sometimes they can occur quite frequently. These changes are difficult to detect because they tend to occur intermittently. In many cases, by the time an operator gets a report that a routing problem has occurred and starts troubleshooting, the problem has disappeared. Thus, an effective system to help operators monitor such anomalies is crucial to provide high quality connectivity service.

In this paper, we present Rtanaly (RouTe ANALYSIS), an open system we developed to detect and measure IGP routing changes. The system (i) detects IGP routing changes in real-time and alerts operators of the changes, (ii) gives the operator a general view on the routing stability of a network by quantifying routing changes over the long term, (iii) estimates the impact of IGP routing changes on the network, and (iv) helps operators troubleshoot problems. Furthermore, by making this system open software, we expect to make it easier for other parties in the research community to perform more measurements on IGP routing changes in different networks, which will allow potential IGP routing problems to be identified more quickly and easily. While we assume Rtanaly will be used in large or relatively large service networks, it can also be deployed in small-scale networks, such as a campus network, to mon-

[†] National Institute of Information and Communications Technology

[☆] Though mechanisms such as graceful restarting to minimize the influence of such changes are being developed.

itor the reachability of the network. Rtanaly is free software and its distribution is based on a BSD-style license. To the best of our knowledge, it is the only open source software at this point that deals with IGP routing changes.

Commercial and academic efforts have gone into developing an IGP monitoring system. Route Explorer and RouteDynamics are commercial products released by Packet Design and Ipsum Networks. However, little is known about the specifics of these products. In Ref. 5), Baccelli, et al. presented the design and the implementation of an OSPFv2⁶⁾ monitoring system, but did not focus on routing changes. In addition, the implementation was neither tested in an operational network nor made public. In Ref. 7), Shaikh, et al. described the design of another OSPFv2 monitoring system. While the architecture of this work is somewhat similar to that of our system, we believe Rtanaly offers several advantages:

- It has been released as open source software, which makes it possible for more extensive measurements to be performed in other networks.
- The impact of observed IGP routing changes can be estimated.
- Results can be clearly visualized, and are accessible through a web browser.
- It provides more support for troubleshooting.
- OSPFv3⁸⁾ and IS-IS⁹⁾, as well as OSPFv2, are supported.

In addition, although not directly dealing with routing changes, techniques using OSPF MIB or raw routing messages were proposed to automatically draw network topology^{10)~12)}.

The rest of this paper is organized as follows. In Section 2, we describe the requirements that an effective system must meet to satisfy network operator needs. In Section 3, we describe the methodology used in Rtanaly. Section 4 presents the architecture and functions of Rtanaly. Section 5 shows the results of measurements done with Rtanaly. We conclude in Section 6.

2. Requirements

Here we summarize the key requirements that should be met by an effective IGP monitoring system aimed at helping network operation. These are based on our own experience gained over a 4-year period of monitoring and comments from other network operators.

Real-time detection First, the system should be able to perform real-time analysis and alert network operators when it detects any IGP routing changes so that the operators can promptly start a diagnosis.

Impact estimation of routing changes Although statistical studies have been done on IGP routing changes^{1),2),4)}, most operators do not know how IGP routing changes affect their networks. An effective system should be capable of quantifying the impact of observed routing changes on the network.

Help in troubleshooting When unexpected IGP routing changes are detected, operators are interested in identifying the problem. Although it is difficult for a monitoring system to pinpoint where trouble has occurred, the system should provide as much related troubleshooting information as possible to operators.

Passive approach A concern of operators regarding the deployment of a routing monitoring system is that the system *must not* have any undesirable side effects on the monitored network. The system should therefore take a passive approach to minimize the possibility of inflicting any damage.

Multi-point monitoring When the monitoring system is disconnected from the remaining network because of a network failure, some link state update messages can be lost. It is desirable to deploy multiple monitoring points so that a complete view of the entire network can be obtained at any time. However, the issue of how to effectively manage data from different locations then arises. A good system should have a function to gather all data around the network in one place in real-time and simultaneously perform the analysis.

Support of different routing protocols Currently most statistical studies have focused on OSPFv2, which can only be used for IPv4 routing. OSPFv3 should also be supported because more and more Internet Service Providers (ISPs) are beginning to provide IPv6 service along with IPv4. In addition, many ISPs, especially large ones, still use IS-IS as an IGP, so IS-IS support should also be provided.

Visualization of results The ability to visualize results is another important factor. Clear visualization will allow operators to gain a better understanding of when and how frequently routing changes are occurring on their networks.

Offline analysis While the main purpose of

the system is to support real-time monitoring, it should also provide an offline data analysis function so that operators can analyze data recorded in the past.

Rtanaly fulfills all of the above requirements.

3. Methodology

3.1 Link-state Routing Protocol

The link-state routing protocol is preferred by most ISPs because of its flexibility, robustness, and efficiency. In a typical link-state routing protocol, the link state update message is flooded throughout a network to disseminate routing information. All routers located in the same routing domain calculate their own routing table based on this information. When there is any change in the network topology, the advertising router will generate a new update message and re-flood it. By comparing a new update message with the previous one, we can figure out what kind of changes have occurred. Note that we do not count refresh updates because they do not reflect any topology change.

3.2 OSPF

In OSPF, link state advertisements (LSAs) are used to disseminate routing information. Currently, five kinds of LSA are used most often in OSPFv2:

Router-LSA In an OSPF routing domain, each router originates a Router-LSA for the area to which it belongs. The Router-LSA describes the collected states of the router's links attached to the area and is flooded throughout a specific area. By analyzing all Router-LSAs, we can learn the link state changes for all routers in an area.

Network-LSA The Network-LSA is used to describe the states of broadcast networks or non-broadcast multi-access (NBMA) networks. A Network-LSA is originated by the designated router (DR) of the network. As is true for a Router-LSA, a Network-LSA is specific to one area. By analyzing Network-LSAs, we can monitor routers that form or break OSPF adjacency in a network.

Summary-LSA The Summary-LSA is originated by area border routers and is used to describe inter-area destinations. Two kinds of Summary-LSA are defined. The Network-Summary-LSA is used when the destination is an IP network and the ASBR-summary-LSA is used when the destination is an autonomous system (AS) boundary router.

AS-external-LSA The AS-external-LSA is originated by AS boundary routers and is used to describe destinations external to the AS.

OSPFv3 was developed for IPv6 routing. In OSPFv3, the following LSAs are defined in addition to those in OSPFv2[☆]:

Link-LSA The Link-LSA is used to provide the router's link-local address to all other routers attached to the link and inform these routers of a list of IPv6 prefixes to associate with the link.

Intra-Area-Prefix-LSA The Intra-Area-Prefix-LSA is used to associate a list of IPv6 address prefixes with a transit network link by referencing a Network-LSA, or to associate a list of IPv6 address prefixes with a router by referencing a Router-LSA.

3.3 IS-IS

In IS-IS, the link state protocol data unit (LSP) is used to flood the routing information. An LSP consists of various triples of type, length, and value (TLV). By analyzing TLVs that are related to the route calculation, we can figure out what kind of routing changes have occurred in IS-IS. Note that not all TLVs affect the network reachability. The TLVs which need to be examined are the IS neighbors TLV, IP interface address TLV, IP internal reachability TLV, IPv6 reachability address TLV, extended IS reachability TLV, and extended IP reachability TLV.

3.4 Estimation of Impact on Network Reachability

When a routing change occurs, the operators will want to know (i) whether the network reachability will be affected by the change, and (ii) how the reachability will be affected if it is. Rtanaly is designed to answer these questions, and it estimates the impact of a routing change by analyzing the intra-domain path (IDP).

We define the IDP as the path a packet traverses from one router to another towards a destination in a routing domain. When any routing change occurs, the shortest-path tree (SPT) of the network for a specific router may or may not change. By comparing the SPTs before and after a routing change, we can figure out whether the SPT has changed, and if it has, how the IDPs from a specific router to other routers have changed. The IDP changes

[☆] In OSPFv3, different names are used for the Summary-LSAs (Inter-Area-Prefix-LSA and Inter-Area-Router-LSA), but the functions are the same.

fall into the following categories:

- (1) The path changes. Its cost may, or may not change.
- (2) The path does not change, but the cost changes. Usually this derives from a change in the cost of a link.
- (3) The path becomes unreachable from a reachable state (a router disappears from the SPT).
- (4) The path becomes reachable from an unreachable state (a router reappears in the SPT).

4. Rtanaly

4.1 Architecture

Rtanaly consists of three components: RA-slaves, the RA-master, and the RA-webstat. As shown in **Fig. 1**, each RA-slave is an IGP data collector that collects raw routing messages and transfers them to the RA-master. The RA-master analyzes the data received from the RA-slaves and generates an alert when it detects an anomaly. The RA-webstat functions on the same node as the RA-master and is the component that operators use to view measurement results regarding observed routing changes through a web browser.

In addition, Rtanaly utilizes the reliable data transmission protocol (RDTP), which we designed to ensure reliable transfer of routing messages from the RA-slaves to the RA-master. Detailed information is provided in Appendix A.1.

4.2 RA-slaves

The RA-slaves collect IGP routing messages at different points of the network and transfer the collected data to the RA-master. To minimize the possibility of inflicting any damage on the monitored network, we use a passive approach to collect the link state update information. As a result, each RA-slave must be connected to a place where the link state up-

date information can be received, such as an Ethernet segment which is configured as a part of the routing domain, or use port mirroring technology[☆]. Note that with this approach, the complete link state database cannot be built up after Rtanaly is started until the refresh messages of all LSAs or LSPs are sent, which can take tens of minutes. However, because this delay occurs only at the very beginning of the measurement phase, we believe the affect is not significant^{☆☆}. In addition, we do not count the first received instance of an LSA as a change because of the lack of previous information.

An RA-slave uses the libpcap library^{☆☆☆}, which is completely passive and widely used to collect packets over a shared link such as an Ethernet link or a mirrored port, to capture the routing messages. There has been research obtaining the entire link state database by SNMP. However, this approach not only increases a router's load, but could miss important changes when more than one change occurs between two SNMP requests.

When an RA-slave is started, it begins data collection and uses RDTP to transfer the collected messages to the RA-master. Sometimes an RA-slave cannot send messages to the RA-master because of a network failure, such as a routing problem or link failure. In this case, the RA-slave will hold the unsent routing messages and restart transferring them when the connection is restored.

4.2.1 Deployment Policy

Because an RA-slave takes a completely passive approach to collect routing messages, deploying an RA-slave in a network segment where IGP messages are flooded is essential. Theoretically, one RA-slave is enough to get all link state updates within an area. However, in some cases, an RA-slave can be cut off from the remaining part of the monitored network and some update messages can be lost. To ensure a complete view of the network status, it is desirable to deploy more than one RA-slave in different subnets of an area. The RA-slave can

Fig. 1 Rtanaly architecture.

[☆] Although this approach imposes some limitations on the deployment of an RA-slave, it can be easily applied in most networks.

^{☆☆} It is possible to obtain the whole link state database from a neighbor through SNMP when Rtanaly is started, but we have not implement this method at this point. We plan to do so in the next version.

^{☆☆☆} Packet capture library, see <http://www.tcpdump.org>

also be deployed in different areas to monitor the network more completely.

4.3 RA-master

The RA-master is the protocol analysis engine. It receives routing messages from RA-slaves and analyzes the messages to find if any routing changes have occurred. The RA-master uses the RDTP to receive the collected messages from the RA-slave(s). It should be capable of receiving and analyzing data from more than one RA-slave. If data transfer between the RA-master and RA-slaves is disconnected unexpectedly, the RA-master starts a new session and waits for an open request from the RA-slaves.

The RA-master can detect a new LSA or LSP. It can also detect an expired LSA or LSP. When analyzing the messages, the RA-master discards duplicated messages and the refresh updates because they do not reflect any routing changes.

When the RA-master finds changes in the link state updates, it can: (i) instantly notify the operators of detected changes, (ii) notify the operators only when changes occur frequently, or (iii) not notify the operators at all. It is up to the operators to decide in which situation they want to be notified. The notification method can be email, syslog, etc. Currently, we only support email notification.

4.4 RA-webstat

The RA-webstat provides a web interface that lets operators view the measurement results, investigate a specific flapping LSA or LSP, and so on. Specifically, it shows the following information.

Statistical results throughout the measurement period These include statistics regarding all OSPF LSAs or IS-IS LSPs. In this way, operators can obtain a general view of the routing changes occurring in a network.

Statistical results for each day These are the daily statistics for OSPF LSAs or IS-IS LSPs which enable operators to know what is happening on a specific day. Real-time analysis results are also provided. In addition, a ranking of LSAs or LSPs sorted by the change number is shown.

Statistical results for a specific LSA or LSP If operators find any unexpected change of an LSA or LSP, they can investigate it in more detail. Both graphical statistics for the specific LSA or LSP and textual results on how it has changed are shown. With these

results, operators can figure out which link or router caused a problem and start further troubleshooting based on this information.

Impact estimation Whether any IDPs from a specific router to other routers has changed will be shown. If an IDP has changed, how it has changed will be displayed.

4.5 Implementation and Availability

As we mentioned in Section 4.2, we use the libpcap library to capture the routing messages. The libpcap library also provides a function which can be used to do offline data analysis. We use RRDtool¹³⁾ to generate the statistical graphs. The RA-slave and RA-master components are implemented in the C language. The RA-webstat is written in Perl script. Version 0.1 of Rtanaly was released in 2004¹⁴⁾. Although we have implemented most of the functions described in this paper, the support of RDTP is still a work in progress. The behavior of Rtanaly has been confirmed on Linux and BSD variants such as FreeBSD. Rtanaly should also work on other UNIX platforms because we provide a “configure” script generated by the GNU autoconf.

5. Measurements Using Rtanaly

In this section, we show the function of Rtanaly through measurements we performed on the WIDE Internet or other networks.

5.1 General Results

Here we show the general result generated by Rtanaly for the WIDE Internet¹⁵⁾, a national academic network in Japan. The measurement began in August 2000 and has lasted for more than four years. **Figure 2** is a graph generated by the RA-webstat. It shows the number of changes observed in all intra-domain LSAs (type 1-4) each day during the period. This gives us a general view of the routing changes occurring in the WIDE Internet. We can see that IGP routing changes occur routinely on the WIDE Internet. Although the change number has been relatively low for most of the period, there have been days when the number spiked to very high levels, sometimes reaching 11,000 per day. Here we need to point out that some of these changes were due to normal network maintenance, but still the number is much higher than what we had expected.

Generated by RA-webstat, **Fig. 3** is a one-day example of LSA changes observed in the Router-LSA of a router located in Tokyo. **Table 1** shows part of the detailed analysis gen-

Fig. 2 Number of changes observed in all intra-domain LSAs.

Fig. 3 Number of changes observed in a Router-LSA in one day.

Table 1 Detailed analysis of a Router-LSA.

Time	Seq. No.	Link ID	Type	Cost
13:19:07	8000482c	-203.178.142.128	3	1
13:19:12	8000482d	+203.178.142.128	3	1
13:19:17	8000482e	-203.178.142.128	3	1
13:19:23	8000482f	+203.178.142.128	3	1
13:19:29	80004830	-203.178.142.128	3	1
13:19:35	80004831	+203.178.142.128	3	1

erated by the RA-webstat which indicates how this LSA changed. In this table, ‘+’ indicates the addition of a link compared with the last instance of the LSA and ‘-’ indicates the opposite. The type of 3 means the link is a stub one. Through this analysis we know that one of the router’s interfaces repeatedly changed its state between up and down. It helps operators quickly find that of which link or router the state is unstable, thus facilitating troubleshooting.

Figure 4 is a graph generated by RA-webstat to show the daily LSP change number observed during a 9-month period in Abilene¹⁶⁾, which is the backbone network of Internet2 and uses IS-IS as the IGP. Abilene is a production network and consists only of 11 routers, but still the change number was quite high from time to time, with the maximum of 4,299 times per day. Our further investigation has shown that on days with such high change numbers, more than 99% of the changes were caused by link problems.

We believe similar phenomena can also be ob-

Fig. 4 Number of LSP changes observed in Abilene.

Fig. 5 Relatively frequent short-term oscillation.

served in other networks and strongly recommend operators to conduct such measurements with Rtanaly.

5.2 Pathological Changes

Here we present some pathological changes we observed using Rtanaly during the measurements. We categorize these changes by their frequency and persistency. All graphs showed here are generated by RA-webstat.

We also provide the most common causes we have found for each change pattern. Operators can also use them for their own troubleshooting.

5.2.1 Relatively Frequent Short-Term Oscillation

Figure 5 is an example of the most typical routing changes which occurred relatively frequently over a short term. We can see that during one day in May 2004, this Router-LSA changed twice in the early morning, declaring the down and up of two links in 10 seconds. It began the up/down advertisements again from noon and continued oscillating for about 4 hours. In our investigation, of all different oscillation patterns, changes like this were observed most often. Usually they were caused by network congestion, but they could also be due to intermittent interface up/down or link problems.

5.2.2 Frequent Short-Term Oscillation

Figure 6 is an example of a serious oscillation. It shows the LSA change number for an L3 switch within 24 hours in 2004. As we can see, this router-LSA kept advertising the up/down of two links at a rate of about 70 times every

Fig. 6 Frequent short-term oscillation.

Fig. 7 Relatively frequent long-term oscillation.

Fig. 8 Less frequent long-term oscillation.

5 minutes, and in total it changed more than 16,000 times in 18 hours. We observed this kind of oscillation for about five times during our measurements, and all of them were caused by misconfiguration where the same router-ID was used on two different routers. Changes leading to this oscillation pattern can also be caused by p2p interface or link problems, and in this case, we will see two changing LSAs.

5.2.3 Long-Term Oscillation

Figures 7 and 8 show one-day examples of relatively frequent and less frequent long-term oscillation. Figure 7 is the result for a router located in San Francisco. We can see that the router originated changing LSAs relatively frequently, and the oscillation lasted for more than five months. Figure 8 is for a router located in Kyoto. We can see that the changes were relatively infrequent, with only several times a day, but the oscillation persisted for about two months. Oscillations like these were usually caused by intermittent interface up/down and link failures.

Although we observed many other interesting events during our measurements, we do not show all of them here due to space limit. More

Fig. 9 Type 1-4 LSA change number on December 20th, 2004.

Fig. 10 SPT changes.

statistical results from Rtanaly are available in Ref. 14).

5.3 Impact Estimation

In this section, we show how Rtanaly estimates the impact of a routing change on the network reachability.

5.3.1 Impact on an SPF

As an example of impact estimation on the SPT of a router, we analyzed the OSPF data collected from the backbone area of the WIDE Internet on December 20th, 2004. As shown in Fig. 9, a lot of OSPFv2 LSA changes occurred that day. The root we used in the SPT calculation was a router near the location where we collected the data.

First, we analyzed the change of node number in SPTs before and after an LSA change. We found that most (85.7%) of the LSA changes did not cause any change in the node number. Fifteen (10.2%) changes caused fewer than five nodes to become reachable or unreachable, and only six of the total 147 changes caused more than five nodes to become reachable or unreachable.

We also analyzed how the SPT was affected by each routing change. We show the number of cost and nexthop changes caused by each LSA change in Fig. 10. From this graph, we can see that most LSA changes (70.7%) did not cause any cost or nexthop change. However, some of the LSA changes (10.2%) caused cost changes to more than 27 nodes, which is about half of the total node number in a complete SPT.

Fig. 11 Type 1-4 LSA change number on May 5th, 2005.

Fig. 12 Impact on IDPs at large.

5.3.2 Impact on IDPs at large

In Section 5.3.1, we showed how Rtanaly estimated the impact on the routing table of a specific router. Usually when a routing change occurs, a network operator also wants to know the overall impact of this change on the reachability of the entire network. Here we present an example on how Rtanaly estimates the impact on IDPs at large.

We selected 19 out of the total approximate 50 routers in the backbone of the WIDE Internet which we thought important, and analyzed the IDPs among these routers[☆] over the routing changes.

Figure 11 shows changes that occurred May 5th, 2005. **Figure 12** shows the number of IDPs affected by each routing change. We can see about two thirds of the total routing changes did not cause any IDP change. We can also see some routing changes has caused almost half of the total IDPs to change. This is because the routing changes has occurred in a core segment of the WIDE Internet, and as a result, one LSA change tended to cause more IDPs to change.

Through the above results, we now know that the impact of routing changes also depends heavily on the routing change itself, but not only on the change number.

6. Conclusion

Although previous work has shown that unexpected IGP routing changes can occur in a network, and sometimes occur quite frequently, network operators have hardly noticed them due to a lack of an effective monitoring and measuring system. In this paper, we have described Rtanaly, which not only detects and alerts operators of IGP routing changes, but also measures the overall routing changes that occur in a network over a long term and estimates the impact of routing changes. Rtanaly uses a passive approach to monitor the link state update messages of IGPs and supports all three widely deployed IGPs. Our long-term measurements on different networks have proven the effectiveness of Rtanaly.

We have released the first version of Rtanaly and plan to release the next one in 2005. RDTP will be implemented in the forthcoming version.

Acknowledgments We would like to express our thanks to WIDE Project and Internet2 for providing valuable routing data of their networks.

References

- 1) Shaikh, A., Issett, C., Greenberg, A., Roughan, M. and Gottlieb, J.: A case study of ospf behavior in a large enterprise network, *Proc. ACM Internet Measurement Workshop* (2002). <http://citeseer.ist.psu.edu/shaikh02case.html>
- 2) Watson, D., Jahanian, F. and Labovitz, C.: Experiences with monitoring ospf on a regional service provider network, *Proc. 23rd International Conference on Distributed Computing Systems*, pp.204–213 (2003).
- 3) Zhang, S.: Case studies in intra-domain routing instability, *31th NANOG meeting* (May 2004). <http://www.nanog.org/mtg-0405/shu.html>
- 4) Zhang, S. and Kadobayashi, Y.: Troubleshooting on intra-domain routing instability, *Proc. ACM SIGCOMM'04 Workshops (NetTs)*, pp.289–294 (2004).
- 5) Baccelli, E. and Rajan, R.: Monitoring ospf routing, *Proc. IFIP/IEEE International Symposium on Integrated Network Management*, pp.825–838 (May 2001).
- 6) Moy, J.: OSPF version 2, RFC 2328 (Apr. 1998).
- 7) Shaikh, A. and Greenberg, A.: "OSPF monitoring: Architecture, design, and deployment experience," *Proc. Symposium on Networked Systems Design and Implementation* (Mar.

[☆] The number of total IDPs is $19 * 18 = 342$.

2004).

- 8) Coltun, R., Ferguson, D. and Moy, J.: OSPF for IPv6, RFC2740 (Dec. 1999).
- 9) Oran, D.: OSI IS-IS intra-domain routing protocol, RFC1142 (Feb. 1990).
- 10) Mansfield, G., Jayanthi, K., Ika, T., Ohta, K., Nemoto, Y. and Kato, N.: Network cartographer, *MULTIMEDIA '96: Proc. 4th ACM international conference on Multimedia*, New York, NY, ACM Press, pp.439–440 (1996).
- 11) Keeni, G.M., Saitoh, T. and Abe, K.: Management of next generation network (jgn ipv6) (in Japanese).
- 12) Taniguchi, K., Noguchi, M., Ohta, K. and Mansfield, G.: Techniques for generation of maps for internet orienteering.
- 13) RRDtool. <http://people.ee.ethz.ch/~oetiker/webtools/rrdtool/>
- 14) Intra-domain routing stability measurement project. <http://rtanaly.koganei.wide.ad.jp/>
- 15) WIDE Project. <http://www.wide.ad.jp>
- 16) Internet2 abilene backbone network. <http://abilene.internet2.edu/>

Appendix

A.1 RDTP

The RDTP is used by RA-slaves and the RA-master to transfer collected routing messages. We took the following considerations into account in its design.

- It must be able to transfer the data collected at each RA-slave to the RA-master in real-time.
- It must ensure that the data sent to the RA-master is completely received. This requires explicit acknowledgments from the RA-master.

If the data cannot be transferred from an RA-slave to the RA-master because of any network failure, the RA-slave must hold the data (even if its connection with the RA-master times out) and periodically send open requests to the RA-master until the connection is restored. When the connection resumes, the RA-slave resends the holding data to the RA-master before sending any newly collected data. This mechanism guarantees that all collected data will be reliably transferred to the RA-master in spite of any network failure and requires sequencing of the data.

The RDTP uses TCP as its transport protocol. As shown in **Fig. 13**, the message sent

Fig. 13 RDTP message format.

from an RA-slave to the RA-master is composed of one byte for the version number (currently 1), one byte for the data type (1 for the libpcap format data), four bytes for the sequence number, and four bytes for the data length followed by the collected data.

(Received July 11, 2005)

(Accepted February 1, 2006)

(Released April 5, 2006)

Shu Zhang joined National Institute of Information and Communications Technology in April 2003. His research focuses on Internet routing, network management and network measurement. He received his bachelor's degree in electrical engineering from Waseda University in 1996, and his master and Ph.D degrees from Nara Institute of Science and Technology in 1999 and 2003, respectively. He is a member of ACM and the WIDE Project.

Katsushi Kobayashi received his B.E., M.E., and Ph.D. degrees in Electro-Communications Engineering, The University of Electro-Communications, Tokyo, Japan in 1987, 1989, and 1994 respectively. During 1994–1998, he was the research associate of Information Processing Center, The University of Electro-Communications. He joined the Communications Research Laboratory (CRL), Tokyo, Japan in 1998. He is currently the leader of Internet Architecture Group at The National Institute of Information and Communications Technology (NICT), Tokyo, Japan. He is a member of IEEE and ACM.