SSM-PING: A PING UTILITY FOR SOURCE SPECIFIC MULTICAST

Pavan Namburi
Department of Computer Science
University of Texas at Dallas
Richardson, Texas - 75080
Email: pavan@student.utdallas.edu

Kamil Sarac
Department of Computer Science
University of Texas at Dallas
Richardson, Texas - 75080
Email: ksarac@utdallas.edu

Kevin C. Almeroth
Department of Computer Science
University of California, Santa Barbara
Santa Barbara, California - 93106
Email: almeroth@cs.ucsb.edu

ABSTRACT
Management has become a key concern for the success of multicast deployment in the Internet. One of the most important management tasks for multicast is to verify the availability of the service to its users. This task is referred to as discovering or testing *reachability*. In this paper, we present a new mechanism, SSM-Ping, for discovering multicast reachability between (remote) end systems in Source Specific Multicast (SSM) enabled networks. SSM-Ping uses the existing Protocol Independent Multicast (PIM) based multicast join mechanism in the network and requires a simple extension to the Internet Group Management Protocol (IGMP). First, we motivate our work by discussing the need for an efficient and effective ping utility for multicast. Then, we present the SSM-Ping operation and discuss a number of deployment and security issues. Finally, we outline our inter-domain scale experimental evaluation approach for SSM-Ping.

KEY WORDS
Source specific multicast, reachability, multicast ping.

1 Introduction

The original IP multicast service model, now called Any Source Multicast (ASM) [4], was designed to provide support for a large group of multicast applications. Over the last decade, it has been well-recognized that this service model has several problems preventing its wide-scale deployment in the Internet [5]. Later on, in order to facilitate multicast deployment in the network, a simplified multicast service model, called Source Specific Multicast (SSM) [7], was defined. With the introduction of SSM, the expectation now is that multicast deployment and the usage of multicast in the Internet can increase significantly.

Some of the main reasons that have been contributing to the lack of wide scale ASM deployment are the complexity of the required protocol architecture and the lack of necessary management tools/systems. Due to the relative simplicity of the SSM service model, the required protocol architecture is not a deployment bottleneck any more. But now, we need to pay more attention to the management needs of SSM to make its deployment successful.

One of the most important management tasks for multicast is to verify the availability of the service to its users. This task is usually referred to as monitoring or testing *reachability* between the sources and the receivers in a multicast group. Multicast is realized through the creation and maintenance of forwarding trees connecting sources and receivers in a multicast group. These trees are dynamically created and maintained by the routers, yet there is no feedback information built into the process. That is, if a tree cannot be built because there is no path to the source, the receiver will never know. Reachability ensures that sources can reach all existing and potential group members. Reachability also implies that receivers have multicast connectivity and can reach all sources. Consequently, verifying reachability becomes very important to maintaining availability and robustness of the multicast service between sources and receivers. Without it, the multicast infrastructure becomes disconnected and essentially unusable.

In this paper, we present a mechanism to test multicast reachability between two end systems in SSM-enabled networks. We call our approach *SSM-Ping*. *Ping* is one of the most basic yet one of the most useful network diagnostic tools constantly used for network management purposes. In unicast, *ping* provides a convenient way of discovering unicast reachability between two systems in the network. On the other hand, multicast *ping* (*mping*) requests are sent to a multicast group address and these requests trigger group receivers to send *ping* responses to the pinging host via unicast. This essentially informs an end system (pinging host) about the fact that there are a number of receivers that received the request on the group address. This information has only very limited use/value and the mechanism is vulnerable to feedback implosion problem. Therefore, this approach is not put into practical use by network operators [13].

Due to the lack of an appropriate diagnostic tool, multicast reachability has traditionally been monitored/verified by using application layer monitoring mechanisms [10, 12]. As we discuss in Section 2, these mechanisms have their own limitations and cannot really provide an effective solution for reachability management. With *SSM-Ping*, we introduce a convenient mechanism to test multicast reachability between a receiver and a source site and therefore fill in an important void for multicast network management. This utility can be of great value for network operators to debug potential multicast reachability problems within and between SSM enabled networks. It is also useful for multicast receivers to verify the reachability between themselves and remote sources.

The rest of the paper is organized as follows. Next section motivates this work. Section 3 discusses the other alternatives we can use for the purpose of reachability. Section 4 presents the SSM-Ping mechanism. Section 5 discusses a number of deployment issues. Section 6 describes

usage scenarios for SSM-Ping. Section 7 presents a number of security issues. Section 8 discusses our evaluations of SSM-Ping. Finally, Section 9 concludes the paper.

2 Motivation

Discovering multicast reachability between multicast enabled hosts/networks has been a difficult task [12]. One explicit mechanism for determining reachability between two end systems is the ping utility. In unicast, ping allows a source/receiver to test bidirectional reachability to the peer. This relationship is shown in the Figure 1. Ping uses ICMP_ECHO and ECHO_REPLY messages to probe the target host. On a successful reply from the target machine, the source is ensured of both liveness and reachability to the target machine. In multicast, due to a lack of an appropriate mechanism, previous work has focused on using application layer information for reachability monitoring. As an example, the Multicast Reachability Monitor (MRM) [2] protocol has been designed to provide reachability monitoring service between remote nodes in the network. But it requires running test multicast sessions between the nodes. In addition, the recent sdr-monitor project [12], used multicast session announcements to monitor multicast reachability among a large group of multicast users. In this approach session announcements available at user locations are used to represent reachability between the announcement originator sites and the user sites.

Finally, the Multicast Beacon [10] has been developed as an active measurement tool to monitor multicast reachability among a large number of participants by using a test multicast session. Even though the above mentioned approaches work to monitor reachability, they have some key limitations: (1) lack of flexible monitoring, (2) lack of heartbeat message control, (3) lack of consistent monitoring and (4) lack of convenience. These drawbacks limit the usage of such type of application layer approaches for reachability monitoring.

3 Approaches to Reachability Discovery

Contrary to previous application dependent approaches, our goal in this paper is to develop a relatively simple tool to verify multicast reachability between remote hosts. As we discussed in Section 1, the mping tool cannot be effectively used to verify reachability.

An alternative approach to verify multicast reachability between a receiver and an *active* source site may be to use a simple unicast based query-response message pair. That is, when a receiver joins an SSM channel (S,G) of a remote source, S, it will expect to receive packets on the (S,G) channel. If the receiver fails to receive multicast data, it may mean two things: (1) either the source is not active, or (2) there exists a reachability problem between the two end hosts. In this case, the receiver can send a unicast-based query to the source, S, to verify if the source, S, is currently active or not. On receiving this query, the source, S, can send a unicast-based response informing the receiver about its activity. Based on this information, the receiver can easily verify its reachability to the source site. On the other hand, if the source is not currently active, then this information does not really help the receiver to know its

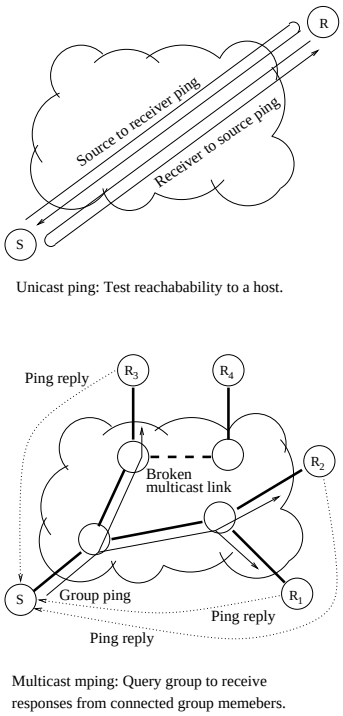

Figure 1. Semantics of the current unicast *ping* (top) and multicast *mping* (bottom).

reachability to the other host. In this case, the source may need to send explicit response information on the (S,G) address to help the receiver to decide about reachability.

One problem with the above approach is that if the (S,G) group has a large number of receivers, the explicit response that the source, S, sends on the (S,G) will go to all receivers unnecessarily. In addition, as the number of receivers trying to verify reachability increases, the amount of traffic introduced into the (S,G) group will be too high. Furthermore, if we want to verify multicast reachability to a remote host which does not currently source any multicast traffic, this approach may not be effective. Because of these disadvantages, the effectiveness of this approach seems to be fairly limited. The next section discusses the approach of SSM-Ping to ensure reachability to the source.

4 SSM-Ping

In this section, we propose a ping utility to verify reachability between SSM-enabled (remote) systems. By using a dedicated multicast group address, say SSM-PING.MCAST.NET, in the SSM address range (232/8), an end system, R, sends an SSM-Ping request to a remote host, S, and waits for a SSM-Ping reply on the (S, SSM-PING.MCAST.NET) channel. The SSM-Ping request is essentially a join request for the (S, SSM-PING.MCAST.NET) SSM channel and is achieved by sending an Internet Group Management Protocol (IGMP) [3] Membership Report by the pinging end system (SSM Receiver, R). The designated multicast router at the pinging site creates a Protocol Independent Multicast (PIM) [6] based join (PIM-Join) message for (S, SSM-PING.MCAST.NET) and forwards it toward the

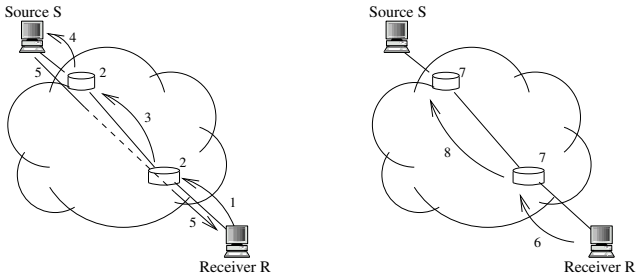

1. IGMP Receiver Report for (S, SSM-PING).
2. Router creates forwarding state for (S, SSM-PING).
3. Router forwards PIM-Join message for (S, SSM-PING).
4. Designated router on S's site sends an IGMP message to S to inform it about the mcping request.
5. S sends mcping reply on (S, SSM-PING).
6. IGMP message to leave the (S, SSM-PING) channel.
7. Router removes (S, SSM-PING) entry from its forwarding table.
8. Router sends a Prune message to upstream neighbor toward S.

Figure 2. Operation of the SSM-Ping utility.

source, S. Each router on the R-to-S reverse shortest path creates a forwarding state entry for the (S, SSM-PING.MCAST.NET) channel and also forwards the join message towards the source S. When the join request reaches the edge router at the source site, this router forwards a message to inform S about the SSM-Ping request. On receiving the SSM-Ping request, the source, S, creates a reply message and sends it to the (S, SSM-PING.MCAST.NET) SSM channel. This message propagates on the previously established multicast forwarding path between S and R and reaches the pinging host, R. On receiving this message, the pinging host creates and sends an IGMP Leave Group message to leave the (S, SSM-PING.MCAST.NET) SSM channel. Consequently, the edge router at the receiver (pinging) host sends a PIM-Prune message to its upstream neighbor on the tree to start flushing the previously created forwarding state for (S, SSM-PING.MCAST.NET) in the network. Figure 2 presents this operation.

The above mechanism requires a new message between the Designated Router (DR) and the source (step 4 shown in 2). According to the current join mechanism, the join messages terminate at the DR and join requests are not conveyed all the way to the source. However in our approach we need to inform source, S, about the incoming ping request (i.e. incoming (S, SSM-PING.MCAST.NET) join message) so that it creates and sends a ping reply. For this reason we introduce a new message to IGMP. Using this new message, the DR at the source site will inform the source about the ping request. This message will only be used when the DR receives a ping request (a join message to the (S, SSM-PING.MCAST.NET) channel) and joins to all other groups will terminate at the DR as usual.

Note that an alternative would be to have the DR act on the incoming ping request. That is, instead of forwarding the ping message to the source, the DR can create a ping response and send it to (S, SSM-PING.MCAST.NET) on behalf of S. This approach would remove the need to modify IGMP. But the disadvantage of this approach is that if there are problems between the DR and S, these problems would fail to be reported by SSM-Ping. In fact, during our implementation efforts we experienced such a problem.

In order to better understand the development and operational issues, we built a test environment between our site at the University of Texas at Dallas (UTD) and the University of Oregon (UO). The network at UO is SSM enabled and the network at UTD is ASM enabled. Figure 3 shows the network layout, with the part of UTD magnified to show additional detail. As seen in the figure, the end host at UTD is connected through a series of switches to the DR which is in turn connected to the UTD border router. As part of our experiments we needed to run a receiver at UO and a source at UTD. The receiver at UO was responsible for pinging the source at UTD and then the source at UTD was responsible for sending a ping reply on the SSM-Ping channel.

During our experiments we observed that the multicast data (ping response) sent by the source in the UTD network was not visible outside the UTD domain. After some investigation we realized that this was due to a misconfiguration on the switches between the source and the DR router at UTD. In fact, our SSM-Ping utility helped us debug and correct an existing multicast reachability problem in our network. This also shows that the ping request should reach all the way to the pinged end system. If the ping request is to terminate at the DR of the pinged source site, potential reachability problems between the DR and pinged host may not be detected. As a result this justifies the need for the new IGMP message that is to be used to convey incoming ping requests to the pinged source.

5 Deployment Issues

An important characteristic of the SSM-Ping utility is that it does not require a significant number of changes to the existing SSM-based multicast service model for its deployment. The only modification required is to include support for a new message type in the IGMP protocol.

The current mechanism allows a receiver to send an SSM-Ping request (i.e. IGMPv3 Membership Report to join an SSM channel). This request is propagated all the way to the edge router at the pinged source site. In order for the source to send a response to an incoming ping request, it needs to be notified about the ping request. As discussed in Section 4, on seeing a join request for SSM-

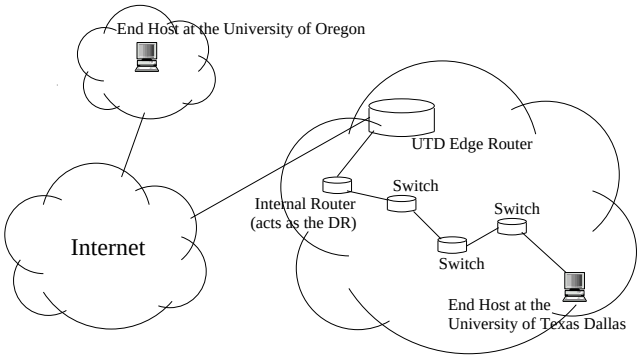

Figure 3. Network layout between UTD and UofO.

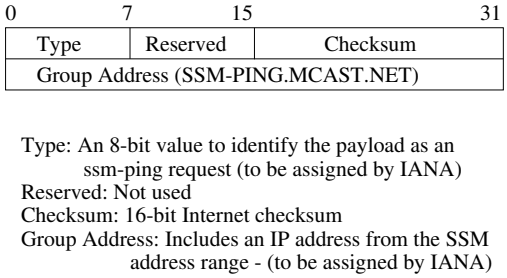

Figure 4. IGMP message format for SSM-Ping notification.

PING.MCAST.NET, the edge router on the source site informs the source about the ping request. For this we propose a new IGMP message type. Figure 4 gives the packet format for this new IGMP message type. The edge routers and the end hosts need to support this new message type in the IGMP protocol.

6 Multicast Management with SSM-Ping

The above mechanism provides sufficient support to test multicast reachability between remote end systems. Network operators can use this service to maintain a robust SSM service within and between SSM-enabled networks in the Internet. In this section we present the utility of SSM-Ping as a tool for multicast management.

6.1 Reachability Testing

The main motive of SSM-Ping is to test if the pinging site (receiver) can receive multicast data from the pinged site (source). In the previous sections we discussed the basic functioning of SSM-Ping tool. That discussion assumes that there is only one pinging receiver at any given point of time. But in reality, there may be more than one receiver initiating a ping requests simultaneously.

A difficulty arises with the SSM-Ping mechanism when there exists simultaneous ping requests from different receiver sites. The problem is that not all ping requests may reach the source site. Recall that SSM-Ping requests are in fact PIM-Join messages propagating in the network towards the source site. Some may join the SSM tree corresponding to the SSM-Ping channel in the middle of network. When such a join message reaches a router which is already on the SSM tree, the router does not continue to forward the request. As a result, the multicast path originating from the receiver site is grafted onto the existing multicast forwarding tree at an intermediate router in the network. In the case of SSM-Ping, this causes a potential problem.

Figure 5 presents an example scenario that visually explains the multi-SSM-Ping problem. In Figure 5-a, R1 sends an SSM-Ping request to a remote source, S. This request propagates in the network toward S. Routers on this path create forwarding state for the channel. Then, in Figure 5-b, source S sends an SSM-Ping response via multicast on the (S,SSM-PING.MCAST.NET) channel and this response passes through the router, X, and reaches R1's site. Finally, in Figure 5-c, R2 sends an SSM-Ping request to remote source, S. At this point, the ping request reaches router, X, and is not forwarded anymore. This is because the ping request is a PIM-Join message and since X is already on the (S,SSM-PING.MCAST.NET) forwarding tree, it does not forward the join message anymore. However, from the ping semantics point of view, this means that the ping request of R2 does not reach the source site and therefore may not receive an answer. Hence, R2 may interpret this as a lack of reachability to S.

In order to resolve this discrepancy, the pinged source periodically sends multiple copies of ping responses on the SSM-Ping channel for a short time interval (e.g. one ping response per second for 10 seconds). This way an SSM-Ping request (PIM-Join request for the SSM-PING channel) that is grafted on the existing multicast forwarding tree for the SSM-PING channel will have a better chance of receiving a response from the source on the existing forwarding tree.

6.2 RTT/Jitter Computation with SSM-Ping

In addition to reachability/connectivity management, one common use of the ping tool in the unicast world is to measure the end-to-end round trip delay (unicast RTT) between two remote end systems. Similar to unicast Ping, SSM-Ping can also be used to measure *multicast RTT* between two remote end systems. Multicast RTT is the delay for the first multicast packet received by the receiver, after it has joined the multicast channel.

Similar to the discussion in Section 6.1, due to the possibility of multiple simultaneous SSM-Ping requests for an end system, the responses returned to the pinging hosts may not always be used to measure the multicast RTT between the two end hosts. However, using simple sequence numbers in the SSM-Ping response message helps pinging hosts to detect for which request message a response message matches. That is, when a receiver pings a source and receives a ping response with sequence number "zero", the receiver is likely the only host pinging the source currently. On the other hand, if the receiver receives a response with sequence number greater than zero, it may indicate the existence of other receivers pinging the source.

The SSM-Ping tool can also be used to estimate jitter. By measuring the inter-arrivals of incoming SSM-Ping responses, the receiver can calculate jitter. This information can be used to estimate the load on the underlying network.

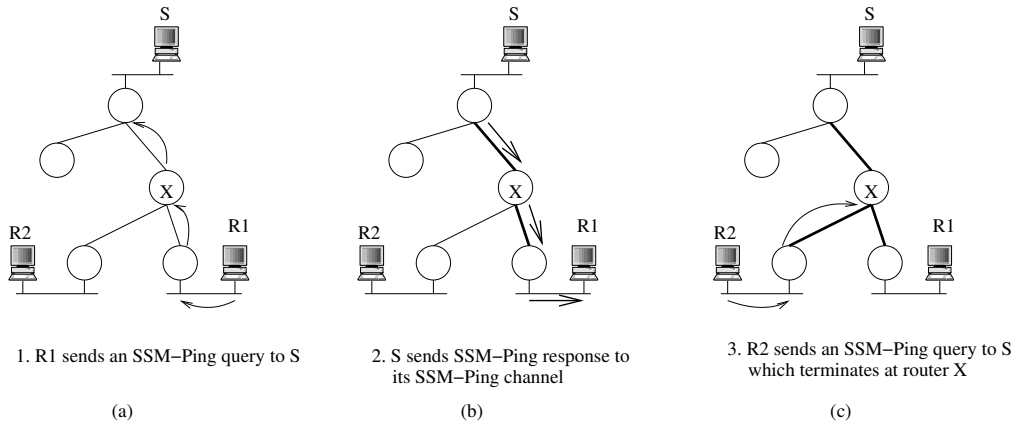

Figure 5. Operational issues.

7 Security Issues

In unicast, ping can be used to initiate denial of service attacks at some third party site. Adversaries can use ping to launch reflector-based denial of service attacks [9] on victim sites. On the other hand, due to the multicast forwarding mechanism, SSM-Ping requests and SSM-Ping responses follow the same path. As a result, SSM-Ping cannot be used for third party denial of service attacks.

A second attempt to attack an SSM source, S, may be to cause S to send redundant packets to arbitrary SSM channels. That is, an adversary may attempt to send join messages to a number of different SSM channels expecting the edge router at S's site to deliver these joins to S. However, the edge routers will only send SSM-Ping Request packets to S for the join messages going to the SSM-PING.MCAST.NET SSM address. Therefore, adversaries cannot use such an approach to cause remote SSM sources to misbehave.

A final security threat that adversaries may use is to send a large number of join requests to a large number of SSM channels on a remote end system, S. Since each join request consumes router memory space (to store forwarding state information), a large number of such join requests may potentially saturate the forwarding state table space in the routers. These attacks are well-known by the multicast community and our SSM-Ping mechanism does not aggravate the security weaknesses of the existing system. Currently, there are several approaches to detect and reduce the effect of such attacks. More specifically, most router vendors have implemented rate limiting mechanisms in their routers to defend against this type of attack. In addition, a specific protocol, called the Multicast Control Protocol (MCOP) [8], is under development by the Internet Engineering Task Force (IETF) to help fight such attacks. Finally, this problem has attracted attention from academia and solutions [11] have been proposed.

8 SSM-Ping Experiments

In this section we describe our efforts to build a network infrastructure to support SSM-Ping queries between two end systems in two remote networks. In our experiments we use a host at the University of Oregon (UO) network as the pinging host and another host in our UTD network as the pinged host. From a reachability point of view, this experiment tests multicast reachability between the host in the UO network as a potential receiver and the host in the UTD network as a potential source. In our test setup, the UO network is SSM-enabled and therefore has the required support mechanism to initiate the ping request for the pinging host. On the other hand, the UTD network does not have all the functionality that we need. That is, when the ping request arrives at the edge router at the UTD network, according to the current IGMP protocol, this router will not forward the ping request to the end host. As a result in order to successfully implement SSM-Ping, we need to account for this missing functionality.

In order to completely implement SSM-Ping, we need the extra bit of functionality at the designated router at the pinged end system. For this purpose we used an open source software router provided by the eXtensible Open Router Platform (XORP) [1] project. The aim of the XORP project is to develop an open source software router, flexible and extensible enough for research use. The XORP router provides support for the PIM-SM protocol. In our experiment we used the XORP software router as the designated router for the pinged end system. Being a software based router we modified the source code of the XORP router to include the extra bit of functionality to inform the source on receipt of a SSM-Ping request message. The source then sends a response message on the (S, SSM-PING.MCAST.NET) channel.

In the unicast world the ping utility is used to measure the round trip time (RTT) to the pinged site. SSM-Ping includes a similar notion. In multicast the delay experienced by the receiver in receiving the first packet can be called the *multicast* RTT. The expectation is that the multicast RTT will be higher than the unicast RTT. As the SSM-Ping requests are propagated towards the source, routers on the direct path creates multicast forwarding states for this ping request. Since state construction occurs via a router's slow path (CPU-based processing), this operation will introduce delay contributing to a larger RTT value. We ran experiments between UO and UTD to record the multicast RTT. The multicast RTT was measured from the time the receiver

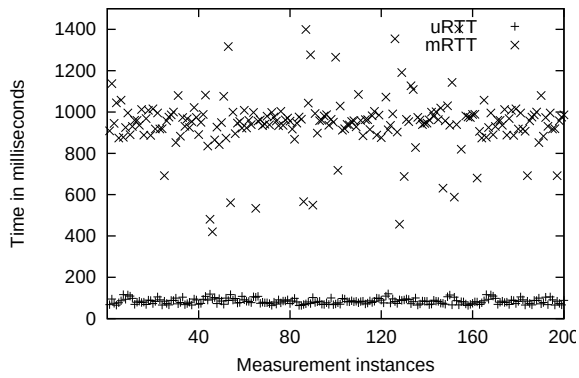

Figure 6. Unicast and Multicast RTT Delays.

sent the SSM-Ping request to the time SSM-Ping response was received. We also measured unicast RTT values along with multicast RTT measurements using the standard unicast ping utility. We ran almost 200 SSM-Ping requests at regular intervals of 10 seconds. Figure 6 shows the comparison between the unicast RTT and multicast RTT. This supports our reasoning earlier, that the multicast RTT is more as compared to the unicast RTT. Note that multicast RTT is different from unicast RTT and multicast RTT gives an estimate for the wait period of a receiver before it receives the first packet.

Based on the experiment, the average unicast RTT is 84ms and the multicast RTT is 953ms. As compared to unicast RTT, the additional processing incurred in multicast RTT is the creation of forwarding states by the routers. According to the path traces, both unicast and multicast paths are same with a end-to-end hop count of 9. If we assume that the propagation delay on the links to be same for both unicast and multicast RTT, the difference between the two delays (i.e. 953 - 84 = 869ms) gives us a good approximate of the total delay incurred by the routers enroute for processing the join message. Hence, on average, 96.5ms (869ms/9) of processing time is introduced by each router. This is a rough estimate of the time taken by each router in processing the join message.

9 Conclusions

In this paper, we introduced an effective utility, SSM-Ping, to support a basic multicast management task, reachability/connectivity testing. First, we motivated our work by discussing the importance of reachability monitoring for the success of SSM-based multicast service deployment efforts in the Internet. Then, we presented our SSM-Ping approach and discussed a number of deployment and security issues with SSM-Ping. Finally, we presented our experiments on evaluating our SSM-Ping utility on a number of important management tasks. Considering the importance of the unicast *ping* utility for end users and network administrators in the unicast world, we believe that our SSM-Ping utility will be a very useful management tool for multicast users and multicast network managers.

10 Acknowledgment

We thank Joel Jaggli and Hans Kuhn of UO for giving us access to their equipment. We also thank Paul Conally and Steve Hodo of the UTD for helping us configure the network. Finally, we thank Pavlin Radoslavov for his help with XORP modifications. This work is partially supported by Cisco Systems through Cisco Systems University Research Program.

References

[1] *eXtensible Open Router Platform (XORP) Project*. http://www.xorp.org/.

[2] K. Almeroth, K. Sarac, and L. Wei. The multicast reachability monitor (mrm) protocol: An instantiation of a multicast management architecture. Technical report, University of California–Santa Barbara, September 1998.

[3] B. Cain, S. Deering, I. Kouvelas, B. Fenner, and A. Thyagarajan. Internet group management protocol, version 3. Internet Engineering Task Force (IETF), RFC 3376, October 2002.

[4] S. Deering and D. Cheriton. Multicast routing in datagram internetworks and extended LANs. *ACM Transactions on Computer Systems*, pages 85–111, May 1990.

[5] C. Diot, B. Levine, B. Lyles, H. Kassem, and D. Balensiefen. Deployment issues for the IP multicast service and architecture. *IEEE Network*, 14(1):10–20, January/February 2000.

[6] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei. Protocol independent multicast sparse-mode (PIM-SM): Protocol specification. Internet Engineering Task Force (IETF), RFC 2362, June 1998.

[7] H. Holbrook and B. Cain. Source-specific multicast for IP. Internet Engineering Task Force (IETF), draft-holbrook-ssm-arch-*.txt, March 2001.

[8] R. Lethonen, J. Soini, J. Majalainen, and H. Vatiainen. Multicast control protocol (MCOP). Internet Engineering Task Force (IETF), Internet Draft, draft-lehtonen-magma-mcop-*.txt, (work in progress), February 2003.

[9] J. Mirkovic and P. Reiher. A taxonomy of ddos attacks and defense mechanisms. *ACM SIGCOMM Computer Communications Review*, 34(2):39–54, April 2004.

[10] NLANR. *Multicast Beacon*. National Laboratory for Applied Network Research, June 2000. Available from http://dast.nlanr.net/Projects/Beacon/.

[11] P. Rajvaidya, K. Ramachandran, and K. Almeroth. Detection and deflection of DoS attacks against the multicast source discovery protocol. In *Journal of Network and System Management, Special Issue on Distributed Management*, 2004.

[12] K. Sarac and K. Almeroth. Monitoring reachability in the global multicast infrastructure. In *International Conference on Network Protocols (ICNP)*, Osaka, JAPAN, November 2000.

[13] K. Sarac and K. Almeroth. Supporting the need for inter-domain multicast reachability. In *Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, Chapel Hill, North Carolina, USA, June 2000.