

End User Level Classification of Multicast Reachability Problems

Pavan Namburi and Kamil Sarac

Department of Computer Science

University of Texas at Dallas

Richardson, Texas 75080

Email: pavan@student.utdallas.edu, ksarac@utdallas.edu

Abstract—In this paper we propose a three step procedure to enable end users to detect and classify multicast reachability problems into one of three groups: (1) connectivity problems, (2) source discovery problems, and (3) multicast join and/or packet forwarding problems. Our approach uses some of the existing mechanisms and requires an additional debugging tool, SSM-traceroute, to classify multicast problems at the end points of the network. In addition, we present a case study in which we use our approach to detect and classify multicast problems for an ongoing multicast application with 40 senders from three vantage points.

I. INTRODUCTION

IP multicast [7] 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. In addition, since the group notion in multicast hides the identities of the receivers, sources may not know who their receivers are and they may not know if their data can reach all the interested receivers. As a result, mechanisms are needed to understand the robustness of the service in the Internet.

One performance metric that is used to measure the robustness of the multicast service is reachability. Reachability gives a measure of the robustness as it is perceived from the user's point of view at the periphery of the network. Reachability can be defined as the availability of multicast data at the group receivers' sites. When a source starts sending to a multicast group, if the packets can make their way to a group receiver site, we say that there is multicast reachability from the source site to the receiver site. In addition, we can talk about the reachability of source to the entire receiver set in a multicast group. As an example, if the source data reaches half of the receiver sites in a multicast group, we say that the source has 50% reachability.

Previous work in the area has focused on monitoring the global reachability characteristics of the multicast infrastructure by using end-to-end application layer information. The main approaches include sdr-monitor [15] and Multicast Beacon [12] projects among several others [2], [3], [4], [14]. These approaches focus on monitoring multicast reachability by using application layer information at the end points of the network. As we will discuss in more detail in the next section,

this approach is useful in detecting the existence of potential multicast reachability problems in the network. However, it cannot really help identify the type of the problems or their root causes. As the multicast service is realized by using a combination of several protocols in the Internet (PIM-SM, MSDP, and MBGP), information collected by these monitoring systems will not be helpful in detecting the protocol responsible for the reachability outages.

In this paper, we introduce a procedure to classify the reachability problems into several different groups and identify their potential root cause. The procedure uses currently available mechanisms and also requires a new debugging primitive to achieve its goal. This work will complement the previous work in the area in that most of the previous work [14], [15] has been on quantifying the amount of multicast problems in the network. The procedure presented in this paper, when completely implemented and available, will help the user to classify these problems. The user can use the tool to debug the problem or at least to understand the source of the problem.

First we divide potential multicast reachability problems into three groups: (1) multicast connectivity problems, (2) source discovery problems, and (3) path establishment and/or forwarding problems. Then, we argue that each of these problems are caused by incorrect/insufficient operation of one of the above mentioned protocols: MBGP for the first problem; MSDP for the second problem; and PIM-SM for the last problem. Finally, by using a simple procedure that we will describe in Section IV, we classify the reachability problems into one of these groups and identify the protocol most likely responsible for them.

We use a variant of the above procedure, using information from the routers to support our findings, to conduct a study to monitor multicast reachability between over 40 diversely located multicast sources and three receiver sites and present a classification of the problems and our reasoning about their root causes. We believe that our monitoring study presents valuable results in understanding the performance characteristics of the various protocols used to build the current multicast protocol architecture.

The rest of the paper is organized as follows. Next section presents the related work. Section III includes an overview of IP multicast service architecture. Section IV discusses a three step approach to aid end user in classifying multicast

related problems. Section V presents SSM-traceroute. Section VI discusses the experiments performed. Finally, the paper concludes in Section VII.

II. RELATED WORK

There has been several tools/systems for conducting end-to-end monitoring of multicast reachability in the Internet [2], [3], [4], [12], [15]. The most widely deployed and used ones are sdr-monitor [15] and Multicast Beacon [12] monitoring systems. These tools are mainly used to help multicast community to get a general understanding about the reachability characteristics of the overall multicast infrastructure. But, they lack the flexibility to reason about the causes of the reachability problems in the multicast infrastructure. In addition, mtrace tool [11] is also used to discover multicast path between a given receiver and a source in a multicast group. Due to their relevance to our work in this paper, we discuss Multicast Beacon and mtrace tools in more detail.

Multicast Beacon [12] has been developed as an active measurement software to monitor multicast reachability among a large number of participants. In Multicast Beacon system, participating end systems use a well-known multicast address to send and receive periodic test multicast packets to and from each other. Multicast Beacon resembles to sdr-monitor in that its monitoring scope is limited by the fixed set of participating multicast end systems in the effort and it is not very useful in identifying the types and causes of multicast reachability problems.

Mtrace [11] is a multicast version of the *traceroute* utility. It is used to discover the multicast path between a given receiver and a source in a multicast group. It returns a snapshot of the set of links used to connect a particular source with a particular receiver. Mtrace attempts to identify the next hop information toward the traced source and uses whatever information is available in the router including existing forwarding state entry or RPF information for the traced source. Once it identifies the information it sends a response back to mtrace originator via unicast. As a result, a successful mtrace verifies the existence of a multicast route (or multicast connectivity) between the receiver and the source. On the other hand, since connectivity does not necessarily mean reachability (i.e., even if a multicast route exists, source data may not reach all the way to the receiver), the information obtained by mtrace cannot really tell us if the receiver is reachable by the source via multicast. Problems due to forwarding errors, inappropriate TTL threshold configurations, or other configuration and interoperability problems may prevent multicast data to reach the receiver. In section V we discuss a new approach to verify multicast reachability between a multicast source and a receiver.

III. OVERVIEW OF IP MULTICAST SERVICE

Currently, IP multicast service is realized by using three protocols including (1) a protocol to construct multicast forwarding trees called Protocol Independent Multicast – Sparse Mode (PIM-SM) [8], [9], (2) the multicast equivalent of the Border Gateway Protocol (BGP) for advertising reverse paths

towards sources called the Multi-protocol Border Gateway Protocol (MBGP) [5], and (3) a protocol for disseminating information about active sources called the Multicast Source Discovery Protocol (MSDP) [13]. In addition, end-hosts use the Internet Group Management Protocol (IGMP)[6], [10] to communicate join and leave requests with the edge routers (i.e., designated multicast routers or DRs) in their subnets.

In PIM-SM, a router is identified as a Rendezvous Point (RP) to connect sources and receivers to each other in a domain. When a receiver joins a group G, its designated router (DR) forwards a PIM Join(*,G) control message toward the RP and all the routers on the DR-to-RP path establishes forwarding states for the group G. This way PIM creates a shared forwarding path between the RP and the DR of all joining receivers. On the other hand, when a source S starts sending multicast data to group G on (S,G), its DR forwards these packets to RP which then forwards them toward the receivers of group G by sending data on the shared tree (*,G). The DR routers at the receiver sites, on receiving (S,G) packets on the shared (*,G) tree come to know about an active source S and then optionally switch to (S,G) source specific tree by sending a PIM Join(S,G) control message toward S. This way the DR switches from the shared tree to the source specific tree of the source S.

In the above discussion, the source S is assumed to be in the same domain as the receivers. Alternatively if the source S is in some other domains, the RP of that domain is responsible for announcing the existence of this active source to RPs in all other domains so that the RPs in these other domains can inform their local receivers about the existence of a remote source for the group. This is achieved by using the MSDP protocol where an RP on receiving (S,G) data from a local source S creates MSDP Source Announcement (SA) messages and sends it to its peering RPs in neighboring domains. All the RPs in different domains form an MSDP overlay network to flood SA messages to entire multicast infrastructure to inform potential receivers of multicast groups about the existence of remote sources for their groups. At the end, the DRs at these remote receiver sites issue PIM Join(S,G) control messages to join the source specific groups of these remote sources and get their data.

IV. CLASSIFYING MULTICAST REACHABILITY PROBLEMS

In this section we present a procedure to classify multicast reachability problems into three groups and point out potential reasons causing the problems in each group. First we divide potential multicast reachability problems into three groups: (1) multicast connectivity problems, (2) source discovery problems, and (3) path establishment and/or data forwarding problems. Then, we argue that each of these problems are caused by incorrect/insufficient operation of one of the above mentioned protocols: MBGP for the first problem; MSDP for the second problem; and PIM-SM for the last problem.

In a typical multicast application scenario, a receiver joins the multicast group address and expects to start receiving packets that are sent to the group address by the active sources

in the group. Assume that we have a receiver R who wants to join a multicast group G to get data on the group address. For this, first the receiver R uses IGMP protocol to inform the local DR router about its request to join multicast group G. If everything goes well, the receiver R will soon start receiving multicast packets from the network. On the other hand, if the receiver does not receive any data, there are two possibilities: (1) there is no active source sending to the multicast group G or (2) there is a source S who is sending to (S,G) address but due to reachability problems R cannot receive these packets. In our work, we are mainly interested in the second scenario that there exists a source S that is currently sending to (S,G). At this point, we enumerate potential reachability problems as follows:

- 1) **Multicast connectivity problems:** One possibility is that the routers in the receiver's local domain do not have a multicast route toward the domain that the source S resides. Since multicast route availability is communicated using MBGP protocol, we conclude that this is a potential problem with MBGP.
- 2) **Source discovery problems:** Another possibility is that the RP at the receiver's domain does not have any information about S being an active source for the group G. Since MSDP is the protocol used to communicate Source Announcement messages, we conclude that this is a potential problem with MSDP.
- 3) **Multicast join and/or data forwarding problems:** A third possibility is that the local domain has a multicast route to S and the local RP has an entry for S in its MSDP table and the RP has already sent out a PIM-Join(S,G) toward the source S but still R cannot receive (S,G) packets from the network. This situation may be related to problems in forwarding PIM-Join message toward the source S or it may be related to forwarding errors for the actual (S,G) packets enroute to R's site. In order to distinguish this case from the previous two cases, we call these errors as multicast join and/or forwarding errors and relate them to PIM-SM. In the paper whenever we categorize a problem related to PIM-SM, we essentially encompass problems related to both forwarding of PIM-Join and multicast data packets.

After enumerating the potential multicast reachability problems, in the next step we are interested in classifying these problems into one of the three groups above. Our motivation in this classification is to identify the multicast routing protocol that is responsible for the multicast reachability problems. Since our research background has been in monitoring multicast reachability using the sdr-monitor tool, with this classification we will be extending our monitoring efforts one step further by introducing an ability to classify the detected problems into one of the above three groups.

The procedure that we use to classify multicast problems into different groups is as follows. The typical scenario is that there is an active source S that is currently sending to a multicast group G. A remote receiver R wants to join the

group G and receive (S,G) data originating at S. We follow the below steps to classify potential problems:

- 1) First, the DR at R's site issues a PIM-Join(*,G) toward the domain RP. If R starts receiving (S,G) data, this means that there is no reachability problems and S's multicast data can reach R's site.
- 2) If the receiver R cannot receive any data, then it issues a PIM-Join(S,G) toward the source S's site. If the receiver R starts receiving (S,G) data, this means that there exists a multicast route to S's site and the new join process was successful. Given the fact that (*,G) join failed and (S,G) join succeeded, we conclude that there is a potential MSDP problem between the two domains. Please remember that MSDP is the protocol that is used by the local RP at the sender S's domain to inform remote RPs about the active source S. In this scenario, the failure to receive (S,G) data after the initial (*,G) join indicates that the RP at R's site did not know about S as being an active source for the multicast group G.
- 3) On the other hand, if the receiver R cannot receive (S,G) data even after joining the source specific group of S, it may mean two things: (1) there is no multicast route to S's domain and therefore the local routers do not know how to forward the join request or (2) there is a problem with join procedure and/or multicast data forwarding. In order to detect the problem we introduce a new user level debugging tool called SSM-traceroute below.

V. SSM-TRACEROUTE

In this section, we introduce a new multicast debugging tool, SSM-traceroute, that can be used to verify multicast reachability between a source and a receiver. As discussed in Section II, mtrace is useful in verifying the multicast connectivity but it cannot always help us decide on reachability (i.e., connectivity does not always mean reachability).

SSM-traceroute uses a dedicated multicast group address, SSM-TR.MCAST.NET in the source specific multicast address range of 232/8. SSM-traceroute is run from a receiver toward a source to collect the multicast path information in between. The main difference between SSM-traceroute and mtrace is that the former uses the underlying multicast service to communicate the request and response messages and the latter uses a unicast based response.

In SSM-traceroute, a receiver R causes its DR to issue a PIM Join(S,SSM-TR.MCAST.NET) to join the source specific group of a remote source S. This can be done by issuing an IGMPv3 Membership Report or if IGMPv3 is not available, the SSM-traceroute client can initiate PIM neighbor relation with its DR and sends a PIM-Join(S,SSM-TR.MCAST.NET) request to its DR. The DR then forwards this join request toward the source S. This is visually represented in Figure 1-a.

Then SSM-traceroute uses a TTL-based approach to cause each router on the R-to-S multicast path to send SSM-traceroute responses on the (S,SSM-TR.MCAST.NET) channel. The first SSM-traceroute request is sent out with an initial

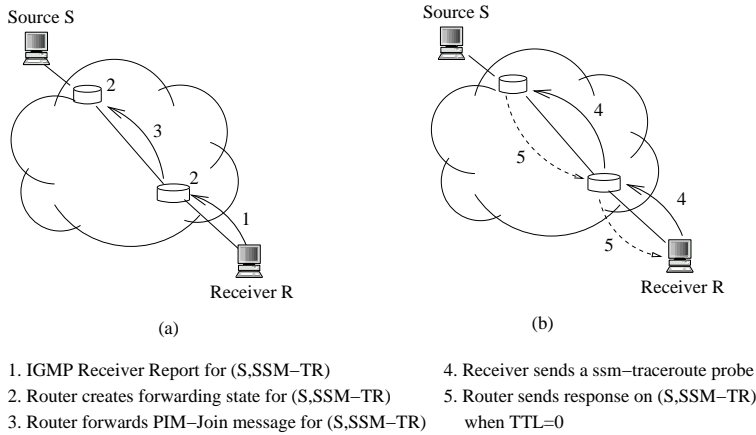


Fig. 1. Operation of SSM-Traceroute.

TTL value of 1 and the TTL is incremented at each round afterwards (see Figure 1-b). This request propagates towards the source along the join path and each router decrements the TTL along the path. When TTL reaches 0 at an on-tree router, the router sends an SSM-traceroute response to (S,SSM-TR.MCAST.NET) channel spoofing the IP address of the source S. The router includes its own IP address into a protocol field in the response message. Source spoofing allows the SSM-traceroute response to propagate on the existing multicast forwarding path toward the receiver. When the receiver receives an SSM-traceroute response originating from an on-tree router X, it indicates that there is no reachability problems between the router X and the receiver. Eventually, if receiver stops receiving SSM-traceroute responses before reaching the source S but after an on-tree router Y, it indicates that there is a potential reachability problem at/after Y. As a result, we argue that SSM-traceroute can be used to detect and locate multicast reachability problems between a source and a receiver. Finally, both SSM-traceroute request and response messages can be implemented by introducing new messages to IGMP protocol.

In our problem classification procedure, we use SSM-traceroute to identify problems as MBGP problems or PIM-SM problems. First, we make sure that the receiver does not have a local connectivity problems. This can be achieved by running SSM-traceroute to a domain-local source site. Then, the receiver R starts an SSM-traceroute toward the source S. If the receiver cannot receive SSM-traceroute responses from routers beyond its DR, we conclude that this is due to a potential MBGP problem. This is so because if there were a multicast route toward the source S, then the initial PIM Join(S,SSM-TR.MCAST.NET) would be forwarded toward the source site S and the subsequent SSM-traceroute requests would cause on-tree routers to send responses back. Alternatively, if the receiver gets some number of SSM-traceroute responses back, it would indicate existence of a multicast route toward the source and hence the proper functioning of MBGP. In this case, if the receiver cannot receive SSM-traceroute responses all the way from the source S, then we conclude that the problem is a potential multicast join and/or data forwarding

problem and attribute it to PIM-SM.

VI. EXPERIMENTS

In this section, we present our sample experiments on the utility of the proposed approach in classifying multicast reachability problems. For this, we use the existing Multicast Beacon active monitoring system as follows. Currently, there are over 40 participants in Multicast Beacon project. These participants are actively sending test multicast messages to the well known Multicast Beacon group address 233.4.200.21, port 10002. The Multicast Beacon web site continuously presents reachability information among these sites. In our case, we use the web site to identify the active participants and use our end user based approach to classify potential reachability problems between these sites as multicast senders and three individual sites that we have access to as multicast receivers. The three sites are University of Oregon (UO), University of California Santa Barbara (UCSB), and University of Texas at Dallas (UTD).

Remember that our approach includes three steps: issuing (*,G) join, issuing (S,G) join, and using SSM-traceroute which is yet to be developed user level debugging tool proposed in this paper. Due to the unavailability of SSM-traceroute, our experiments will mostly depend on the first two steps. However, for the case of UTD receiver, we will use support from the local RP router in the form of collecting MBGP and MSDP routing table information for the Multicast Beacon sources. We will use MSDP information to verify the correctness of our two step approach and use MBGP information as a placeholder for the third step of our procedure.

One practical issue with our approach is the ability of issuing (S,G) joins at our receiver sites. Out of the three sites, UO network is IGMPv3 enabled network and therefore includes support for issuing (S,G) joins easily. On the other hand, both UCSB and UTD networks are IGMPv2 enabled which does not support (S,G) joins directly (i.e., receivers cannot directly communicate the source information to their DR routers using IGMPv2). As we mentioned previously, this difficulty can be overcome by having the end system

| Source | UTD | UCSB | UO | Source | UTD | UCSB | UO | Source | UTD | UCSB | UO |
|-----------------|-----|------|----|-----------------|-----|------|----|-----------------|-----|------|----|
| 129.78.157.172 | - | - | - | 129.128.125.62 | ✓ | ✓ | ✓ | 206.167.204.18 | - | - | - |
| 132.246.2.20 | ✓ | ✓ | ✓ | 204.174.103.32 | - | - | - | 132.246.130.26 | - | - | - |
| 142.55.1.205 | - | - | ✘ | 129.128.25.72 | ✓ | ✓ | ✓ | 129.128.25.181 | ✓ | ✓ | ✓ |
| 216.239.127.230 | ✘ | ✘ | ✘ | 63.105.122.14 | - | - | - | 192.108.35.16 | ✓ | - | - |
| 128.123.3.74 | ✓ | ✓ | ✓ | 128.118.146.51 | ✓ | ✓ | ✓ | 128.118.146.52 | ✓ | ✓ | ✓ |
| 128.118.57.33 | ✓ | ✓ | ✓ | 130.160.4.113 | - | - | - | 128.111.252.50 | ✓ | ✓ | ✓ |
| 128.111.2.2 | ✓ | ✓ | ✓ | 137.110.147.70 | ✓ | ✓ | ✓ | 132.239.253.141 | ✓ | ✓ | - |
| 128.227.212.96 | ✓ | ✓ | ✓ | 131.193.77.102 | ✓ | ✓ | ✓ | 141.142.2.168 | ✓ | ✓ | ✓ |
| 141.142.64.5 | ✓ | ✓ | ✓ | 128.223.157.25 | ✓ | ✓ | ✓ | 192.236.37.104 | - | ✓ | ✓ |
| 155.101.3.111 | ✓ | ✓ | ✓ | 128.83.6.240 | ✓ | ✓ | ✓ | 160.36.188.124 | - | - | - |
| 198.82.169.70 | ✓ | ✓ | ✓ | 198.82.169.72 | ✓ | ✓ | ✓ | 130.215.32.94 | ✓ | ✓ | ✓ |
| 130.215.5.21 | - | - | - | 130.215.201.81 | ✓ | ✓ | ✓ | 193.166.3.92 | ✓ | - | - |
| 192.31.96.42 | ✓ | ✓ | ✓ | 203.181.248.186 | - | - | - | 203.181.249.74 | - | - | - |
| 205.189.33.130 | - | - | - | 138.18.250.6 | ✓ | ✓ | ✘ | 194.80.35.36 | ✓ | - | - |
| 195.194.24.19 | - | ✓ | ✓ | | | | | | | | |

TABLE I
RECEPTION TO (*,G) JOINS AT DIFFERENT CAMPUSES

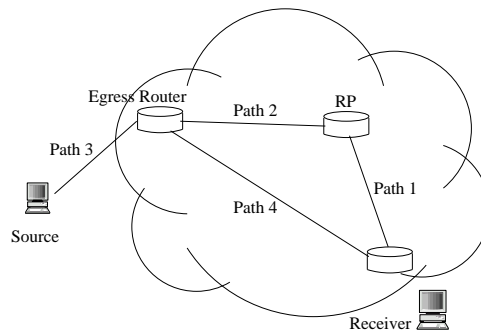


Fig. 2. (*,G) and (S,G) Join Paths.

| Source | MSDP | MBGP | PIM-SM related | Source | MSDP | MBGP | PIM-SM related |
|----------------|------|------|----------------|-----------------|------|------|----------------|
| 129.78.157.172 | | ✓ | | 206.167.204.18 | | ✓ | |
| 204.174.103.32 | | ✓ | | 132.246.130.26 | | | ✓ |
| 142.55.1.205 | | | ✓ | 63.105.122.14 | | ✓ | |
| 130.160.4.113 | | ✓ | | 160.36.188.124 | | | ✓ |
| 130.215.5.21 | | | ✓ | 203.181.248.186 | | | ✓ |
| 203.181.249.74 | | | ✓ | 205.189.33.130 | | | ✓ |
| 194.80.35.36 | | | ✓ | 216.239.127.230 | ✓ | | |

TABLE II
CLASSIFICATION OF PROBLEMS AT UTD

create PIM neighbor relation with its DR which can be easily achieved by sending periodic PIM Hello messages on the subnet by the receiver. Then the (S,G) join can be initiated directly from the receiver by creating a PIM Join(S,G) and passing it to its DR. The DR router then forwards it toward the source S according to PIM protocol. In our experiments, we used an open source software router provided by eXtensible Open Router Platform (XORP) [1] project. The XORP router provides support for the PIM-SM protocol and includes client interface to initiate source specific joins. In our experiments we used the XORP router to implement the above mentioned missing functionality at UCSB and UTD receivers.

Table I lists the IP addresses of Multicast Beacon sources that we used in our experiments. From each of our three receiver sites, we first issued (*,G) joins for the Multicast

Beacon group and then issued (S,G) joins for each of the active sources listed in the table. In the table, ✓ indicates that the receiver was able to receive multicast data from the corresponding source after a (*,G) join as well as after a (S,G) join. As a result, these cases corresponds to a properly working multicast service between the source and the receiver. A '-' character indicates that the receiver did not receive any multicast data from the corresponding source after both (*,G) join and (S,G) join. At this point, without any further information, we cannot really know if this is a connectivity problem or others. Finally, ✘ indicates that the receiver did not receive multicast data from the corresponding source after a (*,G) join but did receive multicast data from the same source after a (S,G) join. This case suggest that there exists a reachability problem between the source and the receiver

and this problem is most likely an MSDP problem. We now elaborate on this case.

When (*,G) join does not result in data reception but (S,G) join does, one can think about several reasons for the outcome: (1) a potential PIM-SM problem between the DR at the receiver site and its RP caused (*,G) join to fail before reaching the RP, (2) a potential PIM-SM problem between the RP and the source S caused (S,G) join issued by the RP to fail before it reached the source S, or (3) due to a potential MSDP problem the local RP did not know about the source and therefore did not join its (S,G) source specific group.

Consider the topology in Figure 2 where we explicitly mark important segments of the paths from Path1 to Path4. The first reasoning above can be shown wrong as follows: During our experiments, we used over 40 sources and for majority of these sources the (*,G) joins returned multicast data from these sources. This suggests that there does not exist any PIM-SM problem on Path1. Similarly the second reasoning above can be shown wrong as follows: Following from the first case, the fact that (*,G) joins returned multicast data from the majority of the sources suggests that there does not exist a PIM-SM problem on Path2. In addition, the fact that (S,G) join returned multicast data for the source S suggests that there does not exist a PIM-SM problem on Path3. These two observations suggest that the second reasoning above cannot be correct. Hence, we conclude that the problem is an MSDP problem. In order to verify the correctness of our reasoning, we used the only case in our experiment in which the source 216.239.127.230 was not reachable via (*,G) join but was reachable via (S,G) join and verified that the local RP did not have an MSDP cache entry for this source as an active source for Multicast Beacon group address. Even though we had several other such cases with UO and UCSB, due to lack of routing table information from the corresponding RPs we could not use these cases to verify our reasoning.

Finally, since SSM-traceroute is not currently available, we instead consulted the MBGP routing table at UTD to categorize problems either MBGP or PIM-SM related problems. Table II shows the final classification for UTD receiver. The listed sources are from Table I which failed on (*,G) joins from UTD. We see five instances of failures due to MBGP and eight instances related to incorrect functioning of PIM-SM either due to protocol problems or configuration problems as categorized in Section IV. After cross checking the sites having PIM-SM problems with Multicast Beacon web site [12], we see that two of these sites (130.215.5.21 and 194.80.35.36) have local connectivity problems and others seem to have reachability to some other beacon sites beyond their own network.

VII. CONCLUSIONS AND FUTURE WORK

In this paper we proposed a three step mechanism for end users to detect and classify multicast reachability problems between themselves and remote sources. Our mechanism uses currently available multicast join primitives and requires a new user level multicast debugging tool, SSM-traceroute, that we

introduce in this paper. Currently, mtrace is used to detect and locate potential multicast connectivity problems between sources and receivers. We believe that our SSM-traceroute tool is a more effective debugging tool as it can detect and locate multicast reachability problems between sources and receivers. By using our three step procedure, we presented an experimental study in which we have detected and classified multicast reachability problems between over 40 remote sender sites and three vantage points. As a result, we believe that the mechanism presented in this paper improves end users' ability to monitor and manage multicast service. We plan to extend this work to build a more powerful debugging tool that will automate multicast problem classification as presented in this paper. We will also extend it to support reachability verification and classification to remote sites without requiring to have an active source at these sites.

VIII. ACKNOWLEDGEMENT

This work is partially supported by Cisco Systems through Cisco Systems University Research Program.

REFERENCES

- [1] *eXtensible Open Router Platform (XORP)*. Available from <http://www.xorp.org>.
- [2] Ehab Al-Shaer and Yongning Tang. Mrmon: Multicast remote monitoring. *IEEE/IFIP Network Operations and Management Symposium*, April 2004.
- [3] Ehab Al-Shaer and Yongning Tang. Smrm: Snmp-based multicast reachability monitoring. *IEEE/IFIP Network Operations and Management Symposium NOMS*, April 2002.
- [4] 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.
- [5] T. Bates, R. Chandra, D. Katz, and Y. Rekhter. Multiprotocol extensions for BGP-4. Internet Engineering Task Force (IETF), RFC 2283, February 1998.
- [6] S. Deering. Host extensions for IP multicasting. Internet Engineering Task Force (IETF), RFC 1112, August 1989.
- [7] S. Deering and D. Cheriton. Multicast routing in datagram internetworks and extended LANs. *ACM Transactions on Computer Systems*, pages 85–111, May 1990.
- [8] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, G. Liu, and L. Wei. PIM architecture for wide-area multicast routing. *IEEE/ACM Transactions on Networking*, pages 153–162, Apr 1996.
- [9] 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.
- [10] W. Fenner. Internet group management protocol, version 2. Internet Engineering Task Force (IETF), RFC 2236, November 1997.
- [11] W. Fenner and S. Casner. A 'traceroute' facility for IP multicast. Internet Engineering Task Force (IETF), draft-ietf-idmr-traceroute-ipm-*.txt, July 2000. Work in progress.
- [12] M. Kutzko and T. Rimovsky. NLNAR multicast beacon project. <http://dast.nlanr.net/Projects/Beacon/>.
- [13] D. Meyer and B. Fenner. Multicast source discovery protocol (MSDP). Internet Engineering Task Force (IETF), RFC 3618, October 2003.
- [14] P. Rajvaidya and K. Almeroth. Analysis of routing characteristics in the multicast infrastructure. In *IEEE Infocom*, San Francisco, CA, USA, April 2003.
- [15] K. Sarac and K. Almeroth. Monitoring reachability in the global multicast infrastructure. In *International Conference on Network Protocols (ICNP)*, Osaka, JAPAN, November 2000.