Statistical Characteristics of Multicast Traffic on a National Backbone Network

Tao He, Xing Li, Jian Qiu
Department of Electronic Engineering
Tsinghua University, Beijing, 100084, China
Telephone: +86-10-62792515
hetao@serv.edu.cn,xing@cernet.edu.cn
qiujian97@mails.tsinghua.edu.cn

Hui Zhang, Zhichun Li
CERNET Network Research Center
Tsinghua University, Beijing, 100084, China
Telephone: +86-10-62784301
hzhang@cernet.edu.cn,lizc@serv.edu.cn

***Abstract*— IP multicast catches Chinese researchers' eyes recently as the deployment of non-tunnel multicast routing protocols throughout the CERNET mature. But characteristics of multicast traffic still need to be understood. Using our developed passive monitoring system, we observe multicast traffic on links connecting peer networks to our native multicast backbone network. First of all, we analyze of collected multicast traffic data on CERNET. Then we pursue a traffic traces collection over seven months. Comparing the experimental results with the analysis of unicast, we describe multicast traffic characteristics: packet size distribution, stream size distribution and address space distributions. We analyze the distribution of source numbers per group, peep into the content of multicast traffic and propose a hierarchical structure for carrying out our measurement and monitoring.**

I. Introduction

Over the last decade, a myriad of new Internet applications have evolved that require transfer of high bandwidth media streams to a large numbers of users. Traditional Internet protocols are mostly unsuitable in handling such high bandwidth traffic flows. Several revolutionary developments have been taken place. Among these, one solution is multicast communication. Multicast is mechanism for one-to-many and many-to-many delivery of data over the Internet in a bandwidth efficient manner. During the first five to six years of its development multicast existed as a virtual and experimental network on top of the existing Internet. This topology called the Multicast Backbone(MBone). However recently multicast has evolved into a mature network service that providers are now deploying. In order to bridge the gap between the initial deployment experiments and the availability of multicast as a robust network service, there needs to be a full complement of multicast measurement, monitoring and management.

China Education and Research Network (CERNET), which provides high speed interconnection among universities, research and education institutions throughout China, began offering native multicast to the research and education community in 2001. Research about multicast started in parallel [5], including reliable multicast, monitoring video conference and multicast traffic measurement. Additional details about CERNET will be given in Section III.

This paper presents our multicast traffic measurements and analysis taken on the CERNET. We first introduced a new traffic stream profiling methodologies, which is not based on packet-level traffic measurement. In [6], Using STM-4 passive monitors, R and K observed multicast traffic and presented multicast-specific characteristics. But considering a long time measurement from an IP Backbone, it is a big issue in capturing, storing and analyzing such huge packet-level traffic data. We adapted our unicast traffic measurement methodology mentioned in [4] for monitoring multicast traffic and It is cheaper and simpler than similar system, such as IPMON system in [7].

The rest of this paper is organized as follows. Section II summarize related work. Section III provides an overview of the CERNET network to provide context for the work. Section IV presents a detailed traffic analysis. Finally, Section V and VI concludes the paper and gives suggestions for future work.

II. Relative Work

In this section, we look at the early multicast measurement and monitoring tools for various multicast network monitoring tasks. We divide previous tools according to three different input data: router information

data, passive captured data and active-produced data. In the following paragraph, we'll find out that the ideal suite of multicast measurement tools does not exist today, because much of the effort [8] being to devoted to multicast measurement is aimed for developing multicast protocols, collecting information from multicast-enabled routers, and monitoring the connectivity among users in end networks.

Starting with the initial deployment of the Multicast Backbone (MBone) in 1992, mrinfo could report on the tunnels and multicast-enabled interfaces for a router or end-host running multicast routing code. Information returned by mrinfo includes the set of tunnels and/or interfaces on which multicast is enabled or disabled.

Mtrace (Multicast Traceroute) returns a snapshot of the set of links used to connect a particular destination. A trace query is passed hop-by-hop along the reverse path from the receiver to the source, collecting hop addresses, packet counts, and routing error conditions along the path, and then the response is returned to the requester. Additional information that can be obtained includes loss rates along the links, and the number of multicast packets flowing across each hop per second for that particular address. The mtrace tool is one of the best ways of discovering the flow of multicast packets through a network [21].

The Multicast Beacon [19] is an active measurement tool used to monitor the performance of current multicast transmission on the network. It does this by injecting a steady stream of probes into a multicast group, and measuring performance information when the probes arrive at other beacons. It has two components: server and client. A set of measurement clients send small probe packets to a particular multicast session, and also receive packets from the session in order to determine session transfer performance. Each probe packet contains the originating Beacon Client's name, a sequence number, and a time stamp so that the receiving client is able to calculate the packet statistics. The Beacon Clients' information are collected by the central Beacon Server.

Mhealth and MultiMON are two opposite approaches to present information about multicast traffic. The former gives the evaluation of the end-to-end performance of a particular multicast group, and the later presents information on all of the multicast traffic flowing on a particular LAN [21].

MRM and HPMM implement a new multicast monitoring system which aimed to provide support for both intra- and inter-domain multicast monitoring tasks. But both of them need additional protocol and may introduce scalability problems [3].

III. CERNET Overview

CERNET is the first and largest nationwide education and research computer network in China. Its backbone consists of over 60 STM-16 and STM-1 links, interconnecting 10 region-level nodes and 38 province-level nodes. The traffic volume ranges from tens of MB/s on STM-1 province access links to more than 1Gb/s on STM-16 national backbone links.

The CERNET backbone IP network provides connectivity over a geographically wide area. The backbone consists of a set of regional nodes connected by high bandwidth links, which are typically Gigabit Ethernet (GE), connecting to region access aggregation routers which provide access service for downstream networks. On the other hand, the BGP border routers which connect CERNET to Internet outside China and other major ISPs in China are generally connected to backbone via GE links too. In case of CERNET topology, we can therefore easily deploy multiple monitoring agents across those GE links to manage the whole network as well as measure network traffic. For more information please check [20].

IV. Multicast Traffic Analysis

This section discusses multicast traffic characteristics of our traffic traces. We first give details of the tests whose results are reported in this paper. Second, we present the overall multicast traffic profile. Finally, more attention are payed to the attributes of multicast source and group.

A. Test Details

Using our multicast monitoring system, we capture all multicast packet headers from our monitor point. After had been processed packet-level information, multicast traffic traces were aggregated so we got stream-level data. The stream-level data is coarse-grained and little in volume which requires less storage space on disk. And it is also simple in format which reduces complexity in post-processing. By aggregating packet-level data into stream-level data, we only need about 20G Bytes disk space to store the final data instead of about 100G packets which could occupy about tens of Terabytes disk space.

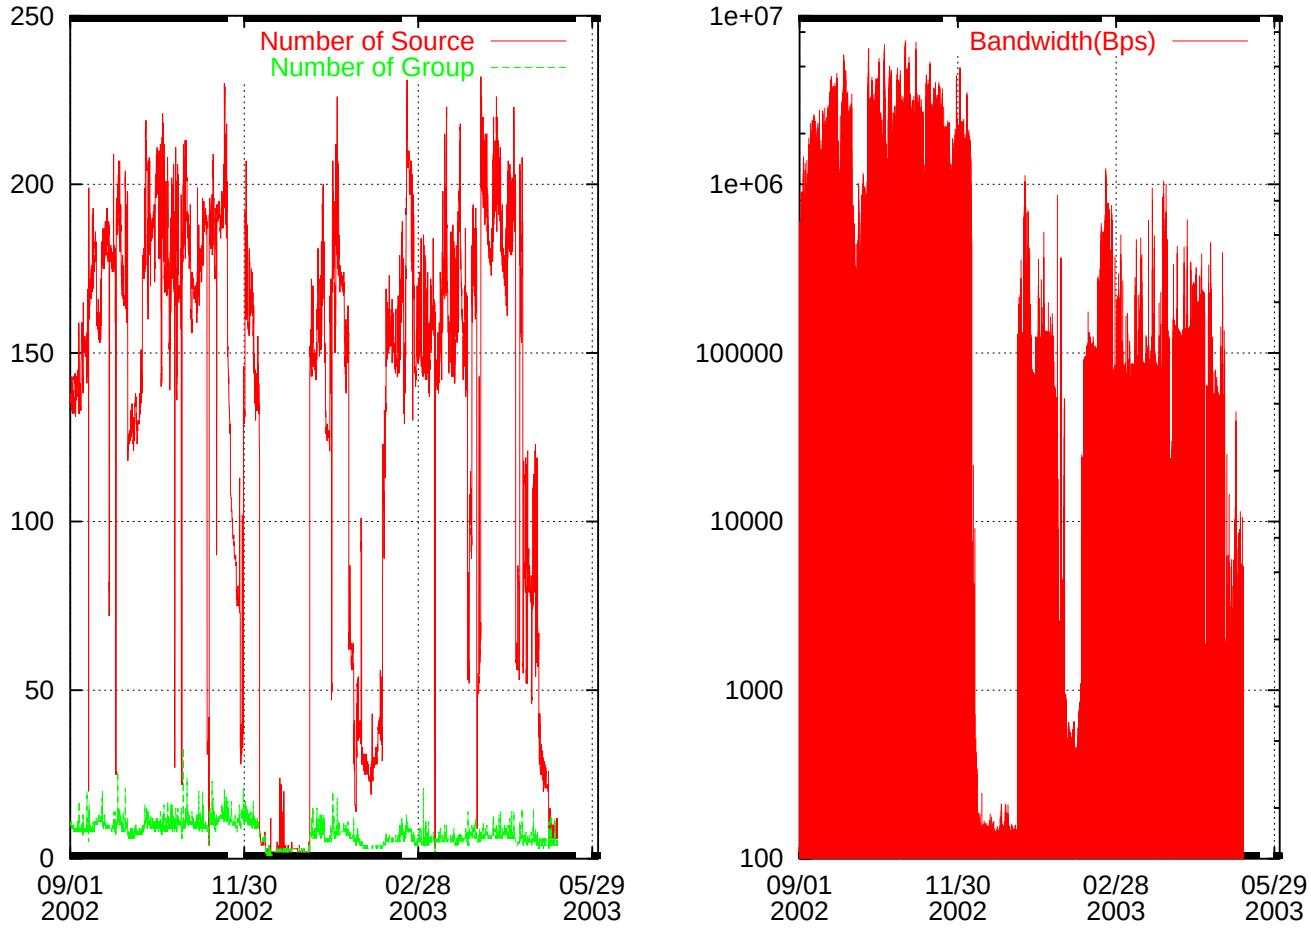

Fig. 1.

B. Overall Multicast Traffic Profile

The right graph in Figure 1 show traffic volume in Bytes per second, multicast source number variation and group number variationobserved at the measurement point. Based on special source address belonging to CERNET, similar graphs are also depicted in Figure 2. The Bytes per second graph show the volume in the link, while the source number and group number variation graphs depict the long-term disturbance. Comparing these two figures, we could draw a conclusion. Even though the traffic volume which produced by those source addresses which belong to CERNET occupies little in total, the group numbers in which hosts in CERNET participate and to which sources in CERNET send multicast traffic are approximately equal in amount. Another explicit change between two figures is the source number variation. Opposite to macroscopic stability of total source number we observed, source number in CERNET descended in the last ten-day of December 2002. And both figures show that strong disappearance of traffic volume.

Being compared with unicast traffic in our network, multicast traffic shows little correlation with it. But both of them showed us the same phenomena in traffic variation graphs. From the graphs in Figure 1 and Figure 2, we could get the indication of abnormal network traffic which peak among plate line. After analyzed the peak of line, we concluded that the sudden change of stream number in multicast traffic was caused by huge connections took place between multicast source address and group address – even though the two addresses may not be real.

C. Multicast Source and Group

Next, we pay more attention to the most popular blocks: the public 224/8 and the GLOP 233/8, the two address blocks occupy almost 90% traffic volume, source number and group number in total. In Figure 4 and 3 and 5, we use red line to present traffic in 224/8 and green

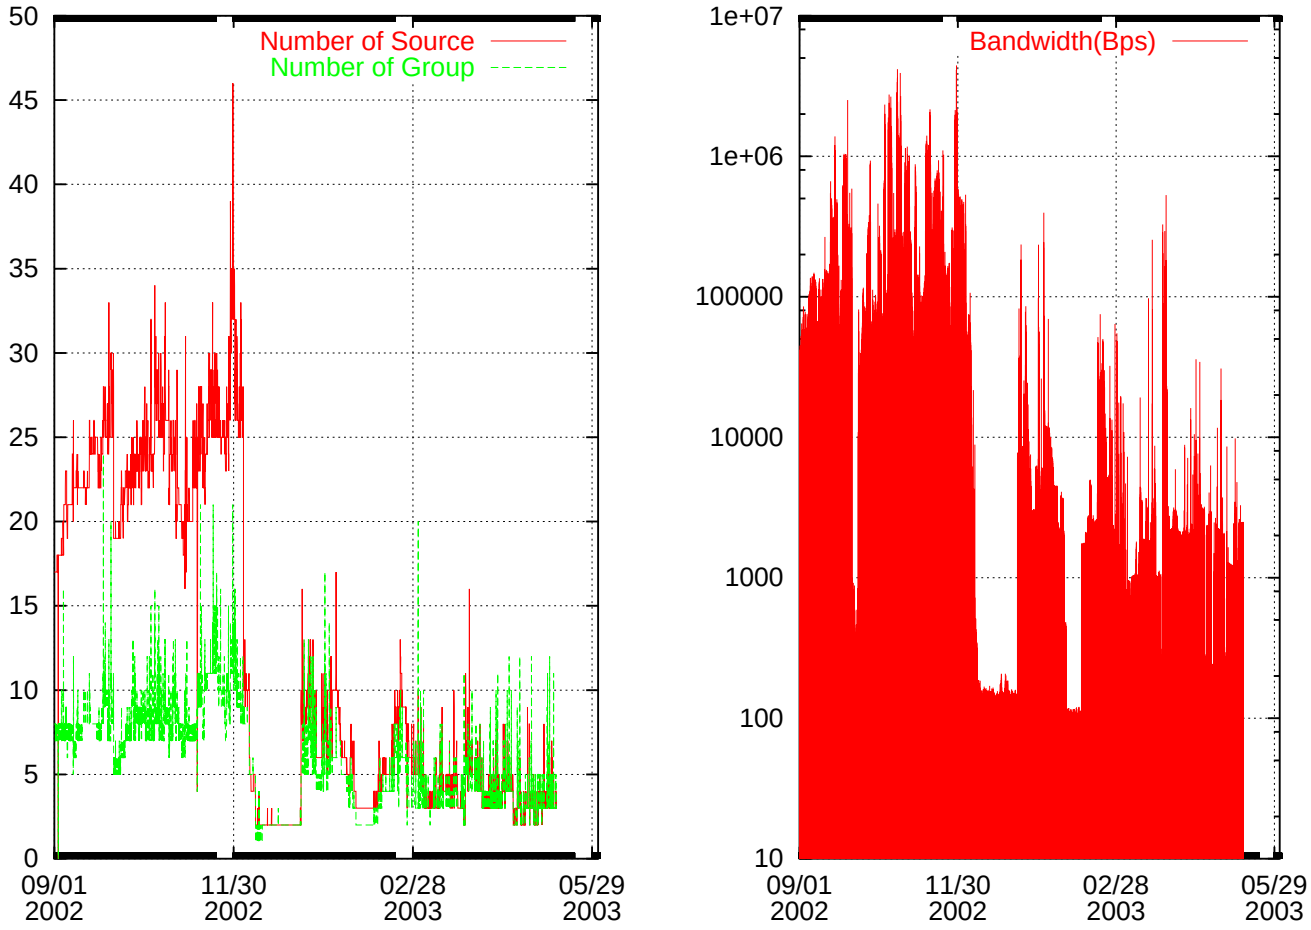

Fig. 2. Source address belonging to CERNET

line to present traffic in 233/8. From the upper graph in these figures, we find out that though the group number in 233/8 is smaller than that in 224/8, the traffic volume in 233/8 is larger than that in 224/8. We could conclude that the source number is more important in generating multicast traffic than group number.

From the Figure 6, we review the group address distribution from the view point of traffic volume. Most multicast traffic only take up in 224/8 and 233/8. We are strongly impressed by the sparse usage of group address observed in the backbone.

We also give the packet size distribution in Figure 7. The graph in figure exhibits strong mode at 60 bytes with smaller modes at 500 and 1250 bytes. It is different from what was given in R and K's work [6], and also different from what was shown in unicast packet size distribution. For this reason, we think that it will bound router performance if multicast demand increases and multicast packets traces explode. Another reason what surprise us is that those modes what are most common in unicast packet size distribution disappear, including modes at 40 bytes due to TCP acknowledgement segments and at 1500 bytes Maximum Transmission Unit (MTU) of Ethernet-attached hosts. Figure 7 depicts the cumulative packet size distribution in each direction. Like unicast traffic, the multicast traffic graph shows great symmetry.

Over the seven months monitoring period, we find over 50% of the groups observed ever had multiple, simultaneous sources. considering that the overwhelming majority of groups are single-source, and that one of the factors impeding wide spread adoption of multicast is the complexity incurred supporting multiple sources, we think that there should be a set of more effective debugging tools to support the deployment of multicast as an important service throughout CERNET.

V. Future Work

As with [10], we describe the CERNET passive monitoring system that is capable of supporting Gigabit Ethernet data rate, and the advantages of this system are

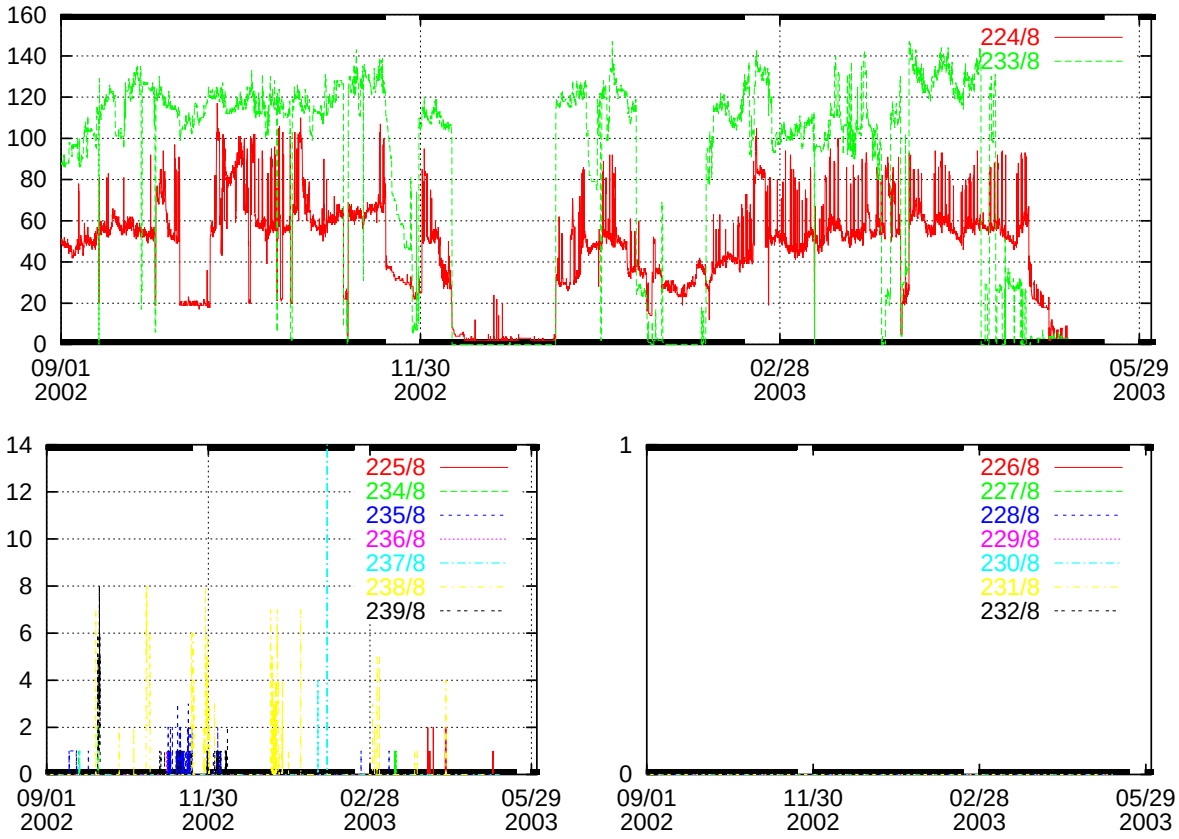

Fig. 3. Souce number statistics divided by group address

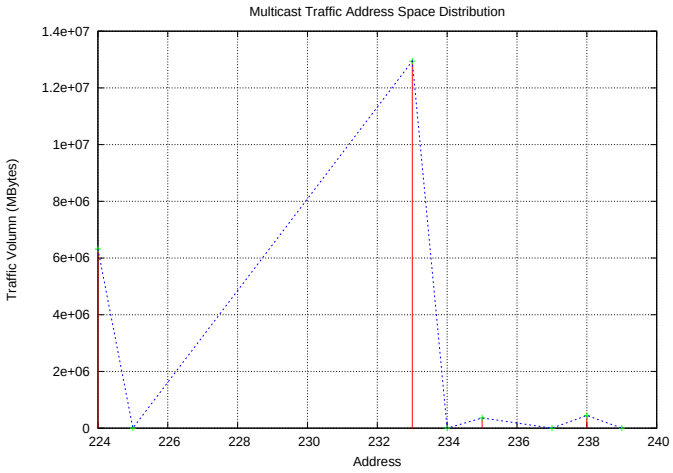

Fig. 6. 224/8 and 233/8 occupy most in total

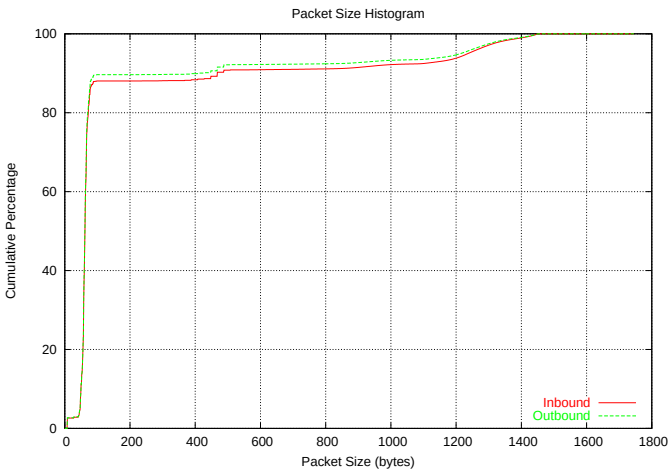

Fig. 7. Packet size distribution

as follows:

- The system equips a powerful collection engine for real-time multicast traffic packet capture on high-speed links.
- It performs multicast traffic packet pattern and stream-based traffic analysis on-line and off-line for network management.
- It could adjust dynamically the packet filter and classification rules, and stream filter rules for accurately identify multicast traffic which we concern in the network with great flexibility.

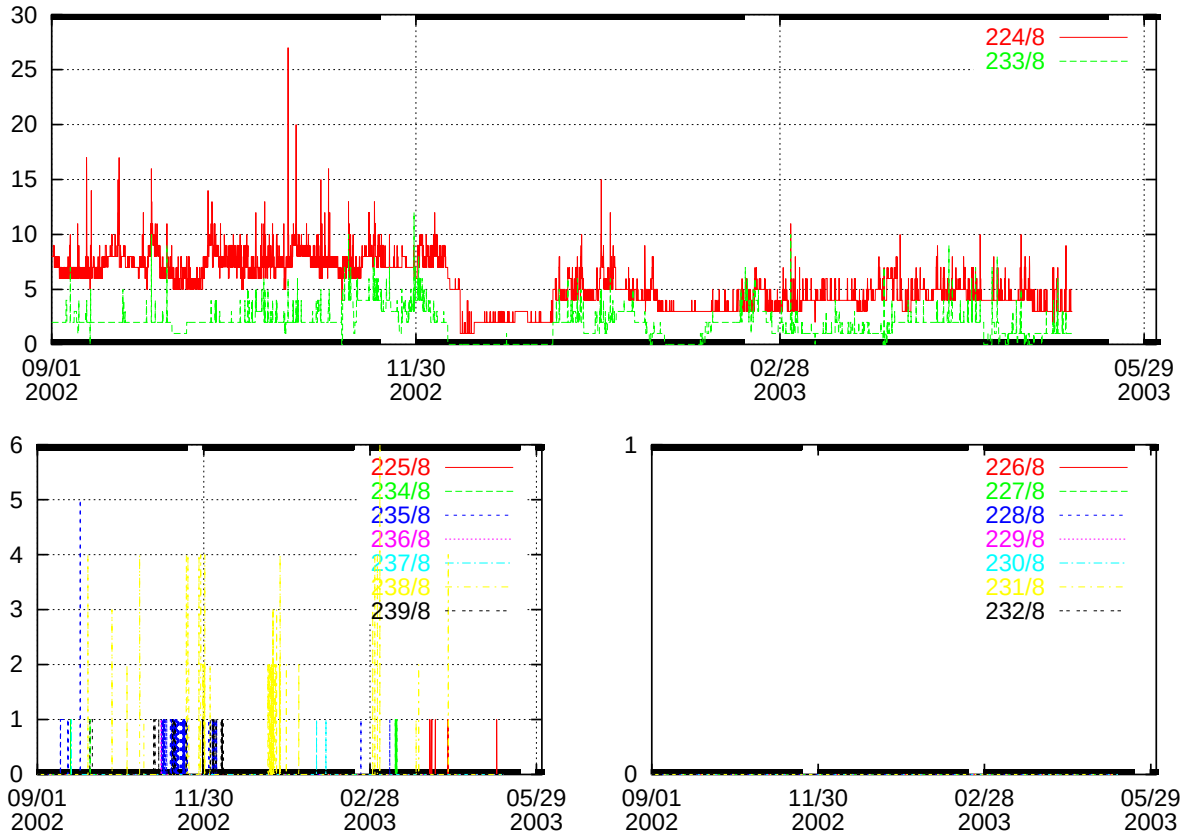

Fig. 4. Group number statistics divied by group number

We are still devoting to enhance our monitoring system's capability of flexible measurement. However, the major part of our future work is to deploy much more monitoring system throughout the CERNET to facilitate our research about multicast traffic characteristics.

VI. Conclusion

In this paper, we presented analysis of over a half-year multicast traffic measurement. Using our passive monitoring system, we collected data over seven-month period with the equipment of only a 1 TB disk array. It is cheaper and simpler to deploy our monitoring system than requiring other additional special hardware. Then a new methodology was introduced for analysis of multicast traffic taken on CERNET. We depicted the characteristics of multicast traffic with the facility of passive monitoring. What we got are different from other researchers', including packet size distribution, group number distribution and group space utilization. It could be reasonable that the deployment of multicast on the CERNET is growing up. Some extraordinary and interesting phenomena will catch our eyes.

After extracting adequate information from passive monitoring traffic, we'll ask for help from router and other multicast-specific protocols. Future multicast research on CERNET will include deploying distributed agents [10] for facilitating our research, collecting and processing multicast-relative data coming from routers and other methods, collaborating all kinds of multicast traffic data to explore the characteristics caused by special transfering mechanism.

Acknowledgement

The work presented in this paper is currently funded by National Natural Science Foundation of China under contract number 60103005 and by National High Technology Research and Development Program of China under contract number 2001AA142080.

References

[1] R. Beverly, G. Miller, K. Thompson, Multicast Performance Measurement on a High-Performance IP Backbone, Computer Communications (ComCom), Vol. 24, pp. 461-472, March 2001.

[2] K. Claffy, H.-W. Braun, and G. Polyzos, A parameterizable methodology for Internet traffic flow profiling, IEEE Journal on Selected Areas in Communications 13, 1481-1494.

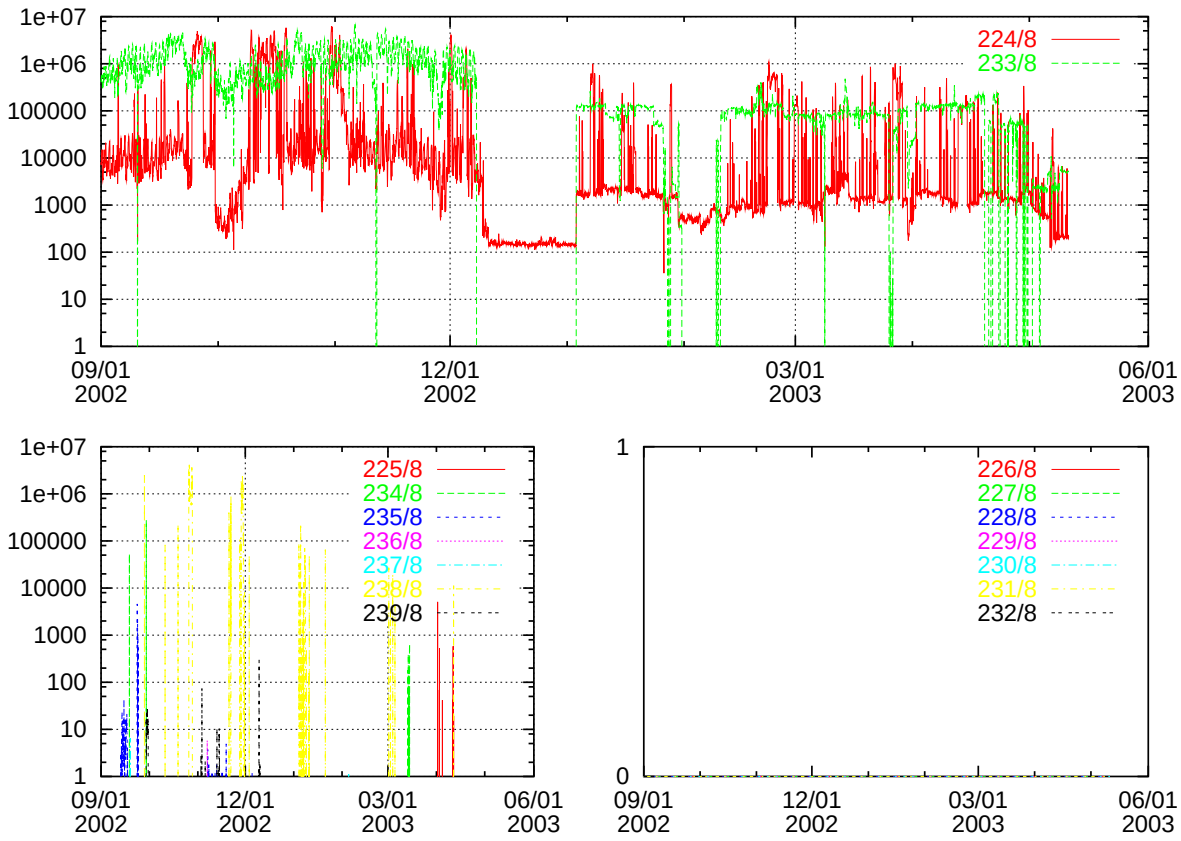

Fig. 5. Traffic volume statistics divied by group address

[3] K. Sarac and K. Almeroth, Supporting multicast deployment efforts: A survey of tools for multicast monitoring, Journal of High Speed Networking–Special Issue on Management of Multimedia Networking, March 2001.

[4] T. He, H. Zhang, X. Li and Z. C. Li, A Methodology for Analyzing Backbone Network Traffic at Stream-Level, ICCT 2003, Beijing China, April 2003.

[5] C.X. Bao, Designment and Debugging of Multicast in NSFC-NET, APAN 2002, Shanghai China, Aug 2002.

[6] R. Beverly, K. Claffy, Wide-Area IP Multicast Traffic Characteristics, IEEE Network, Jan/Feb 2003.

[7] Fraleigh, C. and Moon, S. and Lyles, B. and Cotton, C. and Khan, M. and Moll, D. and Rockell, R. and Seely, T. and Diot, C, Packet-Level Traffic Measurements from the Sprint IP Backbone, IEEE Network, 2003.

[8] K. Almeroth, Managing IP multicast traffic: A first look at the issues, tools, and challenges. IP Multicast Initiative White Paper, August 1999.

[9] B. Mah, Measurements and Observations of IP Multicast Traffic, Technical Report UCB/CSD-94-858, University of California, Berkeley, CA, December 1994.

[10] H. Zhang and X. Li, A Scalable High Performance Network Monitoring Agent for CERNET, submitted to PDCAT03.

[11] T. McGregor, H. W. Braun, J. Brown, The NLANR Network Analysis Infrastructure, IEEE Communications, Vol. 38, No. 5, May, 2000.

[12] C. Fraleigh, C. Diot, B. Lyles, S. Moon, P. Owezarski, D. Papagiannaki, F. Tobagi, Design and Deployment of a Passive Monitoring Infrastructure, Passive and Active Measurement Workshop (PAM) 2001, Amsterdam, The Netherlands, April, 2001.

[13] Z. C. Li, H. Zhang, Y You, T. He, Linuxflow: A High Speed Backbone Measurement Facility, Passive and Active Measurement Workshop (PAM) 2003, La Jolla, California, USA, April, 2003.

[14] W. Stallings, SNMP, SNMPv2, and SNMPv3, and RMON 1 and 2, Addison Wesley, 3rd edition, 1999.

[15] S. McCanne, V. Jacobson, The BSD Packet Filter: A New Architecture for User-Level Packet Capture, In Proceedings of the Winter 1993 USENIX Conference, pp. 259-269. USENIX Association, January 1993.

[16] NetFlow services and applications, http://www.cisco.com/.

[17] CoralReef web page, http://www.caida.org/tools/measurement/coralreef/.

[18] Tcpdump web page, http://ee.lbl.gov/.

[19] Beacon web page, http://dast.nlanr.net/Projects/Beacon/.

[20] CERNET, http://www.edu.cn/.

[21] CAIDA tool collections, http://www.caida.org/tools/taxonomy/multicast.xml