Receiver-initiated Group Membership Protocol (RGMP): a New Group Management Protocol for IP Multicasting

Wanjiun Liao and De-Nian Yang
Department of Electrical Engineering
National Taiwan University
Taipei, Taiwan
Email: wjliao@cc.ee.ntu.edu.tw

Abstract Internet multicast is an important networked service for many existing and emerging applications. The dominant mechanism for group management of IP multicasting is the Internet Group Management Protocol (IGMP). IGMP is based on a query/reply model and refreshes group membership periodically. IGMP has been evolving through three versions. IGMP v1 develops the basis of the query/reply group management model and suppression mechanism, IGMP v2 reduces leave latency that IGMP v1 suffers from, and IGMP v3 adds the capability of source filtering but removes the suppression mechanism of IGMP v1/v2. In this paper, a new group management protocol called Received-initiated Group Membership Protocol (RGMP) is proposed. Both source filtering and membership report suppression are supported. An RGMP host actively refreshes group membership in the neighboring multicast routers. No querier, and hence no query messages and timers, is required for periodically probing the presence of known groups. An individual host maintains a refresh timer per group. The refresh timer is reset once the suppression rule holds for a received report message, where the report may be a join, leave, state change, or a periodical refresh message. The receiver-initiated, self-synchronized refresh timer makes the RGMP suppression mechanism superior to that of IGMP v1/v2, which applies only for periodical refresh report message. As a result, the protocol overhead compared to IGMP v3 is significantly reduced, over a wide variety of service scenarios. In addition to reduced protocol overhead, RGMP is robust, scalable and adaptive to serve as a group management protocol.

Keywords: IP multicast, IGMP, group management protocol, RGMP

1. Introduction

Being an important subject both in research and development [1][2], IP multicast is key in many existing and emerging Internet applications, including bulk data dissemination, resource discovery, replicated database update, real-time video conferencing and media-on-demand, just to name a few. A multicast data packet contains a class D group address[1] in the destination address field of the IP header, and is delivered to destination group members with the same "best-effort" delivery as unicast IP data transmission. It is not necessary for a host that sends data to a multicast group to be a group member. Individual hosts are free to join or leave a multicast group at any time. No restriction is placed on the physical location and the number of group members, and on the number of groups a host can participate.

The provision of multicast data delivery on the Internet is realized through two mechanisms: local group management and global multicast routing. The local group management mechanism enables multicast routers to learn the presence of group members on their directly attached networks, while the global routing mechanism enables multicast routers to exchange information for the determination of multicast delivery trees to forward multicast datagrams across the Internet. Important multicast routing protocols include DVMRP [3], MOSPF [4], CBT [5], and PIM [6]. While a wide variety of global routing protocols is in place, there is only one dominant protocol for local group management, namely, the Internet Group Management Protocol (IGMP). IGMP adopts a query/reply model for communications between group hosts and multicast routers, and refreshes group membership periodically

[1] A class D multicast address identifies a multicast group.

to cope with the dynamic join and departure of group members and to increase robustness to cope with the best-effort delivery nature of IP-based networks.

The key objective of group management protocols is to provide the best support of dynamic group membership for a wide range of Internet applications and service scenarios. With the provision of dynamic group membership, an individual host is free to join or leave a group dynamically without affecting others in the group. The presence or absence of members in a group should be learned by the immediately neighboring multicast routers within an acceptable time period, thereby reducing the possible join or leave latency, respectively, for group participation. Join latency here refers to the time elapsed between a host joining a group and the host starting to receive data packets for the group; leave latency means the time elapsed between the last member leaving a group and the neighboring multicast router detecting no more member in the group. Long join latency introduces long waiting time for group participation, particularly for the first member joining the group; long leave latency wastes system resources with the forwarding of undesired datagrams through the networks.

IGMP has been evolving through three versions, namely, IGMP version 1(v1) [7], IGMP version 2(v2) [8], and IGMP version 3(v3) [9]. All three versions of IGMP follow the same query/reply model, and each, incrementally, achieves a partial scope of the generic group management protocol for IP systems. IGMP v1 develops the basis of the query/reply model for the management of dynamic group membership, and reduces join latency by sending an unsolicited membership report from a host to the neighboring multicast routers upon a new member joining a group. IGMP v2, based on IGMP v1, reduces leave latency by incorporating two types of query (a general query which is the same as the one used in IGMP v1, and a group-specific query) and two query intervals (a longer query interval, 125 sec, for general query and a shorter interval, 10 sec, for group-specific query). IGMP v3, based on IGMP v2, adds the capability of source filtering but removes the suppression mechanism of IGMP v1/v2. Source filtering is the ability for an individual host to specify the reception of packets sent to a multicast group *only* from a list of source addresses (include mode) or to *explicitly* identify a list of the sources the host does not want to receive from a multicast group (exclude mode). Both IGMP v1 and v2 can be treated as a special case of IGMP v3 with all group members requesting for a wildcard source filter (namely, placing no source filter on the group and receiving data from all sources). Protocol overhead of IGMP is due to the exchange of control messages between hosts and multicast routers. In terms of the number of control messages, for both IGMP v1 and v2, protocol overhead is proportional to the number of groups; protocol overhead for IGMP v3 is proportional to the number of hosts on a network participating in any group, due to no report suppression. Although IGMP v3 is backward compatible to v1 and v2 and it supports source filtering, IGMP v3 does not automatically adapt to applications or services scenarios favorable to IGMP v1/v2. These phenomena can be observed from the simulation results in the performance evaluation section.

In this paper, a new group management protocol for IP systems called Received-initiated Group Membership Protocol (RGMP) is proposed. Both source filtering and suppression are supported. An RGMP host actively takes responsibility to refresh group membership in the neighboring multicast routers. No querier, and hence no query messages and various timers, is required for periodically probing the presence of known groups. An individual host maintains a refresh timer per group. The refresh timer is reset once the suppression rule holds for a received report message, where the report may be a join, leave, state change, or a refresh message. The receiver-initiated, self-synchronized refresh timer makes the RGMP suppression mechanism superior to that of

IGMP v1/v2, which applies only for periodical refresh report message. RGMP is robust, scalable and adaptive. Protocol overheard of RGMP is much less than that of IGMP v3, irrespective of group size, group number, number of hosts participating in any group, percentage of hosts on a network having source filtering and group change rate of a host.

The rest of the paper is organized as follows. Section 2 describes RGMP in details. Section 3 shows the simulation results. Finally, the conclusions are made in Section 4.

2. Receiver-initiated Group Membership Protocol (RGMP)

The Receiver-initiated Group Membership Protocol (RGMP) is a protocol for IP systems to report group membership to the neighboring multicast routers. Instead of following the query/reply model as in IGMP, group members actively refresh their membership on expiry of refresh timers. In this section, protocol characteristics are described first, followed by the details of the mechanism.

2.1 Characteristics

The characteristics of RGMP are summarized as follows. Improvements in performance compared to IGMP will be validated by simulation in a later section.

1. Robustness: RGMP adopts soft state to improve robustness, refreshing group membership and source filtering periodically to cope with the best-effort delivery nature of IP systems.
2. Source filter with suppression: RGMP supports both source filtering and suppression. For both IGMP v1 and v2, no source filter is supported; for IGMP v3, no suppression is provided. The source filter mechanism allows host members to customize their preferences for the reception of data from sources; while suppression avoids report message implosion from host members.
3. Scalability: With the support of the suppression mechanism, report message implosion is avoided. Thus, protocol overhead does not increase as the number of groups increases (as in IGMP v1/v2), or as the number of hosts increases (as in IGMP v3).
4. Receiver-initiated refresh timer: The receiver-initiated approach eliminates the querier, query messages and timers employed by IGMP, and hence simplifies the mechanism. With IGMP, a multicast router uses various timers to learn the presence of a group, and a host maintains a random timer to suppress report transmission only if a suppression mechanism is employed. An RGMP refresh timer associated with a host, however, serves the combined purpose of membership refresh and suppression. Moreover, an IGMP join report does not provide suppression due to the query/reply model. The RGMP suppression mechanism, on the other hand, is triggered by a member report from new join, state change, or periodical refresh, resulting in much larger reduction in protocol overhead messages.
5. Self-synchronized refresh timer: With IGMP, control messages are usually exchanged periodically, activated by a single query sent by a router. This causes a periodical burst of overhead messages, and thus increases the possibility of packet collisions, especially as the number[2] of groups increases. In contrast, an RGMP refresh timer is reset per group (or per group-source) on a self-synchronization basis. A join, state change, or the first periodical refresh report in a group suppresses the transmission and resets (or synchronizes) the refresh timer of each member host in the group. As a result, packet transmissions are distributed smoothly over time.
6. Adaptivity: No source filtering is supported by IGMP v1/v2 so that IGMP v1/v2 can be treated as a special case of IGMP v3 with wildcard source filter for all hosts. IGMP v1/v2, however, allows report suppression. The protocol overhead increases mainly as the number of groups on a LAN increases, irrespective of the number of members in a group.

[2] or the number of hosts in IGMP v3.

Thus, the mechanism is best suited for those application scenarios in which hosts are distributed in a small number of groups and each group has a large group size. IGMP v3 has no suppression. The protocol overhead increases as the number of hosts on a local network increases. Thus, the mechanism is best suited for application scenarios in which each host participates in many groups and has source lists associated with each participating group. IGMP v3 performs worse when applied to application scenarios favorable to IGMP v1/v2. RGMP, on the other hand, performs well for both scenarios, adapts well to different scenarios as appropriate, and incurs relatively low protocol overhead while allowing hosts on a LAN with mixed application scenarios.

2.2 Mechanism: an overview

The Receiver-initiated Group Membership Protocol (RGMP) is a group management protocol for IP multicasting. No querier probes the presence of groups periodically as in IGMP. Querier, general query timer, and group specific-query timer are eliminated from the multicast routers. State information stored in a multicast router is on a per group basis. An individual host maintains a refresh timer per group. The report message is the only message used by a group member to communicate with multicast routers. According to the usage, report messages can be classified into three types: join/leave message, state-change message, and periodical refresh message. All report messages are unsolicited and may be sent by hosts upon new join, last departure, state change (e.g., change filter mode, add source IP addresses, or delete source IP addresses, etc), or expiry of the refresh timer. State information maintained by a host includes group id (i.e., multicast address), filter mode, source list, refresh timer, and suppression flag. The group id is a group address, and the filter mode can be either include or exclude. The source list consists of a list of source addresses (id). Contrary to IGMP v3, each element in an RGMP source list is associated with a source flag, used for suppression. The interpretation of a source flag depends on the type of filter mode: (1) When the filter mode is "include," if a source element is associated with the "on" flag, it must have been refreshed by a previous report message within a refresh interval; otherwise, the flag is off. (2) When the filter mode is "exclude," if a source element associated with the "on" flag, the element is waiting to be refreshed.

For example, Figure 1 shows a membership list maintained by a host participating in groups 1, 2 and 3 on a local network. In group 1, source filter mode is **exclude{}**, meaning that the host accepts any coming sources from this interface to group 1. Filter mode of **exclude{A}** usually has a larger accepted source lists than that of **include{A}**, where A is a legal source list x, y, z. In group 3, source elements *a* and *c* have their source flags on while source *b* has the flag off, indicating that both *a* and *c* have been suppressed by previous report messages, but *b* has not. Once the flag of element *b* turns to "on" from being refreshed by a report message (i.e., all elements in the source list have their flags on) before the expiry of the refresh timer t1, the suppression flag of group 3 will be set to on.

The refresh timer controls the time interval to refresh group membership in multicast routers. On expiry of a refresh timer, a refresh report is sent. The suppression flag records the suppression status for each participated group. If the suppression flag is currently "on," it indicates that the corresponding group has been suppressed by previous received reports; otherwise, the suppression flag is "off." The value of a refresh timer is determined by the status of the suppression flag. If the suppression flag is ON, indicating that the group has been suppressed by previous reports, a random delay is set from the range of [T2, T3); otherwise, the suppression flag is off and a random delay is selected from the range of [T1, T2), where $0 < T1 < T2 < T3$[3]. For example, in Figure 1, t1 of group 1 or group 3 is set

[3] To have a fair comparison with IGMP v3, T1 is set to 115, T2 is set to 125 and T3 is set to 135 in our simulation.

to a value in (115, 125), and t2 of group 2 is in (125, 135). The purpose of using two different timer values for two types of suppression flag is to ensure that groups with flag "on" will not send a report to multicast routers because their timers always expire after those with flag "off," and will be suppressed.

Group ID	Filter mode	Suppression flag	Refresh timer	Source list
1	Exclude	OFF	t1	{}
2	Exclude	ON	t2	{(a, ON)}
3	Include	ON	t2	{(a, ON), (b, OFF), (c, ON)}

Figure 1. An example membership list of a host

Operation overview

(1) Host

A host transmits a report to multicast routers upon joining or leaving group(s), changing states, or when the refresh timers expire. Sending a report message comes with three state settings per group in the host: the suppression flag is set to off, the refresh timer is reset accordingly, and all source flags are reset (i.e., all source flags are off at the end). The exact operation of source flag reset depends on the type of filter mode. If the filter mode is "include," all source flags currently on are set to off; otherwise, the filter mode is exclude, and hence all source elements with source flags currently on are removed. A single report message consisting of all *involved* groups is sent. For reports sent upon joining a group, "involved" refers to "interested," for state changes, "involved" refers to "modified," and for departure, "involved" refers to "unsuppressed" (i.e., those with suppression flags off).

A host may receive a report message before the expiry of the refresh timer. If the received report is a new join or refresh message, the suppression rule is applied. Otherwise, when the received report is a leave (a report message with **include{}** for a group) or state change message in the same group, if the suppression flag is off, the refresh timer is reset to a small value (say, 0-1 sec); otherwise, the refresh timer is reset to a slightly longer random delay (say, 1-2 sec). A refresh message, again, is sent as the refresh timer expires.

(2) Multicast router

A multicast router passively handles report messages sent by hosts. The operations performed by routers are similar to what IGMP v3 routers do except all queriers and the related mechanisms are removed to simplify the design of the routers.

Suppression rule

Assume that host H1 joins group G1 with filter mode Mode-A and source list Source-A. When a report message sent by other host in the same group G1 is received, with filter mode Mode-B and source list Source-B,

1. when Mode-A = include,
   (a) when Mode-B = include, $\forall$ x in {Source-A $\cap$ Source-B}, the source flag of x in Source-A is set to on.
   (b) when Mode-B = exclude,$\forall$ x in Source-A, if x $\notin$ Source-B, the source flag of x is set to on.

   For both (a) and (b), if all the elements in Source-A have their flags on, the suppression flag of G1 is set to on, the refresh timer is reset accordingly, and the source flags of all elements in Source-A are set to off.

2. when Mode-A = exclude,
   (a) when Mode-B = include, let Source-D = $\varnothing$. $\forall$ x with source flag on in Source-A, Source-D = Source-D $\cup${x}. If Source-D $\neq \varnothing$, Source-A = Source-A – {Source-B $\cap$ Source-D}. If no element in Source-A has its source flag on while Source-D $\neq \varnothing$, suppression flag of G1 is set to on, and the refresh timer is reset accordingly.
   (b) when Mode-B = exclude,
      (1) If $\forall$ x in Source-A, flags are off, Source-D = {Source-A $\cup$ Source-B} – Source-A. If Source-D $\neq \varnothing$, $\forall$ x in Source-D, set the source flag of x to be on, Source-A = Source-A $\cup$ Source-D; otherwise, the suppression flag of G1 is set to on, and the refresh timer is reset accordingly.
      (2) If $\exists$ x in Source-A with source flag on, then $\forall$

x with source flag on in Source-A and x ∉ Source-B, Source-A = Source-A – {x}. If there is no element in Source-A with its source flag on, the suppression flag of G1 is set to on, and the refresh timer is reset accordingly.

3. Performance Evaluation

To supplement the comparisons of IGMP and RGMP, we examine the protocol overhead in terms of bandwidth requirements[4] on a local network. The measured results vary with the following metrics: (1) the number of hosts participating in any groups on a LAN (host number), (2) the number of hosts in a group (group size), (3) the number of different groups on a LAN (group number), (4) group join/departure rate of a host (join/leave rate), (5) the percentage of hosts with source filter on a LAN, and (6) protocol overhead distribution over the timeline. The simulation was performed over three hours, and under three assumptions: (1) no packet loss, and hence no retransmission, for message exchange, (2) no communications delay between hosts and hosts/routers on the reception of query/report messages, and (3) the following two processes are assumed to be Poisson (with λ=1/5400 per sec): the join/departure process of a group, and the join/departure process of a source within a group.

Since both IGMP v1 and v2 do not support source filtering, the major comparisons will be made between RGMP and IGMP v3. To have a fair comparison with IGMP v3, RGMP will use the same message format as that of IGMP v3 in the simulation. Performance comparisons of each metric are depicted by two figures: one without source filtering and the other with source filtering. Each host may select a list of sources from 15 different source IP addresses per group. For those without source filtering, all group members have the same filter mode and source list, and for those with source filtering, an individual group member is free to select a filter mode and a source list from 15 different source IP addresses. For example, host 1 may have source filter of **include {1,2,4}**, and host 2 may have source filter of **exclude {1}**. Protocol overhead of IGMP v1 and v2 are included in figures without source filtering, mainly for reference.

3.1 Host number

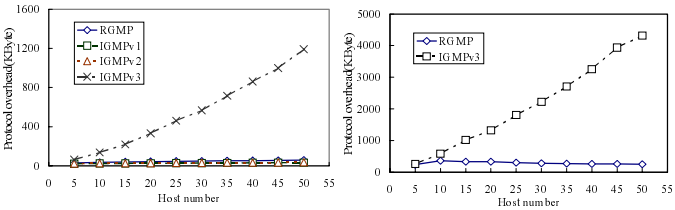

(a) Without source filtering (b) With source filtering

Figure2. Host number vs. protocol overhead

Results are measured over an individual host participating, at most, in 30 group numbers, with host number on a LAN that participates in any group varying from 1 to 50. Figure 2 (a) shows the results without source filtering, and Figure 2 (b), with source filtering. Both figures illustrate that the overhead of IGMP v3 increases as host number increases, while that of RGMP is almost invariant to the increase of host number on a LAN. In addition, the protocol overhead of RGMP is very close to that of IGMP v1/v2 as shown in Figure 2(a), although an IGMP v1/v2 query/report message has a fixed-size packet of 8 bytes, and an RGMP report message using the same format as that of IGMP v3 is relatively large.

3.2 Group size

Results are measured over 30 hosts with each participating at most in 30 groups, with group size on a LAN varying from 0 to 27. Figure 3 (a) shows the results without source filtering, and Figure 3 (b), with source filtering. Considering the one without source filtering during a refresh period, protocol overhead of IGMP v3 for both cases is approximately equal to 8 bytes * group size (or host number) * group numbers, while that of RGMP is approximately equal to 8 bytes * group

[4] In this section, protocol overhead was measured in terms of bandwidth requirements, rather than the number of control packets.

numbers.

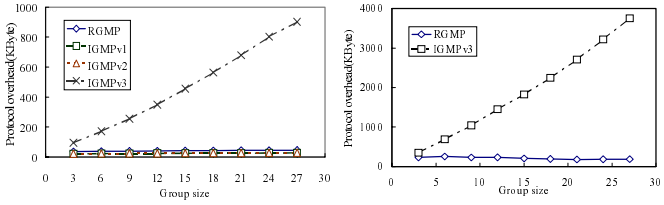

(a) Without source filtering (b) With source filtering

Figure 3. Group size vs. protocol overhead

3.3 Group number

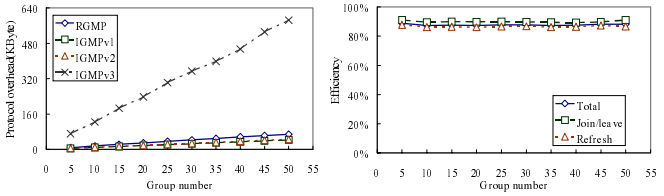

(a) Protocol overhead, no SF (b) Efficiency, no SF

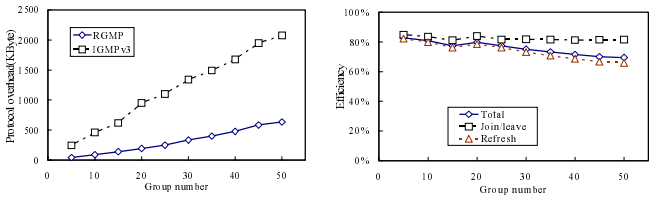

(c) Protocol overhead, SF (d) Efficiency, SF

Figure 4. Group number vs. protocol overhead

Results are measured over 20 hosts participating in any group, with group number on a LAN varying from 1 to 50. Figure 4 (a) and 4 (b) shows the results without source filtering, and Figure 4 (c) and 4 (d), with source filtering. Figure 4 (b) and Figure 4 (d) show the efficiency of RGMP compared to IGMP v3 without and with source filtering (SF), respectively, where efficiency is defined as follows:

Efficiency = 1-(RGMP overhead / IGMP v3 overhead)

The higher the efficiency, the better the performance improvement of RGMP over that of IGMP v3 for the specific comparison item. For example, x% of efficiency means that bandwidth requirement of RGMP is only (100-x)% of that IGMP v3.

Figure 4 shows that the overhead of both IGMP v3 and RGMP increases as group number increases. However, the overhead of IGMP v3 increases more rapidly than that of RGMP as group number increases. Efficiency of RGMP to IGMP v3 stays fairly flat in both cases. This is because the host number is fixed, as group number increases, the number of groups an individual host participates in increases, and hence the size of a single report message of IGMP v3 increases accordingly. Efficiency of RGMP over IGMP v3 due to report suppression is thus significant. Group size, on the average, is about 10 (there are 20 hosts in total participating in any group, and the probability that a host participates in any group is 0.5). Considering the one without source filtering during a refresh period, protocol overhead of IGMP v3 is approximately equal to 10 * 8 bytes * group numbers, and that of RGMP is approximately equal to 8 bytes * group numbers, giving an overhead ratio of 10:1.

3.4 Join/departure rate of a group

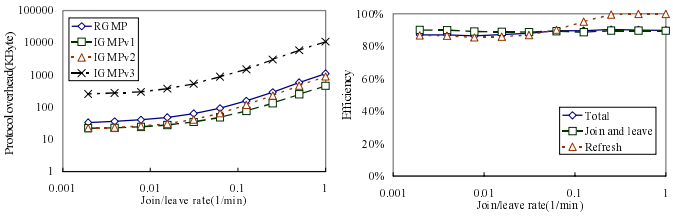

(a) Protocol overhead, no SF (b) Efficiency, no SF

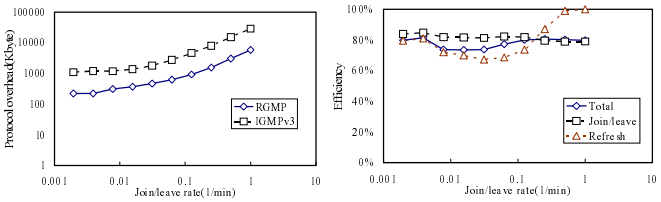

(c) Protocol overhead, SF (d) Efficiency , SF

Figure 5. Join/leave rate vs. protocol overhead

Results are measured over 20 hosts with each at most participating in 30 groups, with average join/departure rate of a group varying from 0.001 to 1 per minute. Figure 5 (a) and 5 (b) show the results without source filtering, and Figure 5 (c) and 5 (d), with source filtering.

It can be observed that as λ is very high, there is almost no periodical refresh message sent because join/leave messages reset refresh timers all along, and thus, efficiency of periodical refresh reaches 100%. As a result, protocol overhead of periodical refresh messages

dominates the overall performance when λ is small, and that of join/leave messages dominates the performance as λ increases.

3.5 Percentage of hosts with source filtering on a LAN

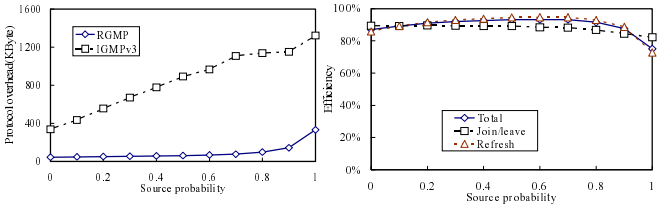

(a) Protocol overhead (b) Efficiency

Figure 6. Source filtering Percentage vs. protocol overhead

Results are measured over 20 hosts with each at most participating in 30 groups and each at most having 15 different sources per group, with the percentage of hosts with source filtering on a LAN varying from 0% to 100%. Surprisingly, when 80% of hosts with source filtering, efficiency of RGMP over IGMP v3 is above 90%. Even when all hosts on a LAN have source filtering (namely, 100% of hosts with source filtering), efficiency of RGMP over IGMP v3 is still over 70%.

3.6 Protocol overhead distribution over time

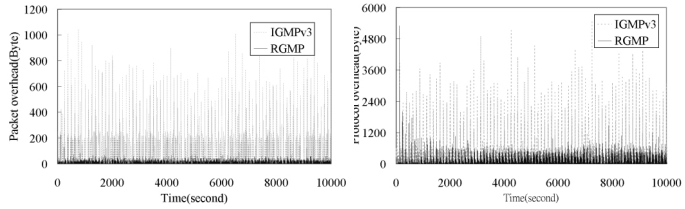

(a) Without source filtering (b) With source filtering

Figure 7. Protocol overhead distribution over time

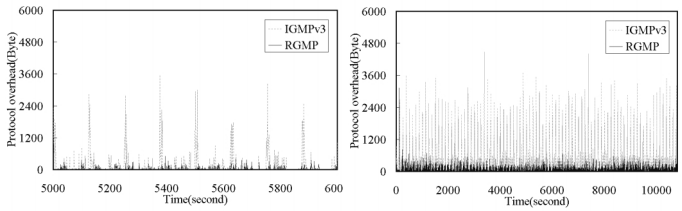

(a) Without source filtering (b) With source filtering

Figure 8. Protocol overhead distribution over an interval

Figure 7 shows protocol overhead distribution over time, and Figure 8 shows a distribution over an interval. Both figures depict that RGMP consistently has less total protocol overhead over the simulation time, and the report messages are distributed smoothly over time rather than as periodical bursts as in IGMP, thanks to the receiver-initiated and self-synchronized refresh timer mechanism.

4. Conclusions and Future Work

In this paper, we have proposed a new group management protocol, called RGMP, for IP multicasting. The protocol characteristics were described, and the mechanism of the protocol was presented. Comparisons to IGMP were made to highlight the advantages of the proposed approach in terms of performance, adaptivity, scalability, and capability to serve as a group management protocol of IP multicasting for a wide variety of Internet services.

Acknowledgement

This work is supported by the National Science Council, Taiwan, under grant number NSC 88-2219-E-002-007.

References

[1] M. H. Ammar, G. Polyzos, and S. Tripathi. “Special issue on networked support for multipoint communications,” IEEE JSAC, vol. 15, April 1997.

[2] J. C. Pasquale, G. C. Polyzos, and Xylomenos, “The Multicasting Problem,” ACM Multimedia Systems, vol. 6, o. 1, 1998, pp. 43-59.

[3] S. Deering, C. Partrige, and D. Waitzman, “Distance Vector Multicast Routing Protocol,” IETF RFC 1075.

[4] J. Moy, “Multicast Routing Extensions for OSPF,” Communication of the ACM, vol. 37, no. 8, pp 61-66, Aug. 1994.

[5] A. Ballaradie, J. Crowcroft, and P. Francis, “Core Based Tree (CBT) – An Architecture for Scalable Inter-Domain Routing Protocol,” ACM SIGCOM ’93, Oct. 1993, pp. 85-95.

[6] S. Deering, D. Estrin, D. Fairnacci, V. Jacobson, C. Liu, and L. Wei, “An Architecture for Wide-Area Multicast-Routing,” ACM SIGCOMM ’94, pp. 126-135.

[7] S. Deering. “Host Extensions for IP Multicasting,” *IETF RFC 1112*, Aug. 1989.

[8] W. Fenner. “Internet Group Management Protocol, Version 2,” *IETF RFC 2236*, Nov. 1997.

[9] B. Cain, S. Deering, and A. Thyagarajan. “Internet Group Management Protocol, Version3,” *IETF Internet draft*, Feb. 1999. (Work in progress)