

Receiver-Initiated Group Membership Protocol (RGMP): A New Group Management Protocol for IP Multicasting

Wanjiun Liao and De-Nian Yang

Abstract—This paper proposes a new group management protocol called Received-initiated Group Membership Protocol (RGMP) for IP multicasting. The dominant group management protocol on the Internet to date is the Internet Group Management Protocol (IGMP). Unlike IGMP based on a query/reply model, an RGMP host actively takes responsibility to refresh group membership on the neighboring multicast routers. Each RGMP host maintains a “refresh” timer per group. The refresh timer is reset once the suppression rule holds true for a received report message, where the report may be a join, departure, state change, or refresh message. The RGMP refresh timer is adjusted in a way to be adaptive and self-synchronized. This receiver-initiated, self-synchronized approach makes the RGMP suppression mechanism superior to that of IGMP v1/v2, because the latter can be applied only to periodical refresh messages. As a result, RGMP protocol overhead is significantly reduced over a wide variety of service scenarios compared to IGMP v3. In addition to the reduced protocol overhead, RGMP is robust, scalable and adaptive to serve as a group management protocol.

Index Terms—Group management protocol, IGMP, IP multicast, RGMP.

I. INTRODUCTION

BEING an important subject both in research and development [1], [2], IP multicast is key to 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. In IP multicast, a host is not necessarily a group member in order to send data to the group. Individual hosts are free to join or leave a multicast group at any time. No restriction is placed on the physical location of a host, the size of each group, and the number of groups in which each host can participate.

IP multicast has been being considered a key technology in multimedia content distribution over the Internet since its proposal by S. Deering [3] in 1988. However, to date, this technology has only been of interest to the research community, and

as not been widely deployed over the Internet. The major challenges to its deployment include the scalability problem (i.e., too many forwarding states at routers), and the security problem. The security problem is further decomposed into source filtering (i.e., source authentication) and access right management (i.e., receiver authentication). The scalability and receiver authentication problems of IP multicasting have been active research topics (e.g., [6]–[12] for scalability, and [13], [14] for access right). However, there are not many discussions on source filtering for IP multicasting. In this paper, we focus on source filtering from the perspective of local routers.

Multicast data delivery on the Internet can be collaboratively provided with 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; the global routing mechanism enables multicast routers to exchange information in order to determine multicast delivery trees through which multicast datagrams are forwarded across the Internet. Important multicast routing protocols to date include DVMRP [9], MOSPF [10], CBT [11], and PIM [12]. While a wide variety of global routing protocols have been in place, the Internet Group Management Protocol (IGMP) [13]–[15] is the only dominant protocol for local group management. Based on a query/reply model, IGMP refreshes group membership on multicast routers periodically. This query/reply mechanism allows the system to cope with the dynamic join and departure of group members, and to increase robustness when dealing with the best-effort delivery nature of IP-based networks.

Group management protocols aim at providing the best support of dynamic group membership for a wide range of Internet applications and service scenarios. With the provision of dynamic group membership, each host is free to join or leave a group dynamically, without affecting others in the same group. The challenge for local multicast routers is to recognize the presence or absence of members in a group within an acceptable time period, so as to reduce the possible join or leave latency 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 from the group. Leave latency means the time elapsed between the last member leaving a group and the neighboring multicast router detecting no more members in the group. Longer join latency introduces longer waiting time for group participation, particularly for the one

Manuscript received September 24, 2002; revised February 23, 2004. The preliminary version of this paper has been published in the Proc. IEEE International Conference on Network Protocols (ICNP)'99. This work was supported by the National Science Council, Taiwan, under Grant Number NSC 92-2213-E-002-064.

W. Liao is with the Department of Electrical Engineering, National Taiwan University, Taipei, Taiwan and also with the Graduate Institute of Communication Engineering, National Taiwan University (e-mail: wjliao@cc.ee.ntu.edu.tw).

D.-N. Yang is with the Department of Electrical Engineering, National Taiwan University, Taipei, Taiwan.

Digital Object Identifier 10.1109/TBC.2004.834015

who is the first member to join the group on a LAN. Longer leave latency wastes more system resources on the forwarding of undesired datagrams through the network.

IGMP has been evolving through three versions, namely, IGMP version 1(v1) [13], IGMP version 2(v2) [14], and IGMP version 3(v3) [15]. All the three versions of IGMP follow the same query/reply model, and each, incrementally, achieves the partial scope of being a generic group management protocol for IP multicasting. IGMP v1 develops the basis of the query/reply model for the management of dynamic group membership. Newly hosts send unsolicited membership reports to the neighboring multicast routers as soon as they join a group, which reduces join latency. IGMP v2 is an improvement over IGMP v1. For the reduction of leave latency, IGMP v2 incorporates two types of queries (i.e., 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 for General Query and a shorter interval for Group-Specific Query). IGMP v3 is derived from IGMP v2, but adds the source filtering capability and removes the suppression mechanism of IGMP v1/v2. Source filtering is the ability that an individual host can specify the reception of packets sent to a multicast group *only* from a list of source addresses or to *explicitly* identify a list of the sources the host does not want to receive from a multicast group. Both IGMP v1 and v2 can be regarded as a special case of IGMP v3, with all the group members requesting for wildcard source filters (namely, placing no source filter on the participating groups and receiving data from all the sources). The protocol overhead of IGMP is caused by the exchange of control messages among hosts and their immediately neighboring multicast routers on a subnet. In terms of the number of control messages exchanged, the protocol overheads caused by both IGMP v1 and v2 are proportional to the number of groups on a LAN. However, the protocol overhead of IGMP v3 is proportional to the number of hosts on a network participating in *any* group due to no report suppression. While backward compatible to v1 and v2 and supporting 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 shown in the performance evaluation section.

This paper proposes a new group management protocol called Receiver-initiated Group Membership Protocol (RGMP) for IP multicasting. Both source filtering and suppression are supported in RGMP. An RGMP host actively takes responsibility to refresh group membership on the neighboring multicast routers. No querier, and hence no query message nor multiple timers, is required to periodically probe the presence of known groups. Each host maintains a refresh timer per group. The refresh timer is reset once the suppression rule holds true for a received report message, where the report may be a join, departure, state change, or a periodical “refresh” message. The RGMP refresh timer is adjusted in a way to be adaptive and self-synchronized. The receiver-initiated, self-synchronized

refresh timer makes the RGMP suppression mechanism superior to that of IGMP v1/v2, because the latter can be applied only to periodical refresh messages. RGMP is robust, scalable and adaptive. As compared to IGMP v3, RGMP has much less protocol overhead, irrespective of group size, group number, the number of hosts participating in any group, the percentage of hosts on a network having source filtering, and the rate hosts change groups.

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

II. RECEIVER-INITIATED GROUP MEMBERSHIP PROTOCOL (RGMP)

This section describes the Receiver-initiated Group Membership Protocol (RGMP) in detail. We first overview IGMPv3. Then, we summarize the protocol characteristics and describe the details of the protocol operation. The performance improvement to IGMP is verified through simulation described in the next section.

A. IGMP Version 3

IGMP v3 improves the previous two versions of IGMP with the support of source filtering. A filter mode can be either an “include” or an “exclude.” With this mechanism, a host may request to receive multicast packets only from some specific sources (when in the include mode), or from all but some specific sources (when in the exclude mode). Thus, it helps multicast routers learn which source lists each group is of interest to receive. The state information on a host or a multicast router mainly includes a group id, a filter mode, and a source list. The state may be changed due to the host changing from one group to another, or changing its source lists or its filter modes. A change from one group to another can be regarded identically as leaving an old group and then joining a new group. Periodically, a host refreshes its membership on receipt of a General Query from the multicast router. Unlike IGMP v1 and v2, IGMP v3 does not support suppression of report transmission, avoiding possible complicated merging operations per (group, source lists) pair. An IGMP v3 host reports multiple groups with a single compound message. To lower leave latency while supporting source filtering, the Group-and-Source-Specific Query is employed in addition to the General Query and the Group-Specific Query used in v2. There are various timers maintained on the routers, including various query timers, group timers and source timers. IGMP v3 defines rules for routers to determine an appropriate filter-mode per group upon the reception of reports. It also specifies source specific forwarding rules based on filter modes.

B. The Characteristics of RGMP

RGMP is a group membership protocol for IP multicasting. Unlike IGMP adopting a query/reply model, RGMP hosts actively refresh their membership on multicast routers. The characteristics of RGMP are summarized as follows.

1. Robustness

Adopting the soft state mechanism to improve robustness, RGMP periodically refreshes both group membership and source filtering on multicast routers to cope with the best-effort delivery nature of IP systems.

2. Source filter with suppression

RGMP supports both source filtering and suppression. The RGMP source filtering mechanism allows members to customize their preferences of data reception from different sources. The suppression mechanism avoids the implosion of report messages from host members. Both IGMP v1 and v2 do not support source filtering; IGMP v3 does not provide suppression.

3. Scalability

RGMP supports the suppression mechanism to avoid the report implosion problem. Thus, RGMP protocol overhead does not increase as the number of groups increases (the problem IGMP v1 and v2 have), or as the number of hosts increases (the problem IGMP v3 has).

4. Receiver-initiated refresh timer

The receiver-initiated approach eliminates the need to have a querier, query messages and timers employed by IGMP. Thus, RGMP simplifies the design of multicast routers. In IGMP, a multicast router uses various timers to learn the presence of groups. Each host maintains a random timer to suppress report transmission if suppression is supported. RGMP associates a host with a refresh timer which serves the combined purpose of membership refreshing and suppression. Meanwhile, an IGMP join report does not allow suppression due to its query/reply model. The RGMP suppression mechanism, on the other hand, can be triggered by any kind of member reports, including new join, state change, or periodical refresh. Thus our approach results in much larger reduction in protocol overhead messages.

5. Self-synchronized refresh timer

IGMP control messages are usually exchanged periodically, activated by a single query sent by a router. This causes a periodical burst of overhead messages, thus increasing the possibility of packet collisions, especially as the number¹ of groups on a LAN increases. In contrast, the RGMP refresh timer is reset on a self-synchronization basis. A join, state change, or the first periodical refresh report for a group may suppress further transmission and reset (or synchronize) the refresh timer of each member in the group. As a result, the packet transmission is distributed more smoothly over time.

6. Adaptivity

Since IGMP v1 and v2 do not support source filtering, they can be treated as a special case of IGMP v3 with wildcard source filters for all hosts. IGMP v1/v2, however, provides report suppression. Their protocol overhead increases mainly as the number of groups on a

LAN increases, irrespective of the number of members in each group. Thus, this mechanism is best suited for the application scenarios in which hosts are distributed to a small number of groups and each group has a large group size. IGMP v3 does not support suppression. The protocol overhead increases as the number of hosts on a local network increases. This mechanism is best suited for the application scenarios in which each host participates in many groups and each host has a source list associated with each participating group. The IGMP v3 performs worse when applied to the application scenarios favorable to IGMP v1/v2. RGMP performs well for both types of scenarios. It adapts well to different scenarios as appropriate, and incurs relatively low protocol overhead while allowing hosts on a LAN with mixed scenarios.

C. The Mechanism: An Overview

RGMP is a group management protocol for IP multicasting. No querier periodically probes the presence of known groups as in IGMP. Thus, the Querier, the General Query timer, and the Group-Specific Query timer used in IGMP all can be eliminated from multicast routers. Each host maintains a refresh timer per group. The report message is the only message type for group members to communicate with multicast routers. According to the usage, reports can be classified into three categories: join/departure, state-change, and periodical refresh messages. All the report messages are unsolicited and may be sent by hosts upon new join, last departure, state change (e.g., change filter modes, add source IP addresses, or delete source IP addresses, etc), or refresh timer expiry.

1) *RGMP States:* The state information maintained by an RGMP host includes a group id (i.e., multicast address), a filter mode, a source list, a refresh timer, and a 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, each of which is associated with a source flag to indicate the suppression status of the corresponding source. A source flag may be an "ON" or an "OFF." The interpretation of a source flag is based on the type of the associated filter mode:

- (1) If the filter mode is an "include," an "ON" source flag indicates that the associated source element has been refreshed by a previous report message within the refreshing interval; otherwise, the flag must be "OFF."
- (2) If the filter mode is an "exclude," an "ON" source flag indicates that the associated source element is waiting to be refreshed.

The refresh timer determines the time interval during which the host's membership is considered valid. On expiry of the refresh timer, a host sends a (refresh) report to refresh its group membership on multicast routers. The suppression flag records the suppression status of each participating group for the host. An ON suppression flag indicates that all the elements in the source list of the corresponding group have been refreshed, and

¹or the number of hosts in IGMP v3.

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)}

Fig. 1. An example membership list of a host.

thus the group membership report can be suppressed. Otherwise, the suppression flag is OFF. The status of a suppression flag in turn determines the value of the associated refresh timer:

- (1) An ON suppression flag indicates that the group status has been suppressed by previous received reports. The refresh timer is set to a value randomly selected from the range of $[T2, T3]$.
- (2) If the suppression flag is OFF, a random delay is selected from the range of $[T1, T2]$.

Where $0 < T1 < T2 < T3$ ². For example, in Fig. 1, t1 of group 1 is set to a value selected from (115, 125), and t2 of both group 2 and 3 is from (125, 135). The purpose of using two different timer values for two types of suppression flags is to ensure that the groups with ON flags do not send reports to multicast routers because their timers always expire after those with OFF flags, and are suppressed.

Fig. 1 shows an example membership list maintained by a host participating in groups 1, 2, and 3 on a local network. In the group 1 entry, the source filter is **exclude**{}, meaning that the host can accept any incoming sources to group 1 from this interface. Filter **exclude**{A} usually has a larger accepted source lists than **include**{A}, where A is a legal source list x, y, z. In the group 3 entry, source elements a and c have their source flags ON, while source b has its flag OFF, indicating that both a and c have been refreshed by previous report messages, but b has not. Once the flag of element b has been turned to ON due to being refreshed by a report message by the expiry of the refresh timer t1 (i.e., all elements in the source list have their flags ON), the suppression flag of group 3 should be set to ON. Otherwise, the suppression flag remains OFF.

2) RGMP Operation:

a) Host:

Send a Report—A host may send a report to its immediately neighboring multicast routers when one of the following situations occurs.

- (1) Join a group: a report is sent when the host first joins a group on a LAN.
- (2) Leave a group: a report is sent only when the suppression flag of the departing host is OFF; otherwise, the host leaves silently.
- (3) Periodical refresh: a report is sent on expiry of the refresh timer.
- (4) State change: a report is sent whenever there is a state change for the host.

²To 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 send a report message, a host should set three states per group: the suppression flag is set to OFF, the refresh timer is reset accordingly, and all the source flags are reset (i.e., all source flags are OFF). The exact operation to reset a source flag depends on the type of the associated filter mode. If the filter mode is “include,” all the source flags currently ON are set to OFF. Otherwise, the filter mode is an exclude, and thus all the source elements with a source flag currently ON are removed. A single report message is sent, which includes all the *involved* groups. For the 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).

Receive a Report—A host may receive a report message by the expiry of the refresh timer. If the received report is a new join or refresh message, the suppression rule is applied (described below). Otherwise, the received report may be a departure (a report message with **include**{ } for a group) message or a state change message for the same group. In either case, 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 once the refresh timer has expired. The state information is reset if one of the following conditions holds true: (1) upon receipt of a report before the refresh timer expires, or (2) the suppression rule holds (i.e., when the suppression flag is set to ON).

b) 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 the queriers and the related mechanisms are removed to simplify the design of the routers.

3) Suppression Rule: Suppose that host H1 joins group G1 with a filter mode of Mode-A and a source list of Source-A. If H1 receives a report message with a filter mode of Mode-B and a source list of Source-B from another host in the same group G1, the G1 state of host H1 is updated as follows:

- 1) Mode-A = include,
 - (a) Mode-B = include,

$$\forall x \text{ in } \{\text{Source-A} \cap \text{Source-B}\}, \text{ the source flag of } x \text{ in Source-A is set to ON.}$$
 - (b) Mode-B = exclude,

$$\forall x \text{ in } (\text{Source-A} - \text{Source-B}), \text{ the source flag of } x \text{ in Source-A is set to ON.}$$

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 the elements in Source-A are reset to OFF.

- 2) Mode-A = exclude,
 - (a) Mode-B = exclude,
 - (1) If all the source elements in Source-A have their source flags OFF, Source-D = Source-B – Source-A. If Source-D $\neq \emptyset$, $\forall x$ in Source-D, the source flag of x is set to ON,

and $\text{Source-A} = \text{Source-A} \cup \text{Source-D}$; otherwise, the suppression flag of G1 is set to ON, and the refresh timer is reset accordingly.

- (2) If there are source elements in Source-A which have source flags ON, $\forall x$ with an ON flag in Source-A but $x \notin \text{Source-B}$, $\text{Source-A} = \text{Source-A} - \{x\}$. If in Source-A there is no element with an ON source flag, the suppression flag of G1 is set to ON, and the refresh timer is reset accordingly.

- (b) Mode-B = include,

Let $\text{Source-D} = \emptyset$. $\forall x$ with an ON source flag in Source-A, $\text{Source-D} = \text{Source-D} \cup \{x\}$. If $\text{Source-D} \neq \emptyset$, $\text{Source-A} = \text{Source-A} - \{\text{Source-B} \cap \text{Source-D}\}$. If no element in Source-A has an ON flag while $\text{Source-D} \neq \emptyset$, the suppression flag of G1 is set to ON, and the refresh timer is reset accordingly.

4) *An Operation Example:* Fig. 2 explains how RGMP suppression works. Fig. 2(a) shows a network topology comprised of five hosts and a multicast router. Fig. 2(b)–(e) are the source filtering states of each host in the same multicast group at the different time. Originally (at $T = 0$), only hosts A, B, and C were in this multicast group. At $T = 29$, host C times out and sends a refresh report message to the multicast router and the other hosts. Because the set of sources host B wants to receive is a subset of the sources that host C is interested in receiving, all the source elements of host B are refreshed after receiving the report message. Therefore, the suppression rule holds and host B is suppressed: the suppression flag set to ON and the refresh timer reset to a random number between 125 and 135 seconds. For host A, the flags of sources a and c in the source list are set to ON because both sources have been refreshed by the host C's report.

At $T = 111$, host D joins this multicast group and multicasts a join report to the multicast router and the other hosts. For host A, since source b is a source that host D wants to receive, the source flag of source b is set to ON after receiving this report message. The suppression condition of host A holds true because all its source flags have turned ON. Therefore, host A sets the suppression flag to ON, sets all the source flags to OFF, and resets the refresh timer accordingly. For host B, source e is refreshed and thus its source flag is also turned ON. For host C, only sources a and c are waiting to be refreshed, and thus both are added with ON flags into the source list. Later, if some hosts which want to receive data from sources a and c send report messages before host C's refresh timer times out, host C will be suppressed.

At $T = 131$, host E joins this multicast group. The join report sent by host E refreshes the two sources a and c of host C, and thus these two sources are removed from the source list of host C. Therefore, the suppression condition of host C holds true: host C sets the suppression flag to ON, sets all source flags to OFF, and resets the refresh timer accordingly.

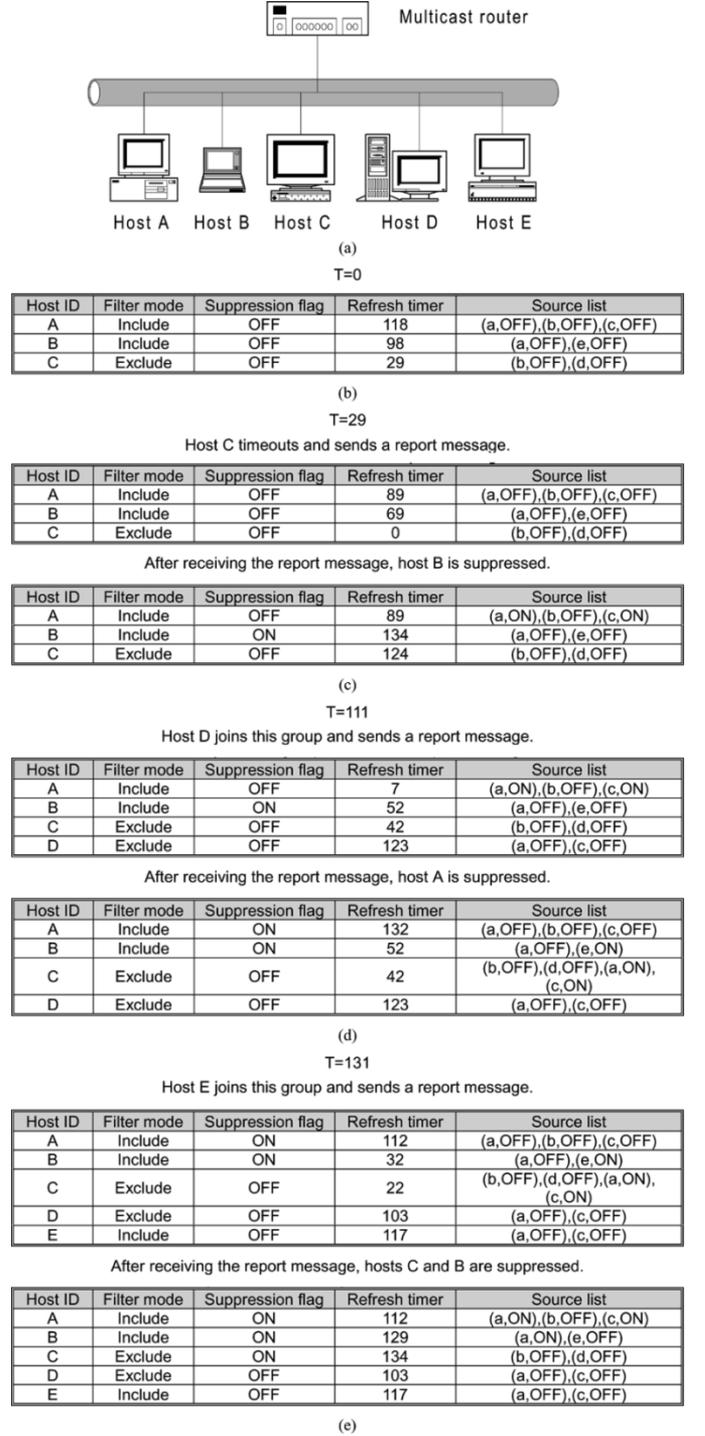


Fig. 2. An operation example.

III. PERFORMANCE EVALUATION

To compare IGMP with RGMP, we conducted simulations with our developed C++ simulator to examine protocol overhead in terms of bandwidth requirement³ on a local network for both protocols. We measured the performance with the following parameters: (1) the number of hosts participating in any

³In this section, protocol overhead is calculated by accounting for all control messages sent by each protocol, and is measured in terms of bandwidth requirements, rather than the number of control packets.

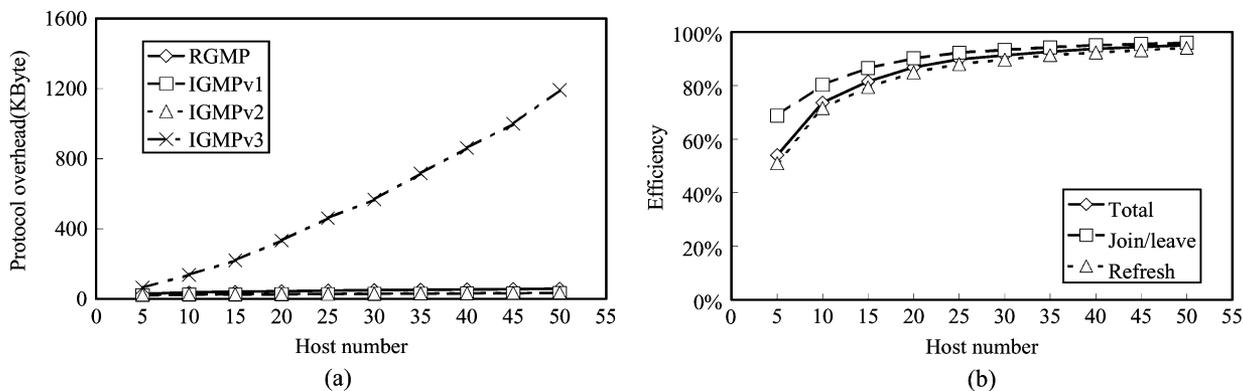


Fig. 3. Host number vs. protocol overhead, without source filter (a) Total protocol overhead (b) Efficiency.

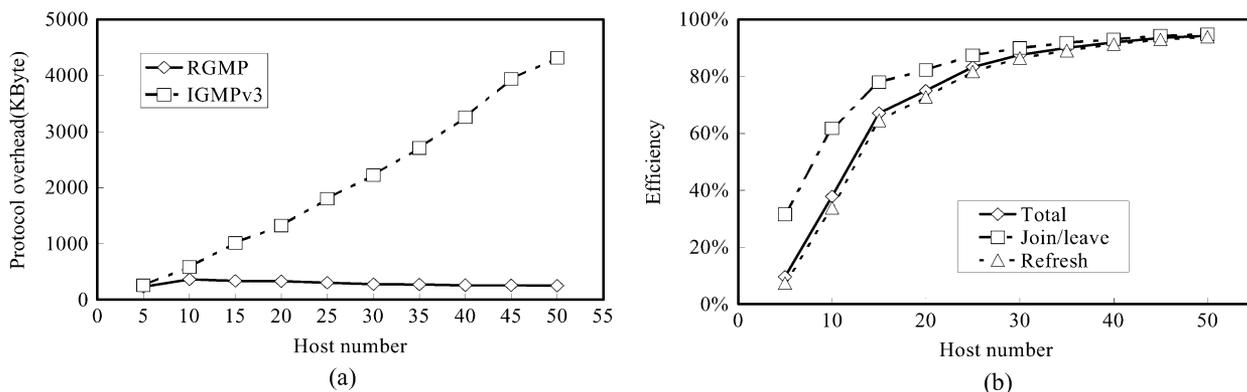


Fig. 4. Host number vs. protocol overhead, with source filter (a) Total protocol overhead (b) Efficiency.

group on the network (i.e., host number), (2) the number of hosts currently participating in one group (i.e., group size), (3) the number of different groups on the network (i.e., group number), (4) the percentage of hosts with source filtering on the network, and (5) the query/reply model vs. receiver-initiated model. The simulation was performed over a three-hour span, and under three assumptions: (1) no packet loss, and thus 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 were assumed to be Poisson (with $\lambda = 1/5400$ per sec): the process that a host used to join or leave a group, and the process that a host used to change a source (i.e., wish to, or wish not to receive the source) within a group.

Since both IGMP v1 and v2 do not support source filtering, we just mainly compare RGMP with IGMP v3. Each comparison is illustrated by two figures: one without source filtering and the other with source filtering. Each host may select a list of sources per group from up to 15 different sources. For those without source filtering, all the members in a same group have the same filter mode and the source list; for those with source filtering, each group member is free to select a filter mode and a source list from 15 different sources. For example, host 1 may have a source filter of `include{1, 2, 4}`, and host 2 may have a source filter of `exclude{1}`. Note that the protocol overheads of IGMP v1 and v2 are included in the figures without source filtering, mainly for reference. To observe the joint and individual

impacts of different message types on each compared parameter, each figure is further decomposed into two subfigures: (a) shows the total protocol overhead of each protocol, including join, departure, periodical refresh, and state change messages, and (b) shows the efficiency of RGMP compared to IGMP v3's, where the efficiency is defined as follows:

$$\text{Efficiency} = 1 - \left(\frac{\text{RGMP protocol overhead}}{\text{IGMPv3 protocol overhead}} \right)$$

The higher the efficiency, the better the performance improvement of RGMP over IGMP v3 for each comparison item. For example, the $x\%$ of efficiency means that the bandwidth requirement of RGMP is only $(100-x)\%$ of IGMP v3's. Note that since state change has similar performance trend to the join/departure of groups for both protocols, we include only the join/departure curve in the figures.

A. Host Number

This experiment was conducted to observe the impact of total host number on a local network on the performance of both protocols. We varied host number from 1 to 50. Each host could participate in up to 30 different groups. Fig. 3 shows the results without source filtering, and Fig. 4, with source filtering. Both figures show that IGMP v3 overhead increases as host number increases, while RGMP overhead is almost invariant to the increase of host number on a LAN. Interestingly, the efficiency of

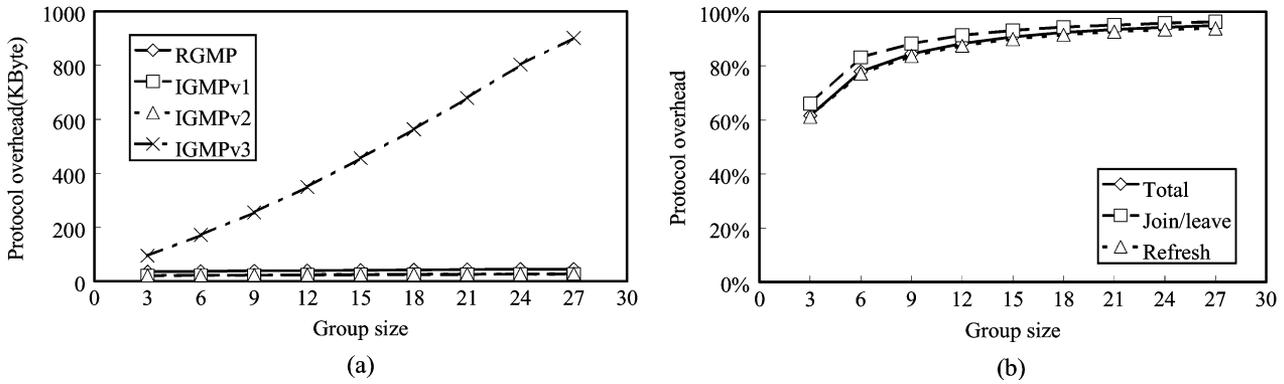


Fig. 5. Group size vs. protocol overhead, without source filter (a) Total protocol overhead (b) Efficiency.

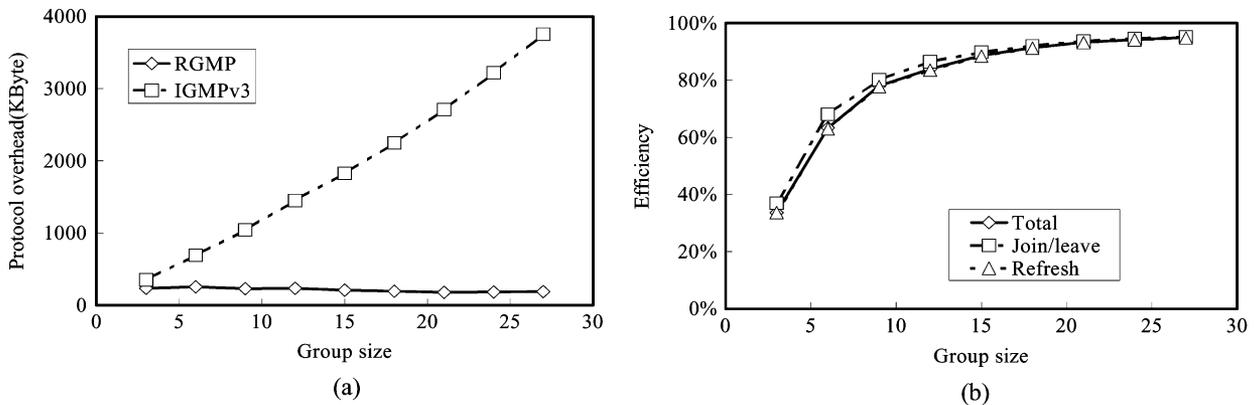


Fig. 6. Group size vs. protocol overhead, with source filter (a) Total protocol overhead (b) Efficiency.

RGMP over IGMP v3 stays relatively high as host number becomes large in both figures (subfigure (d)). In addition, RGMP protocol overhead is very close to that of IGMP v1/v2 as shown in Fig. 3, 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.

B. Group Size

This experiment was conducted to observe the impact of group size on the performance of both protocols. The group size was varied from 0 to 27. We considered 30 hosts on the network, each participating in up to 30 different groups. Fig. 5 shows the result without source filtering, and Fig. 6, with source filtering. Again, the efficiency of RGMP over IGMP v3 is relatively high as group size becomes large. Figs. 5 and 6 are similar to Figs. 3 and 4, respectively. This is because both cases are based on the same number of groups on a network. Thus, a change in group size is similar to a change in host number on the network. Considering the case without source filtering during a refresh period, IGMP v3 protocol overhead is approximately equal to 8 bytes * group size (or host number) * group numbers, while RGMP overhead is approximately equal to 8 bytes * group numbers.

C. Group Numbers

This experiment was conducted to observe the impact of group number on the performance of both protocols, varying group number from 1 to 50. We considered 20 hosts in total, each of which could participate in any group. Fig. 7 shows the result without source filtering, and Fig. 8, with source filtering. Both figures depict that the overheads of both IGMP v3 and RGMP increase as group number increases. However, IGMP v3 overhead increases more rapidly than RGMP's as group number increases. The efficiency of RGMP to IGMP v3 stays fairly flat in both figures. This is because when host number is fixed, as group number increases, the number of groups participated by each host increases, and thus the size of a single report message of IGMP v3 increases accordingly. The efficiency of RGMP over IGMP v3 is thus significant, thanks to report suppression. Group size, on 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 case without source filtering during a refresh period, the protocol overhead of IGMP v3 is approximately equal to 10 * 8 bytes * group numbers, and RGMP overhead is approximately equal to 8 bytes * group numbers, giving an overhead ratio of 10:1.

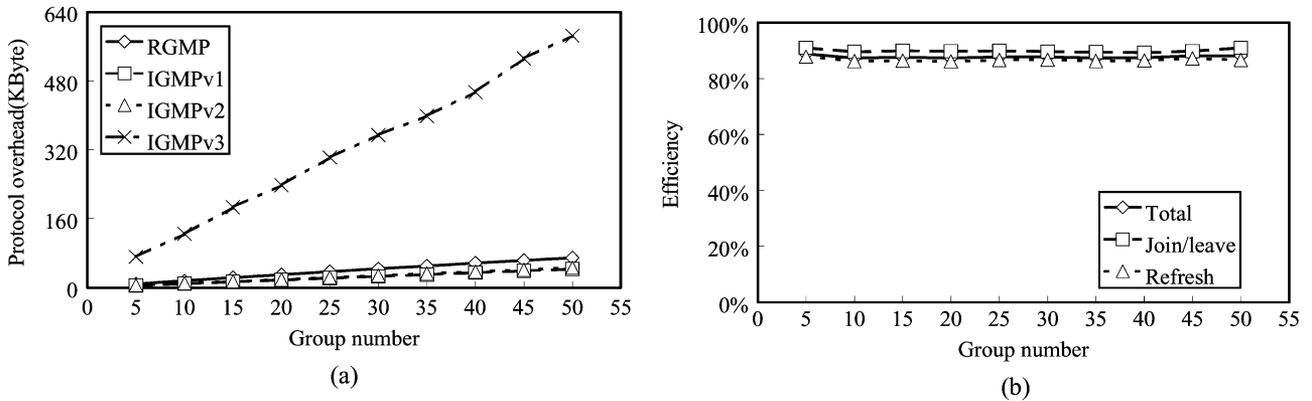


Fig. 7. Group number vs. protocol overhead, without source filter (a) Total protocol overhead (b) Efficiency.

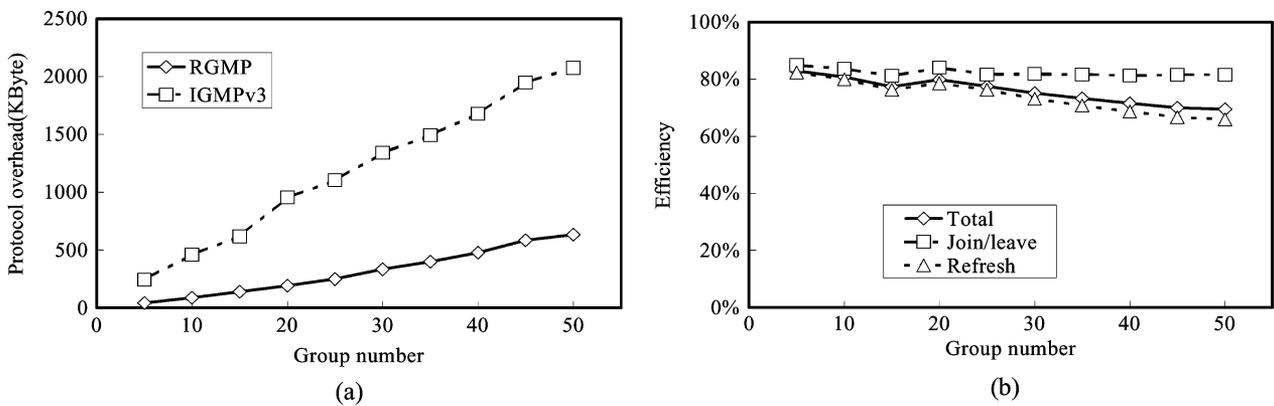


Fig. 8. Group number vs. protocol overhead, with source filter (a) Total protocol overhead (b) Efficiency.

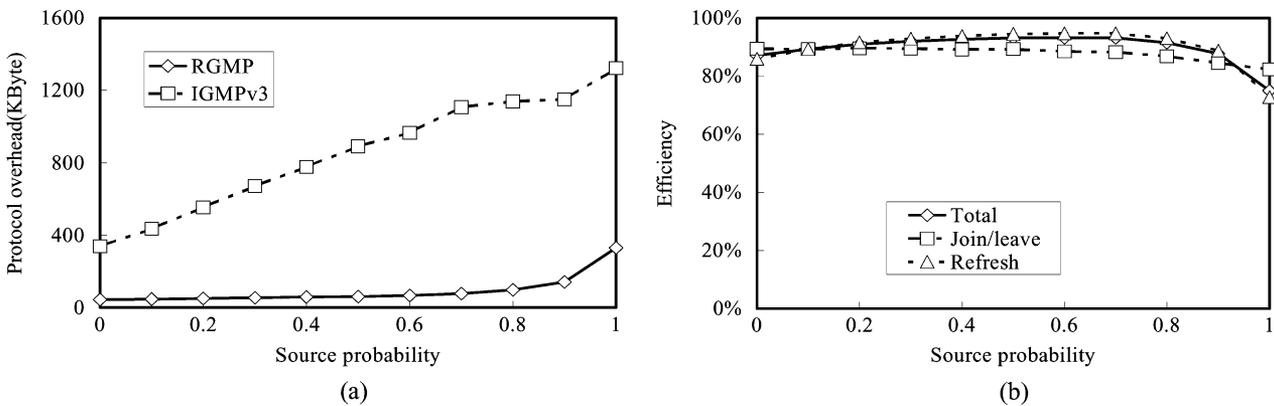


Fig. 9. The percentage of hosts with source filtering vs. protocol overhead (a) Total protocol overhead (b) Efficiency.

D. The Percentage of Hosts With Source Filtering on a LAN

This experiment was conducted to observe the impact of the percentage of hosts with source filtering on a network on the performance of both protocols. We considered 20 hosts in total, each of which could participate in up to 30 different groups and could select up to 15 different sources per group. The percentage of hosts with source filtering on the LAN was varied from 0% to 100%. Surprisingly, as shown in Fig. 9, when 80% of hosts

with source filtering, the efficiency of RGMP over IGMP v3 exceeds 90%. Even when all hosts on the LAN have source filtering (namely, 100% of hosts with source filtering), the efficiency of RGMP over IGMP v3 is still over 60%.

E. Query/Reply vs. Receiver-Initiated Model

This experiment was conducted to compare the query/reply model of IGMP with the receiver-initiated, self synchronized

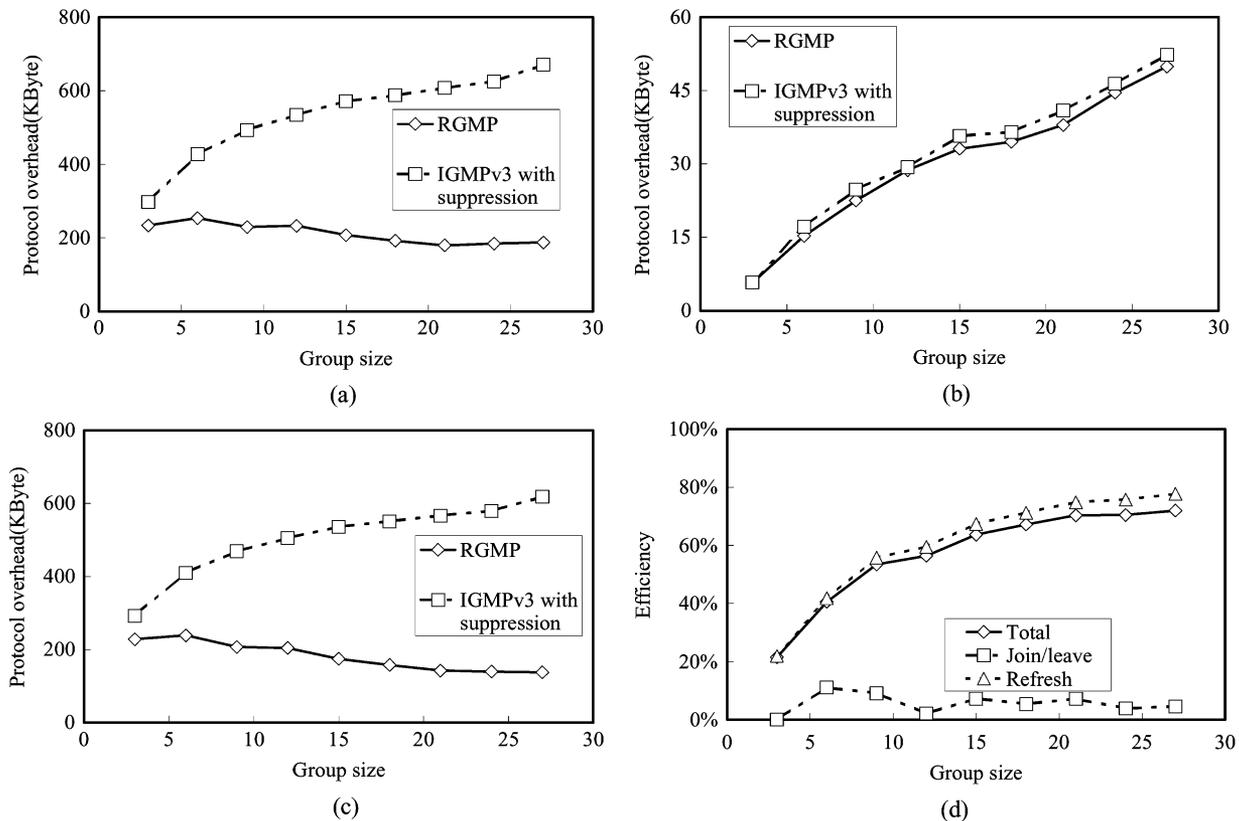


Fig. 10. Query/reply vs. receiver-initiated, self-synchronized model (a) Total protocol overhead (b) Protocol overhead from join or leave (c) Protocol overhead from periodical refresh (d) Efficiency.

model of RGMP. We modified IGMP v3 to include the RGMP suppression mechanism, and observed the protocol overheads of both protocols. Note that here we just show the result with different group sizes. The result was consistent to the others with different parameters. Due to space consideration, we just include the one with different group size. Fig. 10 shows that even if IGMP v3 supports suppression, RGMP still performs much better. The reason is that IGMP adopts a query/reply model. Thus, with IGMP, suppression takes effect only during refreshing periods. RGMP, on the other hand, employs the receiver-initiated refresh timer which is adjusted in a way to be adaptive and self-synchronized. Thus, RGMP allows all kinds of report messages to suppress refresh messages, including first join, last departure, state change, and periodical refresh.

IV. CONCLUSION

In this paper, we have proposed a new group management protocol, called RGMP, for IP multicasting. The protocol characteristics have been described, and the operation of the protocol has been presented in details. The comparison 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.

REFERENCES

- [1] M. H. Ammar, G. Polyzos, and S. Tripathi, "Special issue on networked support for multipoint communications," *IEEE J-SAC*, vol. 15, Apr. 1997.
- [2] J. C. Pasquale, G. C. Polyzos, and Xylomenos, "The multicasting problem," *ACM Multimedia Systems*, vol. 6, no. 1, pp. 43–59, 1998.
- [3] H. Eriksson, "Mbone: The multicast backBone," *Commun. ACM*, vol. 37, no. 8, pp. 54–60, Aug. 1994.
- [4] T. Wong, R. Katz, and S. McCanne, "An evaluation of preference clustering in large-scale multicast applications," in *Proc. IEEE INFOCOM*, 2000, pp. 451–460.
- [5] S. Song, Z. Zhang, B. Choi, and D. H. C. Du, "Protocol independent multicast group aggregation scheme for the global area multicast," in *Proc. IEEE GLOBECOM*, 2000, pp. 370–375.
- [6] A. Fei, J. Cui, M. Gerla, and M. Faloutsos, "Aggregated multicast: An approach to reduce multicast state," in *Proc. IEEE GLOBECOM*, 2001, pp. 1595–1599.
- [7] M. J. Moyer, J. R. Rao, and P. Rohatgi, "A survey of security issues in multicast communications," *IEEE Networks*, pp. 12–23, Nov./Dec. 1999.
- [8] P. Judge and M. Ammar, "Security issues and solutions in multicast content delivery: A survey," *IEEE Networks*, pp. 30–36, Jan./Feb. 2003.
- [9] S. Deering, C. Partridge, and D. Waitzman, "Distance vector multicast routing protocol," in *IETF RFC 1075*, Nov. 1988.
- [10] J. Moy, "Multicast routing extensions for OSPF," *Communication of the ACM*, vol. 37, no. 8, pp. 61–66, Aug. 1994.
- [11] A. Ballardie, J. Crowcroft, and P. Francis, "Core based tree (CBT)—An architecture for scalable inter-domain routing protocol," in *Proc. ACM SIGCOM*, Oct. 1993, pp. 85–95.
- [12] S. Deering, D. Estrin, D. Fairnacci, V. Jacobson, C. Liu, and L. Wei, "An architecture for wide-area multicast-routing," in *Proc. ACM SIGCOMM*, Oct. 1994, pp. 126–135.
- [13] S. Deering, "Host extensions for IP multicasting," in *IETF RFC 1112*, Aug. 1989.
- [14] W. Fenner, "Internet group management protocol, version 2," in *IETF RFC 2236*, Nov. 1997.
- [15] B. Haberman and J. Martin, "IGMPv3 and multicast routing protocol interaction, in IETF Internet Draft, July 2001. draft-ietf-idmr-igmpv3-and-routing-01.txt.



Wanjiun Liao received the B.S. and M.S. degrees from National Chiao Tung University, Taiwan, in 1990 and 1992, respectively, and the Ph.D. degree in Electrical Engineering from the University of Southern California, Los Angeles, California, USA, in 1997. She joined the Department of Electrical Engineering, National Taiwan University (NTU), Taipei, Taiwan, as an Assistant Professor in 1997. Since August 2000, she has been an Associate Professor. Her research interests include the design and analysis of multicast and QoS protocols for

wireless networks and broadband Internet.

Dr. Liao is actively involved in the international research community. She is currently an Associate Editor for IEEE Transactions on Wireless Communications. Dr. Liao has received many research awards. She was a recipient of the Outstanding Research Paper Award in Electrical Engineering at the University of Southern California in 1997. Two papers she co-authored with her students received the Best Student Paper Award in the First IEEE International Conferences on Multimedia and Expo (ICME) in 2000, and the Best Paper Award in the First International Conference on Communication, Circuits and Systems in 2002. Dr. Liao was elected as one of Ten Distinguished Young Women in Taiwan in 2000 and is listed in the Marquis Who's Who in 2001–2003, and the Contemporary Who's Who in 2003.



De-Nian Yang was born in Taiwan in 1977. He received the B.S. degree from the Department of Electrical Engineering, National Taiwan University, Taipei, Taiwan, in 1999. Exempted from the normal entrance exam to the graduate school, he is currently a Ph.D. candidate in the EE department of National Taiwan University. He received the best student paper award in ICME 2000. His research interests include broadband Internet, QoS, and multicasting.