

INTERNET GROUP MANAGEMENT PROTOCOL: AN OBJECT ORIENTED SIMULATION

A.Halim ZAIM

Department of Computer Science
North Carolina State University
NC, 27695, USA

E-mail: ahzaim@eos.ncsu.edu

ABSTRACT

This paper explains a simulator designed to test IGMPv2 which is used by IP hosts to report their multicast group membership termination to be quickly reported to the routing protocol, which is important for high-bandwidth multicast groups and/or subnets with highly volatile group membership.

1. INTRODUCTION

IGMP is a part of the IP layer similarly with Internet Control Message Protocol (ICMP) [1]. As a consequence, like ICMP, IGMP messages are transmitted in IP datagrams. On the other

hand, IGMP has a unique message structure with a fixed-size message type without optional data. Figure 1 shows the encapsulation of an IGMP message in an IP datagram.




Figure 1: Encapsulation of an IGMP message within IP datagram

The format of an IGMP message is given in Figure 2 [2]. The TYPE field can get four different values three of which are the main messages for host-router interaction, and one type of message is used for backward-compatibility. If type is 0x11, it means a

membership query, if it is 0x16, it means a version 2 membership report, if it is 0x17, it means a leave group message and if it is 0x12, it means a version 1 membership report. A Membership query can be in two ways: 1) General query, used to learn which groups have

Received Date : 7.9.2001

Accepted Date: 2.6.2002

members on an attached network, 2) Group-specific query, used to learn if a particular group has any members on an attached network. These

two messages are differentiated by the Group Address.



Figure 2: Format of an IGMP message

MAXIMUM RESPONSE TIME field is used to specify the maximum allowed time before sending a responding report in units of 1/10 second during a membership query message. In other type of messages, it is set to zero and ignored by the receivers.

CHECKSUM is calculated similarly with TCP and IP checksums. It is the 16-bit one's complement of the one's complement sum of the whole IGMP message (the entire IP payload).

GROUP ADDRESS is used by the hosts to indicate which multicast group they belong to. It is set to zero in a general query, and set to the group address in a group-specific query. It is the same with the IP multicast group address of the group being reported or left in cases of a membership report or leave group messages.

In case of an IGMP version one message, the message size may be longer than 8 bytes. In that case, the entity dealing with the message ignores anything passing the first 8 bytes [3].

After giving the general structure of an IGMP message, we can start working on elements of the protocol infrastructure. The next section explains the network representation. Section three gives the structure of the Host class. The fourth section concentrates on Routers. After

explaining the Simulation in section 5, we conclude our paper in section 6 with a Conclusion.

2. NETWORK

IGMP is used by IP hosts capable of receiving multicast traffic to notify any immediately neighboring multicast routers of their memberships. Since this protocol is concerned with the interactions between hosts and routers connected to each other on the same physical network, we abstract this network as a collection of hosts and routers. These individual entities compromising the network are configured during set-up phase of the simulation based on the parameters given in the configuration file [4].

Table 1 presents the format of the configuration file. The first three parameters are the number of routers, the number of hosts and the number of groups in the network. Once we give those numbers, we give necessary information about each group such as popularity, in volatility and out volatility. A particular host's membership in a particular group at any point in time is determined by the applications running in that machine. The applications' average behavior with respect to a certain multicast group can be simulated by parameters like popularity, in volatility and out volatility.

Table 1: Format of the Configuration File

#Routers			
#Hosts			
#Groups			
Group#1	Popularity	In Volatility	Out Volatility
Group#2	Popularity	In Volatility	Out Volatility
Group#3	Popularity	In Volatility	Out Volatility
...
...
Group#n	Popularity	In Volatility	Out Volatility

We view a network of m routers and n hosts as a collection of $m+n$ different entities each having its own unique id (similar of IP address), where routers have ids from 1 through m and hosts have ids from $m+1$ through $m+n$. IGMP messages are encapsulated in IP datagrams before they are transmitted onto the network [5, 6]. The destination IP address in the IP header is one of the following:

All Systems Multicast Group (224.0.0.1)

All Routers Multicast Group (224.0.0.2)

A valid class D address representing a multicast group.

This destination address determines which host(s) and router(s) will receive the datagram. In our simulation, we emulate multicast by contacting each and every intended recipient, host(s) and/or router(s), on the network. We also provide for packet loss in congested routers by simulating random packet drops at routers.

3. HOSTS

Hosts are the first of the two active elements in the system. They do not have any knowledge about the system, their only role is acting as a passive element leaving control of the system to the routers. Hosts need to perform their actions on all interfaces that have memberships associated with them. Each host sets a delay timer per group of which it is a member, each time it receives a General Query. Each timer is set to a different random value. The range of this random value is between 0 and Maximum Response Time. At the setup phase, each host is given a host id number ranging between $nrouter+1$ and $router+nhost$. Their initial states

are also set to NonMember. Then each host creates a Join Event for a future time and put that event to the event calendar.

After the setup phase has been completed, the event handler of each hosts take control. Host Event Handler receives four parameters: the event triggering the event handler, the event calendar used in case of new events to be added, group information and the trace parameter. First, the type of the event is checked. If this is a Query Received event, the host checks whether this is a group specific query or a general query. If this is a Group Specific Query, then it passes control to the function SpecificQuery which handles the query only for that group. Otherwise, it calls the function GeneralQuery, which handles the queries for every group. These two functions checks also version of the incoming query. If this is a Version 1 query, it stays in the IGMP Version 1 Router Present state. Otherwise it keeps the state No IGMP Version 1 Router Present.

The second event is the Version 1 Watch Timer Expiration. This event shows that there is no more version 1 router in the system. Therefore, the host can change to the state No IGMP Version 1 Router Present.

The third event is a Join Group event. This event triggers three more events. Once a host is joined to a multicast group, it changes its state from No Member Present to Delaying Member. At the same time, it sends immediately a version 2 report to see if there are other members of that group. To ensure that this report has been reached to its destination, it sets a report timer. If this report is lost, another report is scheduled.

The last part of the Join Event is scheduling a Leave Event to end that session.

The fourth event is Timer Expiration. In this case, the host state changes to Idle.

The fifth event is the Leave Group event. This event makes the host change to state Nonmember. This event triggers another Join event.

The last event is a Report Reception event. This may be either a version 1 or version 2 report. In that case, the host goes to the Idle state and stops the timer.

The state diagram for a host is given in Figure 3.

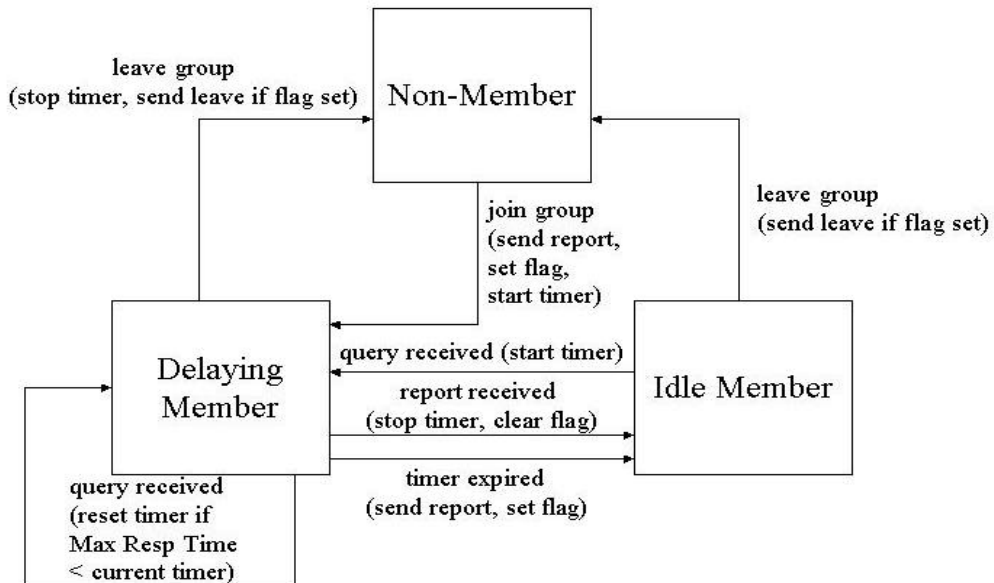


Figure 3: Host State Diagram

As shown in Figure 3, a host may be in one of three possible states with respect to any single IP multicast group on any single network interface:

1. Non-Member State:, when the host does not belong to the group on the interface. This is the initial state for all memberships on all network interfaces; it requires no storage in the host.
2. Delaying Member State, when the host belongs to the group on the interface and has a report delay timer running for that membership.
3. Idle Member State, when the host belongs to the group on the interface and does not have a report timer running for that membership.

There are seven possible actions that may be taken in response to the above events:

1. "send report" for the group on the interface. The type of report is determined by the state of the interface. The Report Message is sent to the group being reported.
2. "send leave" for the group on the interface. If the interface state says the Querier is running IGMPv1, this action SHOULD be skipped. If the flag saying we were the last host to report is cleared, this action MAY be skipped. The Leave Message is sent to the ALL-ROUTERS group (224.0.0.2).
3. "set flag" that we were the last host to send a report for this group.

4. "clear flag" since we were not the last host to send a report for this group.
5. "start timer" for the group on the interface, using a delay value chosen uniformly from the interval (0, Max Response Time], where Max Response time is specified in the Query. If this is an unsolicited Report, the timer is set to a delay value chosen uniformly from the interval (0, [Unsolicited Report Interval]].
6. "reset timer" for the group on the interface to a new value, using a delay value chosen uniformly from the interval (0, Max Response Time], as described in "start timer".
7. "stop timer" for the group on the interface. In all of the following state diagrams, each state transition arc is labeled with the event that causes the transition, and, in parentheses, any actions taken during the transition. Note that the transition is always triggered by the event; even if the action

is conditional, the transition still occurs.

From the perspective of the network interface, a host may be in one of two possible states.

1. No IGMPv1 Router Present: This is the initial state in case that the host doesn't receive any version 1 query.
2. IGMPv1 Router Present: As long as the host receives a version 1 query, it stays in that state. If it doesn't hear from a version 1 query during a predetermined version 1 router present timeout, it goes back to the state No IGMPv1 Router Present. The state diagram is given in Figure 4.

There are two events that can cause state transitions: IGMPv1 query received and timer expires. When the host receives a version 1 query with the max response time field set to 0, it is treated as an IGMPv1 query and it takes the necessary action. In the second case, the timer is set to note the presence of an IGMPv1 router expires.

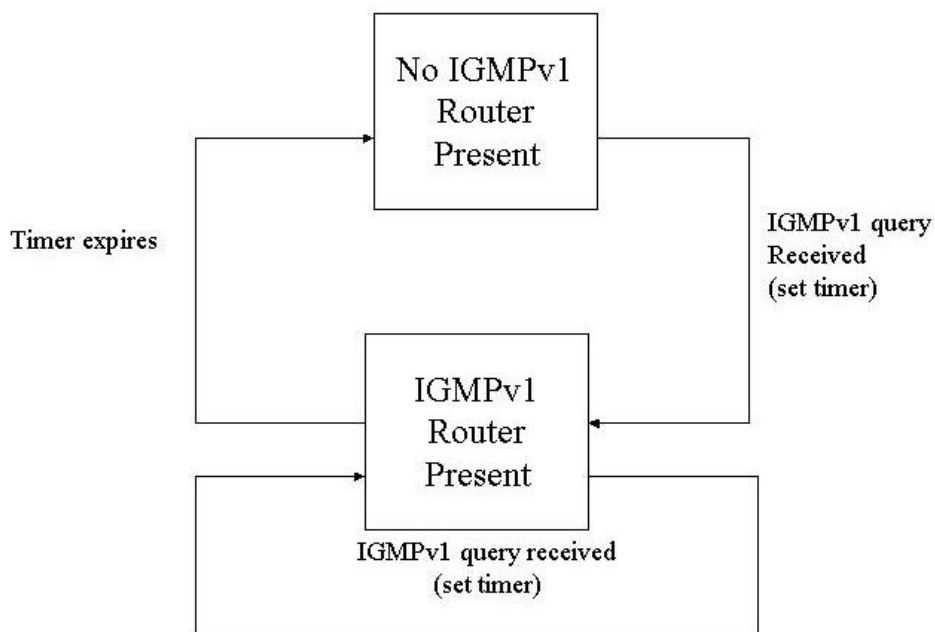


Figure 4: Host State Diagram From the Network Perspective

4. ROUTERS

The router is the second type active element in the system. Unlike hosts, a router has total control on the system. A router uses IGMP messages to keep track of those multicast groups which have members on each of its attached physical networks. Routers need not keep track of which or how many hosts are members; they need only know whether there exist any members on the attached physical network.

A multicast router may assume one of two roles: Querier or Non-Querier. A Querier router sends General Query messages every 125 seconds.

IGMPv2 dictates that at start-up, each router must assume the role of Querier. If a router then sees a Query from a router with a lower IP address, it becomes Non-Querier. It also sets the Other Querier Present Timer, for a time interval of Other Querier Present Interval (255 seconds in our simulation). Each time a Non-Querier sees a Query from a router with a lower IP address, it resets this timer. Should the timer expires, then the Non-Querier takes on the role of Querier. The state diagram of a router is given in Figure 5.

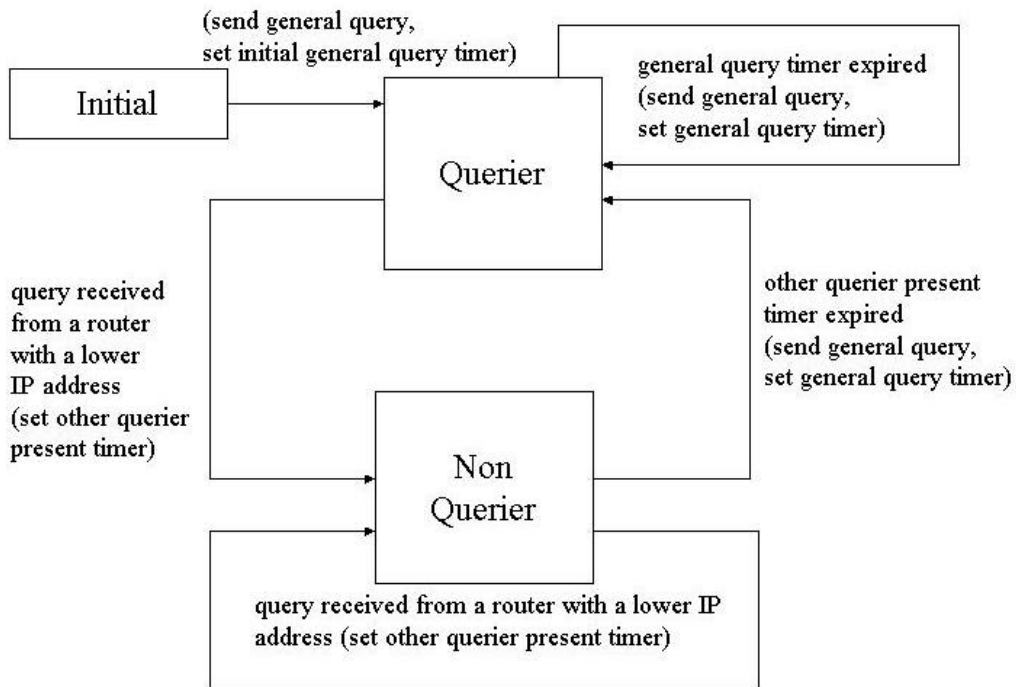


Figure 5: Router State Diagram

A router should start in the Initial state on all attached networks, and immediately move to Querier state. In addition, to keep track of which groups have members, a router, once changed its state to Querier, uses another state diagram as shown in Figure 6.

As seen in Figure 6, a router may be in one of four possible states with respect to any single IP multicast group on any single attached network. The initial state for all groups on the router is the

No Members Present state during which the router sends reports for multicast groups. Once the router receives a membership report from any host on the network, it changes to Members Present state. If the router receives a version 1 membership report from any host, then it changes its state into Version 1 Members Present state. The Checking Membership state is used for cases when the router has received a Leave Group message but has not yet heard a Membership Report for the multicast group.

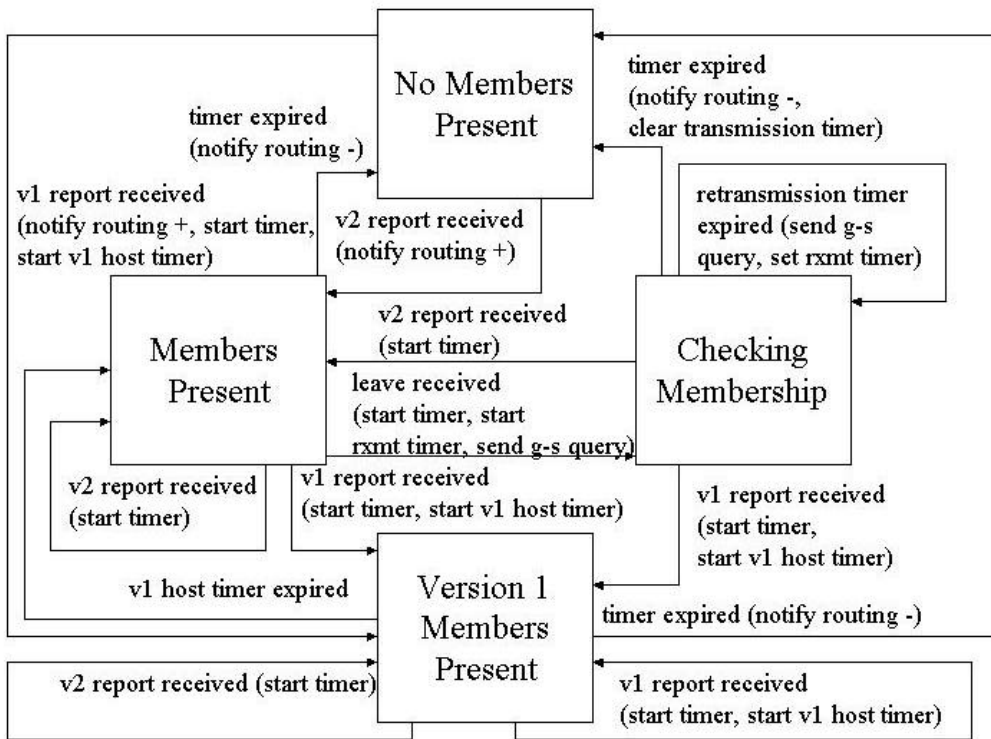


Figure 6: State Diagram of a Router in Querier State

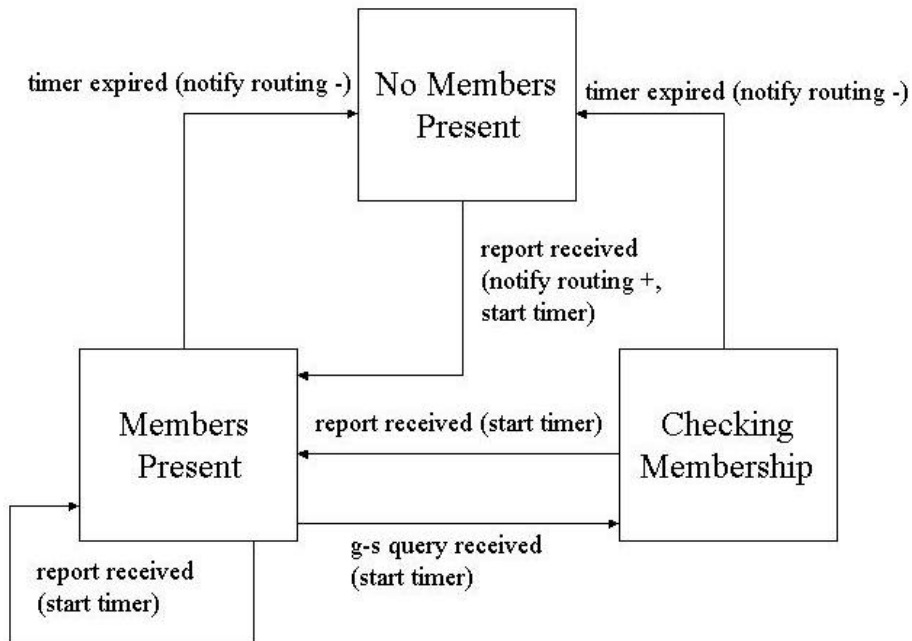


Figure 7: State Diagram of a Router in Non-Querier State

The state diagram for a router in Non-Querier state shown in Figure 7 is similar, but non-Queriers do not send any messages and are only driven by message reception. Note that non-Queriers do not care whether a Membership Report message is Version 1 or Version 2.

4.1. Members

As prescribed by its base class Entity, each router has ID number representing its IP address. Router Ids vary from 1 to n routers (Recall that nrouter is a parameter read from the configuration file at the beginning of the simulation).

Specific to the derived class Router, there are three additional members. First, the member role has the value 1 if the router is Querier, 0 otherwise. Second, the member nGroups holds the number of multicast groups that may or may not have on the attached physical network over the course of the simulation. Third and finally, the member groupList is an array of size nGroups which holds state information for each multicast group. Each element in the array is an hg struct, defined in global.h.

4.2. Methods

At the beginning of the simulation, method SetUp uses information from the configuration file to initialize groupList. ScheduleFirstEvents is then called to place the router's first event in the Event Calendar, InitialToQuerier.

The primary method of the Router class is EventHandler. When an event is removed from the Event Calendar with a scope of either AllSystems or AllRouters, then the router's EventHandler is called upon to handle that event. Based on the type of the event and the role of the router (Querier or Non-Querier), EventHandler takes steps to make the necessary changes in state. For many event types, EventHandler employs the method AdjustState to accomplish this task.

Below, several other methods and their functionality are described briefly:

1. **MakeEvent:** To accomplish such tasks as setting a timer or sending an IGMP message, a router uses MakeEvent to create an event of the desired type and

then schedules it in the Event Calendar. For example, to send a General Query, a router makes an event of type QueryRcvd with groupId set to 0 (the AllSystems group). It then schedules the event in the Event Calendar with a delay of 0 by calling schedule, an Event Calendar method.¹

2. **ValidateGroupId:** Before processing a message, a router uses ValidateGroupId to check whether the message belongs to a host or router on the same multicast group.
3. **NotifyRouting:** NotifyRouting must be called to contact routers located one hop away that the multicast tree must be pruned or extended.

4.3. Router Events

In this subsection, we briefly describe the important events affecting the state of a router, together with the implementation details associated with carrying out these state changes. For the sake of simplicity, we introduce the following shortcut in our descriptions. We shall state only that "a message is sent" or "a timer is set". It should be understood that such actions require the router first to create the event with the appropriate parameters (MakeEvent) and then to schedule the event in the Event Calendar.

The event InitialQuerier occurs at the beginning of the simulation. The member role is set to 1 for Querier, a General Query is sent, and the General Query Timer is set.

The following four events correspond with the arrival of an IGMP message:

1. QueryRcvd (Query Received)
2. V2ReportRcvd (v2 Report Received)
3. V1ReportRcvd (v1 Report Received)
4. LeaveRcvd (Leave message received)

The remaining six events have to do with a timer expiring:

1. GenQueryTimerExpired (General Query Timer Expired)
2. OtherQuerierPresentTimerExpired

¹ We assume transmission delay is negligible, i.e. zero. Hence IGMP messages are scheduled to be received immediately when they are sent.

- (Other Querier Present Timer Expired)
- 3. TimerExpired
- 4. REXMITTimerExpired (Retransmit Timer Expired)
- 5. V1HostTimerExpired (v1 Host Timer Expired)
- 6. UpdateLatency (Update Latency)

5 SIMULATION

We implemented the simulation as a next-event simulation. In this strategy, each event is uniquely identified by its type, time and owner, where type can be one of several possible types of events and owner denote the id (IP address) of the host/router that scheduled the events. We can classify all the possible types of events in three categories as shown in Table 2.

Table 2: Events used in the simulation

Event
IGMPQUERY
IGMPREPORT
IGMPLEAVE
HLEAVEGROUP
HJOINGROUP
HTIMEREXPIRED
HV1WATCHTIMEREXPIRED
RGENQUERYTIMEREXPIRED
ROTHERQUERIERPRESENTTIMEREXPIRED
RTIMEREXPIRED
RREXMITTIMEREXPIRED
RV1HOSTTIMEREXPIRED
RINITIATIONTOQUERIER
RUPDATELATENCY

We now explain each event in detail.

IGMP QUERY: Querier (router) schedules this event to send a General query or a Group-Specific query. This event has an associated Query message which is delivered to all the intended recipients.

IGMP REPORT: A host schedules this event to send a Membership Report as per the protocol specifications. Although more than one host may schedule a report to be sent, some of these reports will have to be removed to simulate report suppression described in the protocol.

IGMP LEAVE: A host leaving the multicast group schedules this event to send a Leave message. It would do so if it thinks that it was the last host to send a membership report.

H LEAVE GROUP: This event is scheduled to simulate a host's leaving a multicast group as dictated by the application(s). When this event is processed, the host may decide to schedule IGMP LEAVE event as the immediately next event.

H JOIN GROUP: This event is scheduled to simulate a host's joining a multicast group as dictated by the application(s). When this event is processed, the host schedules as IGMP REPORT event as the immediately next event.

H TIMER EXPIRED: This event signifies that the timer set by a certain host has expired. The host which scheduled this event then takes appropriate action depending on its state with respect to the multicast group for which this timer was scheduled.

H V1 WATCH TIMER EXPIRED: This event indicates that there is not any IGMPv1 router in the system, because, the host didn't hear any IGMPv1 query for the duration of that timer. In this case, the host changes its state to No IGMPv1 router present state.

R GEN QUERY TIMER EXPIRED: This event signifies that it is time to send a General Query to find out which groups have members on the attached network. The router that had scheduled this event then schedules IGMP QUERY as immediately next event.

R OTHER QUERIER PRESENT TIMER EXPIRED: This event is scheduled by Non-Querier router(s) to help them find out if they should start playing the role of Querier in the absence of one.

R TIMER EXPIRED: This event signifies that the timer set by a router has expired. The router which scheduled this event then takes appropriate action depending on its state with respect to the multicast group for which this timer was scheduled.

R REXMIT TIMER EXPIRED: This event is scheduled by a Querier router. It indicates that the retransmission time for a group specific

membership query expired. Then the router stays on the Checking Membership state and continue sending group specific queries by setting the retransmit timer.

R V1 HOST TIMER EXPIRED: The router use this timer to find out if there are no IGMPv1 hosts on the network. When this event happens, the routers assume that there are no IGMPv1 hosts on the attached network and behave accordingly.

R INITIAL TO QUERIER: Whenever a router comes up it is supposed to move from Initial to Querier state. This transition also requires Query messages to be sent on the network. The routers schedule this event to help them make the state transition from described.

R UPDATE LATENCY: The RFC indicates possible tuning of Leave Latency by adjusting maximum response time in the Query message sent after a Leave message is received. We believe that the routers will need to review maximum response time periodically for each group. This event allows the router to schedule an event to indicate the need for updating the latency.

6 CONCLUSION

In this paper, we explained the main elements of a simulator designed to test the IGMPv2 protocol. The simulator took all the features of the protocol. We designed the maximum response time as a variable so that we can tune the leave latency and the burstiness of IGMP traffic on a subnet. That way, we had the chance

of testing the behavior of the protocol under bursty traffic.

According to RFC2236, a querier router must use IGMPv1 if there is an IGMPv1 router in the system. The decision of whether IGMPv1 or v2 will be used, can not be adjusted dynamically. Because there is no reliable way of doing it. Therefore, in RFC, it is suggested that, the administrators may enable the use of IGMPv1 on their routers. However, if the configuration is not determined explicitly, then the configuration must use default to IGMPv2.

At that point, we thought, it might be possible to use the state diagram suggested for hosts to adjust between IGMPv1 and v2 for routers also. That way, it may be possible to change dynamically the version of the routers.

7 REFERENCES

1. S. Deering, *Host Extensions for IP Multicasting*, RFC 1112, 1989.
2. W. Fenner, *Internet Group Management Protocol, Version 2*, RFC 2236, 1997.
3. W. Richard Stevens, *TCP/IP Illustrated, Volume 1*, Addison-Wesley Profesional Computing Series, USA, 1994.
4. Douglas E. Comer, *Internetworking with TCP/IP, Volume 1*, Prentice Hall, USA, 1995.
5. Douglas E. Comer, *Computer Networks and Internets*, Prentice Hall, USA, 1999.
6. Daniel Minoli and Andrew Schmidt, *Internet Architectures*, John Wiley & Sons, Inc., USA, 1999