

Entitled:

High Performance IP Multicasting Over
Wireless Satellite-Terrestrial Networks

Authors:

J.S. Baras, and I. Secka

Conference:

*17th AIAA International
Communications Satellite Systems
Conference and Exhibit
February 24-26, 1998*

HIGH PERFORMANCE IP MULTICASTING OVER WIRELESS SATELLITE - TERRESTRIAL NETWORKS*

John S. Baras

Isatou Secka

Center for Satellite & Hybrid Communication Networks
Institute for Systems Research
University of Maryland
College Park, MD 20742

Abstract

We describe our recent work on the design and implementation of high performance Internet services over networks consisting of interconnected high data rate satellites including Direct Broadcast Satellite hosts and terrestrial wireless LANs with various capabilities (with rates from 16 kbps to 10Mbps, including LMDS and MMDS systems). The network can use either bi-directional or receive only satellite links for downstream data delivery and wireless and wireline terrestrial or satellite links for the upstream path. A key concept in our work is that of a hybrid terminal, which is a PC connected to a satellite antenna (including just DBS antennas) and to the wireless LAN. The hybrid terminal uses a modem connection for outgoing traffic while receiving incoming information through the VSAT. The hybrid terminal is attached to the Internet through any Internet service provider who supports Serial Line Internet Protocol (SLIP). The traffic from the hybrid terminal is transmitted to the hybrid gateway through IP-within-IP encapsulation, to accomplish asymmetric routing. The hybrid gateway is responsible for decapsulation of traffic from hybrid terminals. It is also responsible for formatting data to suite the satellite transmission.

The asymmetric nature of traffic in most networks, as evident in the Internet, is shifting current networking technology trends more towards the development of hybrid networks. Multimedia traffic with its inherent variability in Quality of Service (QoS) requirements further reinforces this trend. Technologies such as DirecPCTM which allow users to send traffic terrestrially and receive traffic through satellite have demonstrated the efficiency of the broadcast nature of satellite communications as a means of delivering high bandwidth traffic to end

users. Even though the majority of Internet applications rely on point-to-point transmission (unicast), emerging applications such as teleconferencing and information distribution have necessitated the development of an overlay multicast backbone network in the Internet (MBONE) for point/multipoint-to-multipoint data transmission. A major hurdle in multicasting over the Internet is the potential for high bandwidth traffic to cause congestion in the terrestrial backbone. Introducing hybrid terminals within corporate LANs for incoming multicast streams thus would provide an effective means of preserving gateway bandwidth for other outgoing traffic.

We describe our work on IP multicast extensions to the wireless hybrid network described. We describe effective extensions of IGMP, and asymmetric multicast algorithms that exploit the asymmetry to increase the number of users, scale-up and improve the loading of the terrestrial components. This requires an asymmetric multicast routing mechanism. We describe enhancements to existing multicast routing protocols such as CBT to the hybrid environment described here. We provide results on performance of our proposed hybrid multicast algorithms with respect to the following performance metrics: time to join a group; time for a packet to reach every member of the multicast group; performance with large multicast groups.

Multicasting in Hybrid Networks

Traditional multicasting on the MBONE has been used for exchanging information between a group of users in applications such as video or audio conferencing but a major hurdle in multicasting over the Internet is the potential for high bandwidth traffic to cause congestion in the terrestrial backbone. For groups with many members that are sparsely distributed over a wide area, the multicast packets would have to traverse several links before reaching all group members, hence the potential for causing congestion. Some companies may wish to en-

*This work was supported in part by the Center for Satellite and Hybrid Communication Networks under NASA Cooperative Agreement NCC2-528, by Hughes Network Systems and by the State of Maryland under the MIPS program.

gage in multicast conferencing applications but may have limited gateway bandwidth to the Internet. For such users, introducing hybrid terminals within their corporate LAN to route incoming traffic through a satellite link would be a way of preserving the corporate wireline gateway bandwidth for other outgoing traffic. Another motivation of multicasting in hybrid networks is its use in military or medical applications, where individuals in remote areas equipped with hybrid terminals would be able to receive critical high data rate packets.

There are several issues to consider when extending multicast over hybrid networks. First and foremost, a group membership protocol has to be defined for keeping track of group membership information in the hybrid network. The work described in this paper is mainly directed towards developing asymmetric multicast routing techniques for constructing multicast trees at remote LANs, so that all outgoing traffic is directed toward the corporate wireline gateway while incoming multicast traffic comes through a satellite link. The protocol established for this special case (satellite-terrestrial) could then be extended to other hybrid networks.

Construction of a multicast tree gives the ability to both send and receive multicast packets. The motivation for multicasting is to support high-data-rate applications such as video conferencing. In hybrid networks where there is limited bandwidth on the uplink, it is impossible to support such applications. Hence, use of the asymmetric nature of hybrid networks for multicasting data makes sense only on the receiving end. Thus what we are doing, in effect, is **constrained** multicasting where hybrid hosts take advantage of the high bandwidth downlink to receive packets, but are restricted to sending only low-data-rate voice and data packets which can tolerate the degradation of quality.

One of the biggest challenges faced is that the asymmetric nature of traffic, out through the corporate LAN and in through a satellite receiver, creates the potential for the formation of loops, breaking the concept of tree construction completely. Further complications could arise at a multi-homed (multiple routers) local LAN with a hybrid host particularly when more than one router are multicast capable because this would make construction of an internal delivery tree difficult. Generally, Internet routing protocols were developed assuming bi-directional and symmetric links, and may no longer work in the uni-directional environment. For example, routers on the receiving end of a uni-directional link have no means of announcing routes

to feeds at the source of links because they cannot communicate directly with them. A subcommittee, the Uni-Directional Link Routing (UDLR) working group, has been formed at the Internet Engineering Task Force (IETF) to find solutions for dynamic routing problems caused by uni-directional links. The UDLR working group currently focuses on support of alternative uni-directional links on top of a bi-directional internetwork. There are currently two proposed approaches that address this problem. One is based on the modification of the common routing protocols to support uni-directional links. The other one proposes adding a layer between the network interface and the routing software to emulate bi-directional links through tunnels. Both approaches are being studied in order to come up with a solution for dynamic routing in the presence of uni-directional links.

The main objective of this paper is to develop a system-level design of a multicast routing protocol that would allow hybrid hosts in hybrid satellite-terrestrial networks to dynamically receive multicast packets.

Protocol Design

Protocol Architecture

For extending multicast protocols to hybrid networks we used a modified version of CBT, hereafter referred to as Hybrid Core-based Trees (HCBT), and assumed the architecture model described in [1] and [2] for Hybrid Internet Access. In addition, the HCBT architecture assumes the scenario illustrated in Figure 1 where we have:

- N users that want to form a multicast group
- Out of these N users, H are static HHs and (N-H) are terrestrial users on the MBONE
- Out of the H Hybrid Hosts, L are attached on LANs and (H-L) are "stand-alone" hybrid hosts. Note that the LANs also have terrestrial (wireline) access to the MBONE.
- The HHs attached on LANs may be responsible for forwarding packets to other users on the LAN.
- A modified version of IGMP is running between the HHs and the gateway.

For our system, we define two types of traffic, low-data rate or "short length" traffic (e.g. audio, web browsing), and high-data rate or "bulky" traffic

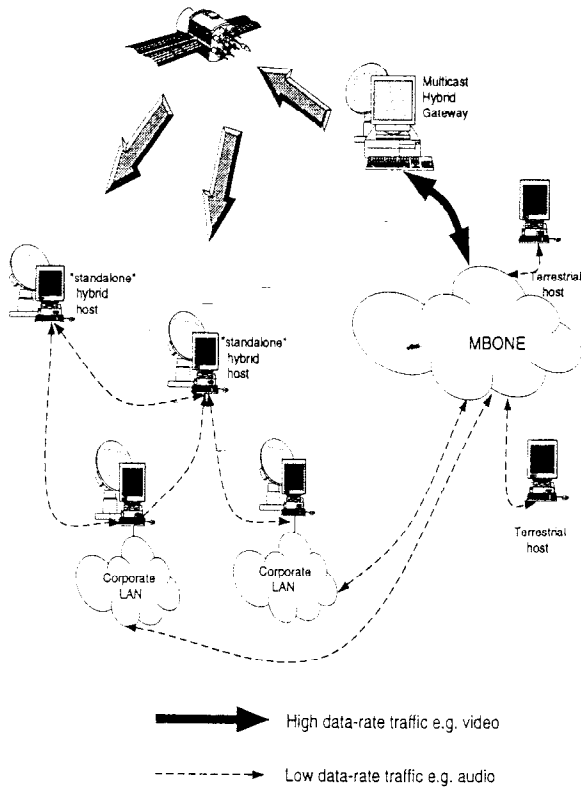


Figure 1: Diagram illustrating the HCBT architecture.

(e.g. video, images, books). All traffic below a certain threshold, T (bits/sec), is considered low-data rate traffic and all traffic with rate above T is considered high-data rate. Likewise, all traffic beyond size, S (bits), is considered "bulky", otherwise it is considered "short length". An intelligent routing scheme will be deployed that routes high-data-rate or "bulky" traffic through satellite and low-data-rate or "short length" traffic through the terrestrial network.

We are proposing that all HHs be required to join multicast trees through a Multicast Hybrid Gateway (MHGW) which is analogous to the Hybrid Gateway in Hybrid Internet Access architecture. It is assumed that the MHGW would be the IGMP querier for all HHs and is thus aware of group membership information of HHs. Necessarily, all multicast traffic to and from the HHs is routed through MHGW. When packets are multicast to a group with HH members, the MHGW would observe the data rate to determine whether to send them terrestrially or via satellite to HHs. If the latter is required, the packets are put on the satellite interface for broadcasting to the HHs.

Since packets put on satellite are broadcast and

would be available to everyone, some authentication mechanism need to be established to allow only HHs that are members of the group to receive multicast packets. Therefore some "key sending process" needs to be included in the IGMP version for Hybrid Networks so that when a HH registers with the MHGW to be a member of a multicast group, the MHGW sends it a "special key" to be used for receiving messages. The alternative to this is for the MHGW to keep track of all group members and unicast a copy of the message to each of them, which obviously wastes satellite downlink bandwidth.

The HHs that are attached on LANs would have an extra responsibility of forwarding multicast packets to and from other hosts attached on the LAN. Therefore, in addition, these HHs would run a proxy to enable them to act as a multicast router for the LAN.

The system architecture defined raises a lot of interesting issues to be addressed. Let us suppose that a group of multicast users are having an audio conference terrestrially (wireline) and in the middle of the conference, a user decides to multicast an image to others. The users equipped with hybrid terminals would be receiving this image through their satellite link instead. Therefore, it is important that certain performance issues, such as which link will act as a bottle neck to the conferencing, be carefully studied. Another interesting question is determining how many HHs can be served by a MHGW with minimal delay because there is the potential for congestion since all multicast traffic is routed through the MHGW.

Motivation for using CBT

The system architecture defined raises a lot of interesting issues to be addressed. Let us suppose that a group of multicast users are having an audio conference terrestrially (wireline) and in the middle of the conference, a user decides to multicast an image to others. The users equipped with hybrid terminals would be receiving this image through their satellite link instead. Therefore, it is important that certain performance issues, such as which link will act as a bottle neck to the conferencing, be carefully studied. Another interesting question is determining how many HHs can be served by a MHGW with minimal delay because there is the potential for congestion since all multicast traffic is routed through the MHGW.

In considering a routing protocol to be used for multicasting in hybrid networks, one has to carefully look at the issues unique to this type of network and make use of its asymmetric nature to min-

imize the overhead introduced by routing. The best approach would be to modify an existing routing protocol to accommodate hybrid networks since this would ensure changes are only made on gateways to HHs. As previously mentioned, the most predominant multicast routing protocol is DVMRP. However, the asymmetric nature of traffic in hybrid networks almost eliminates using any distance-vector-based protocol which only forwards multicast packets if they arrived over interfaces used to reach the source of a packet. Thus, if a HH is the source of a packet, the hybrid gateway would not forward it to other hosts since the packet arrived on a different interface (e.g. terrestrial) from the one used to reach the source (e.g. satellite interface). MOSPF was also eliminated since it uses a flooding based scheme and has high SPT computational costs, thus limiting its use on the Internet. PIM was not considered as an option because of the implementation complexities involved in switching between its two modes of operation. Even though implementation of CBT has not been completed, ongoing work shows that its merits make it well-suited for hybrid networks.

Non-Member Multicast Source: One of CBTs' attractive features is support of non-member sending, which makes it the best choice for resource discovery applications. Data driven protocols such as DVMRP and PIM dense mode are less suitable for such applications since a group forwarding state is established as data flows in all routers from point of source. On the other hand, routers in between a non-member sender and the corresponding CBT delivery tree incur no group-specific overhead for forwarding that sender's multicast data packets; these are encapsulated by the sender's local CBT router and unicast to one of the group's core routers. The core would then decapsulate packets and distribute them over the corresponding delivery tree.

Minimal Delay: The asymmetric nature of traffic has been a major motivating factor in the development of hybrid networks as a means of preserving wireline corporate Internet bandwidth for other outgoing traffic. In the case of satellite broadcast for incoming traffic, delay incurred at the satellite link could be significant. The CBT architecture that routes all multicast traffic towards the cores of the distribution tree suggests that by careful selection of cores, we can minimize delay incurred in CBT trees.

Scalability: Current multicast routing schemes such as DVMRP, MOSPF, PIM dense mode employ

some sort of source-based routing where a multicast tree is constructed per source per group. This type of architecture works well when multicast traffic is densely populated in a region. However, in hybrid networks that mostly span wide areas sparsely, CBT which was designed to suit low traffic distribution areas would work better since there is less protocol messaging overhead involved. Moreover, since only one shared tree is built per group, the number of entries in the CBT routing table is exactly the same as the number of groups thereby providing a considerable reduction in storage space required. It would also be easier to construct the FIB table since each group's members are attached to the same satellite interface.

Interoperability: The CBT operation mode which assumes a region is heterogeneous with routers using different protocols, as is typical of WANs, makes it possible for multicast packets to traverse regions that are not CBT capable. This facilitates Inter-Area routings and compliments the interoperability with other protocols. Already the interoperability of CBT with DVMRP has been defined in [3].

Routing Protocol Independence: Most of existing multicast routing protocols depend on the underlying unicast routing protocol used. For example, DVMRP is based on RIP while MOSPF only runs on networks running OSPF. Because of the spontaneity of applications of multicasting such as conferencing, a server multicasting video packets to hybrid hosts may belong to a network running a different routing protocol. Hence CBT which builds its multicast tree independent of unicast routing protocol would be at an advantage.

Protocol Specifications

For the HCBT architecture proposed, all routing of multicast packets to and from hybrid terminals is done through the MHGW. To make this possible, modifications would have to be done on both the HHs and the MHGW.

The HHs would run a modified version of IGMP to enable the MHGW to learn group membership. In addition, those HHs that act as routers for members on their LANs would have to run a proxy to enable them to act as a "semi-querier" for the LAN and forward membership information to the MHGW. Furthermore, these special HHs would be responsible for multicasting received packets to member hosts on their LAN (either through broadcasting, say on Ethernet, or some other multicasting scheme).

The MHGW has to be CBT capable in order to join the corresponding multicast trees on behalf of the HHs. As specified by CBT, the group joining process will be triggered by the receipt of an IGMP message for a multicast group. The MHGW would then send a join message towards the target core as specified in [4] for attachment to the multicast tree. After receiving an acknowledgment message, the HCBT module would include in its Forwarding Information Database (FIB), an entry corresponding to the tree joined. Since the IGMP message arrives over a different interface from the one where multicast packets have to be forwarded (the satellite interface), slight modifications have to be made to the way CBT operates to ensure that the correct entry is put in the FIB.

The elegance in the proposed architecture lies in its capability to do intelligent routing based on traffic type. To support this feature, the MHGW will have to implement a switching mechanism that routes high-data-rate packets through satellite and low-data-rate packets through terrestrial wireline links. In effect, this would be equivalent to maintaining two separate multicast delivery trees. A simple solution would be to have the MHGW encapsulate all low data rate packets and unicast them to the HHs but obviously this is resource wasteful.

When the MHGW receives a multicast packet, it would consult its multicast routing table to determine the interfaces out of which packets have to be forwarded. If the data rate warrants, it would forward packets to the satellite interface for broadcasting. The HHs would receive packets by listening to the channel for multicast packets sent using a scheme similar to Ethernet multicasting where a mapping is defined between an IP multicast address and the HHs' adapter addresses. Because broadcast packets would be available to all HHs, the MHGW would have to run some authentication scheme to allow only registered group members to receive packets. The authentication mechanism could be included in the IGMP messaging process so that once multicast trees are joined, all the necessary information to send and receive packets is available to HHs.

To establish a reliable multicast delivery mechanism that guarantees "at least once" delivery of multicast packets, MHGW would keep a copy of all packets until an acknowledgment is received from all HHs. Hence, the MHGW would have to keep track of all HHs members for each group. However, this deviates from traditional IP multicast schemes (IGMP) where multicast routers only keep track

of group membership information on their attached networks and not individual members of each group.

Group Membership Protocol

The IGMP used by multicast routers to learn about group membership information on their local subnet, is ill-suited for the satellite-terrestrial hybrid network considered in this paper because some of the assumptions made may no longer hold for this scenario. The IGMP specifically assumes that all hosts within a local subnetwork can hear each other and that routers need not keep track of individual members of each group. In our scenario, HHs form a virtual subnetwork with the MHGW as their gateway. However, HHs have no direct link with each other since the satellite link is uni-directional. Therefore, certain modifications have to be done to the IGMP before it can be used.

The IGMP specifies that a Querier router on the subnetwork periodically (about every 1 second) sends a general query to all hosts on their attached LAN to determine group membership information for each group with directly attached group members. When a host that is a member of the group hears the Query, it sets a random delay timer for each group of which it is a member. When a group's timer expires before another host's report is received, the host broadcasts a membership report on the local subnetwork. If a local host receives another host's report while it has a timer running, it stops its timer and suppresses the report that was about to be sent. In the hybrid network considered, the only logical choice for the Querier is the MHGW. However, if IGMP is used as specified, the HHs within the MHGW's logical subnet would not hear each others' group report since traffic to the MHGW is sent via a terrestrial link and hence would not be able to suppress their own reports. This would lead to an undesirable flooding effect of messages to the MHGW from a HH once a query is issued. The trivial solution would be to have the MHGW broadcast reports received from HHs on the satellite link so that other HHs could hear them. This would involve increasing the random timer delay to account for the time it takes for a report to reach the MHGW and be broadcast.

If reliable multicast delivery is desired and HHs are allowed to suppress their group membership reports, then the MHGW would not have information on the individual membership information of each group, and hence would not be able to guarantee delivery of packets to hosts. In this case, it would be better to remove the query option from IGMP and have all HHs send a membership report to the

MHGW when they join or leave a group. To cover the case of lost packets, the report should be duplicated if an acknowledgment is not received within a specified delay timer. This method would cause problems during startup or end of a multicast session when all HHs try to join or leave group because the MHGW would be flooded with group messages. Therefore this technique is only suitable for groups with a small number of HH members.

On the other hand, if reliability is not desired, the MHGW can still forward reports over the satellite link so that other HHs may suppress their reports. Query-Requests need not be sent since Leave-Reports would also be broadcast. Hence if a Leave-Report is heard by a HH for a group it is still a member of, it sends another Join-Report to the MHGW after its delay timer for that group expires before it receives a Join-Report from another HH.

HCBT Subsystems

Before proceeding with our design specifications, several simplifying assumptions are made that introduce some level of abstraction so that details not immediately essential are delayed until needed. As we proceed, our model would be validated to determine how close it is to the design requirements, and new subsystems added so that the whole abstraction process is re-iterated. We consider the special case of the HCBT architecture illustrated in Figure 2 where:

- there are only “stand-alone” HHs, i.e. HHs are not attached on LANs where they are responsible for routing multicast packets to other terrestrial members.
- there is no intelligent routing at each HH, i.e. there is no differentiation among the different traffic types. Hence, all traffic from HHs goes out terrestrially and all incoming traffic is routed on satellite.
- all HHs that are multicast sources only send low data rate traffic.
- there is only one hybrid gateway serving all HHs.

Multicast Hybrid Gateway Subsystem.

Supporting multicasting in the architecture shown above, requires implementing three new modules at the hybrid gateway; an IGMP module, a Multicast Database (MDB) module, and a Hybrid CBT router module (HCBT).

The IGMP module would run a modified version of IGMP and would be responsible for keeping track

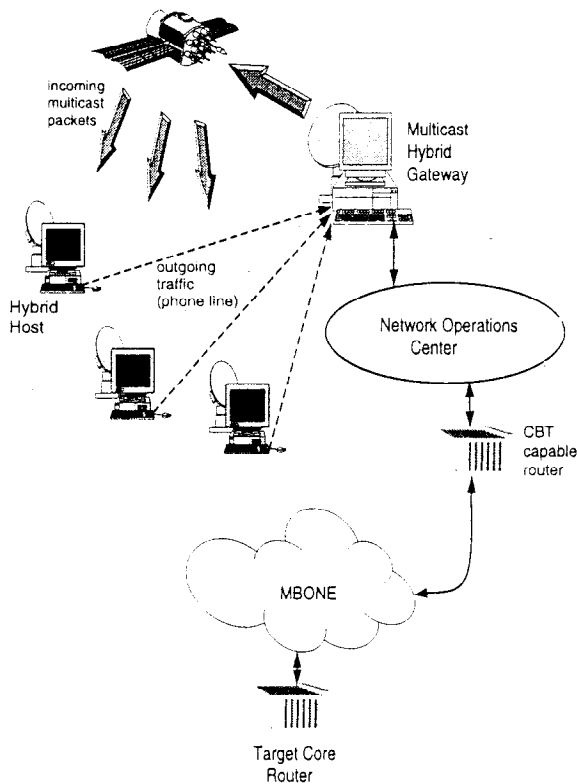


Figure 2: Simplified HCBT architecture.

of group membership information of the HHs. It would query the HHs to determine which HHs are members of multicast groups. When it receives a group membership report from a HH, it would query the MDB to determine whether it has already joined the corresponding tree for that group. Once the corresponding tree has been joined, it would run an authentication process to authorize HHs to receive multicast packets.

The MDB module would maintain and manage a local database of trees joined by the MHGW. It would consist of entries denoting which multicast trees have been joined. Furthermore, for reliable “at least once” delivery of packets, this table will keep track of all hosts that are members of each group. The MHGW will keep a copy of all packets until they are acknowledged by all HHs in the group. It is necessary to separate this module from the HCBT module which contains a FIB with the same information because as we drop some of the assumptions made, it may be necessary to maintain more state information.

The HCBT module will run a CBT router function that enables the MHGW to join multicast trees on behalf of the HHs. It will be responsible for sending join messages towards the core of the tree and

routing multicast packets to and from HHs.

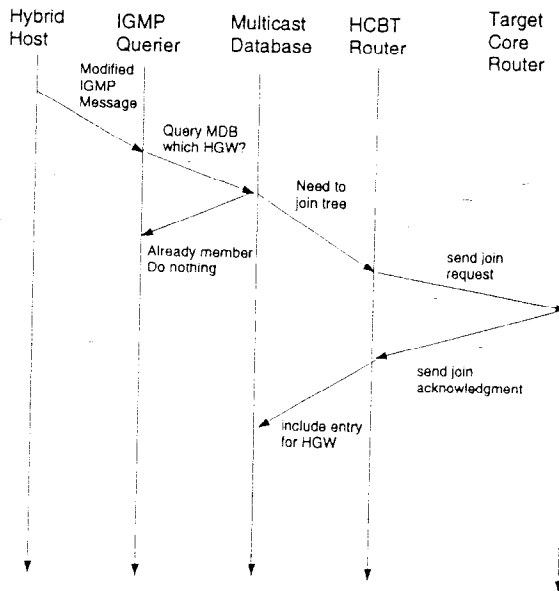


Figure 3: HCBT Tree Joining Process.

On receipt of a IGMP report, the IGMP module will consult the MDB module to determine if it has already joined the corresponding tree of that group for the HGW responsible for the HH. If not, it will inform the HCBT router module of its intention to join the tree. The HCBT module would then send a join message towards the target core as specified in [4] for attachment to the multicast tree. After receiving an acknowledgment message, the HCBT module would include in its Forwarding Information Database (FIB), an entry corresponding to the tree joined and an entry will be added to the MDB specifying the HH as belonging to that group. It should be noted that the IGMP report arrives over a different interface than one where multicast packets are to be forwarded. Therefore, it would be necessary to modify CBT to include the correct interface to which the packet has to be sent. A timing diagram for the tree joining process is shown in Figure 3.

When the HCBT module receives a multicast packet, it would use the FIB information to forward it over the satellite interface if there are HH group members. It would also encapsulate a copy of the packet and send it CBT mode to other interfaces as specified in the FIB since other CBT capable routers on the MBONE could join the delivery tree through it. Because the MHGW will need to keep a copy of all multicast packets until acknowledgment is received from all HH group members, a good buffering

management scheme has to be devised.

Security has been of growing concern especially for multicast applications because it becomes relatively more difficult to distribute group keys to each of the group's receivers than to authenticate a session of a single source and destination. A scalable multicast distribution key has been described in [5] which uses CBT to establish secure multicast groups. The solution allows multicast routers to become Group Key Distribution Centers (GKDCs) after receiving a CBT Join ACK to become part of a multicast tree. Thereafter, the GKDCs are responsible for distributing group keys and key encrypting keys to group members on attached sub-networks. Therefore, we could have the MHGW act as the GKDCS for all HH group members and provide them with authentication keys. Because the keys would be broadcast on satellite, maintaining confidentiality would be difficult and extra precautions such as encryption techniques would need to be taken to ensure that only HH members receive packets.

Hybrid Host Subsystem

There are several functions that need to be implemented in the HH for it to support multicasting. The HH must run an IGMP module that allows it to listen for IGMP queries on its satellite interface and respond (send group reports) using its terrestrial interface. The HH has to be level 2 compliant with IGMP to be able to both send and receive multicast packets. A mapping has to be defined between its satellite IP address and its adapter card to be able to forward packets destined for it up the TCP/IP stack. The host must be able to cache the keys sent to it by the MHGW during authentication so that it could be used for future multicast traffic.

When a HH wishes to be a member of a group, it sends a group membership report on its terrestrial link to the MHGW. The MHGW will then construct a delivery tree if needed, add the HH in the MDB, and then unicast an authentication key to the HH. The HH then listens on the satellite interface for packets destined for that group. When it receives packets, it sends acknowledgments to the MHGW via its terrestrial link.

It is important to note that as HHs join or leave groups, new keys may be broadcast by the MHGW. Therefore, a process running on the HH would need to renew keys for the HH. This process may need to periodically compare the checksum of its current key to that broadcast on satellite. If they are different, then it should trigger a request for new keys from

the MHGW to be sent via unicast to prevent other non-member HHs from receiving key.

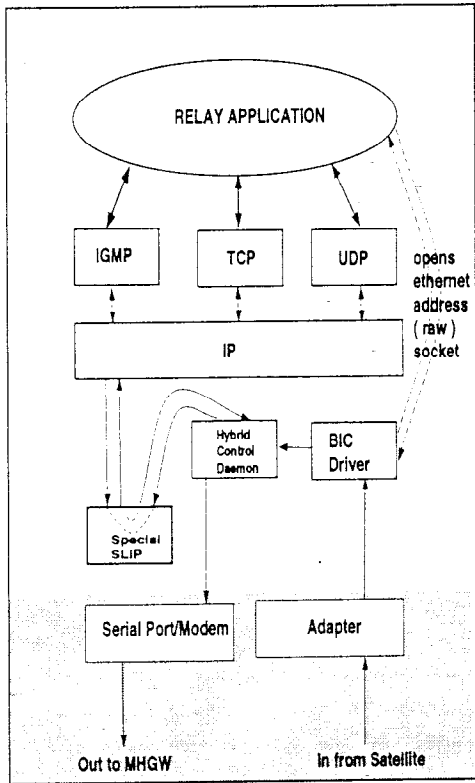


Figure 4: Flow control in the Hybrid Host.

The multicast packets are broadcast to the HHs similar to the way multicast packets are sent to hosts attached to an Ethernet LAN. Hence, one way of receiving the multicast packets would be to make the HH physical interface (adapter) act like a single Ethernet link for the sake of carrying a multicast address. To achieve this, a socket has to be opened through which the relay application running on top of TCP/IP can receive multicast packets. In the only Hybrid Internet Access product, DirecPCTM TurboInternetTM developed by the University of Maryland and Hughes Network Systems for Windows using the architecture described in [2], there is a "special" SLIP driver in the HH that communicates with the two physical networks to make the TCP/IP package believe that is connected to an Ethernet card when it is actually connected to a satellite dish and modem. A hybrid control daemon manages the flow of data between this special driver and a BIC driver. The BIC driver does all the call handling by scanning all packets transmitted over the satellite channel for one with a header corresponding to the IP address of the satellite interface. In addition, the BIC driver performs some error

detection and correction on the packet and buffers the received packet before passing it to the special driver. Similarly for our system, the BIC driver call handler can be modified to support raw sockets and filter out UDP or TCP packets destined for multicast groups that relay applications have joined. Figure 4 shows the flow control in the HH subsystem.

When a HH is the source of multicast traffic, packets are encapsulated through the terrestrial tunnel to the MHGW. At the MHGW, they are decapsulated revealing their true multicast address destination and routed to other group members according to the distribution tree of the group. If group membership includes other HHs, the packets are also broadcast on satellite. Hence, additional filtering has to be done by the BIC driver to discard packets with HH address as the source.

Multiple MHGWs

With all multicast traffic to and from HHs routed through the MHGW, it is inevitable that there would be traffic congestion problems as the number of HH group members grow. Fortunately, provision has already been made in the current Hybrid Internet Access architecture to support multiple hybrid gateways (HGW) where each subnet of HHs are represented by different hybrid gateways. Packets to and from HH are routed first to the hybrid LAN gateway which broadcast it on the Ethernet LAN connecting all HGWs. The HGW takes up all routing tasks for all packets to and from HHs on its subnetwork.

To project this scenario to the multicasting case, the MHGW could be implemented at the HGW with one of the HGWs designated as the IGMP querier (DR-MHGW) responsible for joining multicast trees. When a HGW receives an IGMP group report from a HH, it will include the HH in its MDB module and broadcast a copy of the report on the LAN. The IGMP Querier will pick up the report and consult its FIB to determine if it has already joined the corresponding delivery tree. If not, it will trigger its HCBT module to send a join message towards the core of the tree. Similarly, when a multicast packet arrives, the DR-MHGW will broadcast it on the LAN and forward a copy over all interfaces (including satellite interfaces) as dictated by its FIB. All MHGWs with group members will buffer a copy until acknowledgment are received from all group members on its subnet. HHs with errored or missed packets will request their MHGW for retransmission of packets. This will significantly reduce the buffering management complexity at the MHGW. Also, since the DR-MHGW will be the only one attached

to the delivery tree, all other MHGWs need not run a HCBT module.

Core Selection and Migration

A major problem of CBTs is that shared trees built incur a high traffic concentration on the shared path. Furthermore, the tree built is not always the shortest path tree. It is believed that strategic core placement would help eliminate these problems completely. This would require developing core migration techniques that allow the dynamic transition from an initial CBT tree constructed around a pre-configured set of cores to another tree with different set of cores. The authors of CBT have not completely solved the core placement or core advertisement problem, but have defined a dynamic source migration mechanism in the appendix of [4]. This strategy allows a CBT tree to dynamically reconfigure itself around the source's local CBT router to emulate a shortest path tree.

The network architecture assumed for this solution routes all multicast traffic to and from hybrid hosts through the MHGW. A lot of research has been done on determining core selection methods for multicast routing and how it affects performance in [6], [7], and [8]. Specifically, three performance criteria, bandwidth, delay, and traffic concentration, are considered to investigate the effect core choice has on them. In their evaluation, the authors of [6] considered instances of three different types of scenarios, reflecting distributions and numbers of sources and receivers. An "All Receivers Sources" scenario modeled applications such as video conferencing where receivers are distributed randomly throughout the network and a user is both a source and a receiver. "Single Source, Distributed Receivers" covered applications such as a video broadcast of a lecture or meeting where most members are receivers. Finally, "Localized Receivers" modeled distributed resource discovery applications where sources (clients) are randomly distributed and request information from receivers (servers) via multicast. In addition, core selection methods were classified into one of the following categories in increasing order of information required about the network: arbitrary, random, topology-based, or group-based, where arbitrary requires no information about the network and group-based requires information on both network topology and location of nodes. From the studies in [6] and [7] it was established that the best performance - maximum bandwidth improvement and minimum delay degradation is obtained from a core chosen based on both the network topology and location of nodes (receivers), although the

improvement was not significant for certain distribution scenarios. Furthermore, it was established that the core should be the center of the portion of the shortest path trees that spans all group members and sources.

Traffic distribution in hybrid networks can be best modeled by a "single source, distributed receivers" since it was developed based on the assumption that traffic is asymmetric with most users receiving much more than they are sending. Hence, multicast applications in hybrid networks would mostly be of video broadcasting nature. Since in the HCBT architecture described, the MHGW is responsible for routing of all multicast packets, it acts as a source to the HCs and the rest of the multicast network is hidden from the hosts. On the other hand, CBT mode allows users to unicast all multicast packets towards the core of the group using encapsulation. Once the packet reaches the core, it is decapsulated and forwarded out of all outgoing interfaces. Therefore, to emulate the shortest path tree and minimize delay for HCs, it makes sense to select the MHGW as a core for all groups joined. The MHGW could be configured to be a core for all groups joined. Alternatively, since dynamic core migration has been specified in [4], the MHGW could be configured to trigger a core migration to itself after it joins a tree for a group. The disadvantage of making the MHGW a core is that additional processing power may be required to process CBT protocol messaging. Introducing multiple cores would keep this to a minimum and would also reduce the traffic concentration problems inherent in shared links.

Performance Metrics of Multicast Protocols

To evaluate performance of a multicast protocol, several indicators are used to see how well the protocol performs under different scenarios. For dynamic multicast routing, it is important to determine the latency involved in joining the multicast group, from the time the request is sent by a host to the time the first multicast packet is received. It is desirable that this latency be kept to a minimum.

However, the main performance metric used is the time it takes for each member of the group to receive packets sent, i.e. transfer time. The transfer time depends on the throughput of the multicast session which in turn depends on both the available bandwidth and the probability of packet loss. Thus reducing the transfer time involves using a congestion control scheme to ensure that available bandwidth is not exceeded, and at the same time controlling packet loss in the delivery path. In order to control

packet loss, it is essential to first understand the underlying process and identify the source of losses so that the appropriate error control measures can be taken. Packet loss can be due to transmission or switching errors and buffer overflows at routers and hosts. A lot of studies have been done to determine packet loss correlation in multicast networks, and in [9], it was shown that losses on the MBONE are in fact “temporally” correlated, i.e., most losses occur at receivers and routers and not on links.

Topology of the multicast distribution tree also affects the packet loss characteristics and consequently, the transfer time. Mishra et. al. in [10], a study done to evaluate the effects of topology on reliable multicast routing, conclude that as a general rule, a topology which increases “fanout” of the distribution tree performs better in the asymptotic case.

Traffic concentration on the links in the distribution tree is also used as a performance indicator. Multicast routing protocols that construct shared trees experience a higher concentration when compared to source-based trees [7]. Path cost in terms of the number of links transversed when delivering a packet to all group members also gives a good estimate of the bandwidth used. Other metrics used include overhead traffic of protocol, scalability, and protocol algorithmic complexity.

Because of the high-delay satellite link involved, the most important, metric for the protocol proposed in this paper is the transfer delay in delivering multicast packets to all hybrid hosts since the remainder of the delivery path is terrestrial. We will assume when estimating the delay in subsequent sections that packet losses on the satellite link are insignificant and that most losses occurs at the hybrid host receivers due to overflow of buffers. This is actually quite close to reality because most current and future satellite systems incorporate strong Forward Error Correction (FEC) protection so that up to a certain signal-to-noise ratio (SNR), the satellite channel can be modeled as an on-off switch.

Performance Evaluation

Simulation techniques were used to verify and validate our design instead of an actual prototype because it gives us a quicker methodology for evaluating performance and more flexibility with modifying design parameters once the model is built.

One of the motivating factors of supporting multicasting in hybrid satellite terrestrial networks is to allow companies with limited gateway bandwidth to engage in high-bandwidth multicast applications.

Therefore, a simulation was done to evaluate the bandwidth savings in multicasting over a hybrid network over traditional terrestrial wireline multicasting. Since some of these applications may have a time constraint on the transmission time, further studies were done to find the effects of high-delay satellite link on a multicast session. Finally, we investigate the use of traffic type and size in deciding whether it is advantageous to route multicast packets through the satellite or not.

Simulation Model All simulations were done using Optimized Network Engineering Tools (OPNET), a comprehensive engineering system capable of simulating communications networks with detailed modeling and performance analysis. OPNET features include: graphical specification of models; a dynamic, event-scheduled Simulation Kernel; integrated data analysis tools; and hierarchical, object-based modeling. OPNET’s hierarchical modeling structure accommodates special problems such as distributed algorithm development.

Two OPNET network models were built to simulate two environments: one in which multicasting is done terrestrially, and another in which all multicast packets are routed through a hybrid network over satellite to HH group members. The scenario under consideration is “single-source distributed receivers” typical of applications such as video lecture broadcast, with listeners (HHs) allowed to send only low-data-rate traffic to group since they may have limited uplink bandwidth. It is assumed that there are additional group members in the terrestrial network in the vicinity of the source. The same number of hops are used between the HHs and source in both environments. Standard OPNET TCP/IP processes were modified when appropriate to build the simulation model. For the hybrid network model, some of the processes used in [11], a simulation of a hybrid network, were also modified to support routing of multicast packets. The same network topology assumed in the analysis section was used in the simulation. The simulation parameters used are given in Table 1 and their values are given in the Appendix .

Parameter	Description
Traffic Type	Rate of packet generation for the multicast source (Server) and hybrid host (Source) to model high-data-rate and low-data-rate type traffic.
Traffic Size	Size of transaction requested to model “short-length” and “bulky” traffic.
Service Rate	Service rate of packets destined for HHs at MHCW and Corporate Gateway buffers

Table 1: Simulation Parameters.

Simulation Results

Figure 5 compares the corporate link utilization for each of the two environments considered: multicasting in a terrestrial network (wireline) versus a hybrid network with a satellite downlink. As expected, introducing hybrid terminals in a corporate LAN preserves corporate gateway bandwidth for other traffic. Since all incoming packets for hybrid case are routed through satellite, available bandwidth for other traffic is more than twice the bandwidth available when incoming multicast packets are routed terrestrially.




Figure 5: Corporate Link Utilization Comparison

Figure 6 shows the round-trip-time (RTT) of packets for both hybrid host and terrestrial host group members. As the session length increases, the performance of the terrestrial network considerably declines while that of the hybrid network remains stable. In the terrestrial network, the RTT of packets initially slows down but increases quickly because the corporate gateway is slowed down by the additional packets to be processed. Thus, more packets are transferred as indicated by the increase in throughput at the HHs observed in Figure 7 for the hybrid network case.

Figure 8 shows the effect of traffic type and size on the transfer time of multicast packets. From the figure, it can be seen that the delay is less in the terrestrial network for "short-length" or low-data-rate sessions (see Table 4). Thus under such a scenario, it is not advantageous to route multicast packets through satellite. This clearly demonstrates the need for an intelligent routing scheme at the MHGW, that would allow only high-data-rate or "bulky" traffic to be routed via satellite.




Figure 6: Round-Trip-Time Comparison




Figure 7: Received Segment Sequence Number of HH Packets

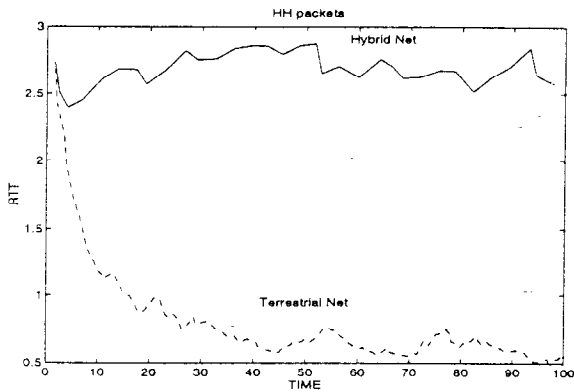


Figure 8: Effect of traffic type on RTT

Figure 9 shows a comparison of the throughput for different MHGW buffer sizes (see Table 5). From the figure, it can be seen that the achievable throughput is higher when a large buffer size is used since this allows the source to send larger amounts of data by advertising a larger window size.




Figure 9: Effect of buffer size on Throughput

Conclusions & Further Research

The asymmetric nature of traffic in most networks, as evident in the Internet, is shifting current networking technology trends more towards the development of hybrid networks. Emerging group communication applications such as video broadcasting and teleconferencing that demand high bandwidth have driven the development of multicast protocols on the MBONE. Thus hybrid terminals can be deployed for receiving IP multicast packets

as a means of preventing congestion on the Internet backbone and preserving Corporate gateway bandwidth.

The design described here entailed implementing a Multicast Hybrid Gateway that would be responsible for keeping track of group membership of hybrid hosts, join delivery trees of multicast sessions, and route multicast packets to hybrid host group members. In addition, this subsystem could also handle authentication or intelligent routing schemes. The IGMP protocol was modified to emulate a virtual link between hybrid hosts and the gateway. Also, the multicast routing protocol employed in the terrestrial part of the network was assumed to be CBT and thus appropriate changes were done in the CBT module of the hybrid gateway.

Simulation techniques were used to demonstrate the bandwidth savings in multicasting in hybrid networks over terrestrial counterpart. Our simulation results agreed with our mathematical analysis (described elsewhere [12] that it is only advantageous to route packets over satellite if there is high-data-rate or "bulky" traffic, thus indicating a need to implement intelligent routing at the gateway. Finally, our results also agree with analytic studies that show that an increase in buffer size also improves performance [12].

Appendix: Simulation Parameter Values

Table 2 and 3 shown below present the important parameter values used in the simulation.

Parameter	Value
Source App. Interarrival Rate	0.001secs/pk
Packet Size	9 Kbits
Hybrid Host App. Interarrival Rate	0.5 secs/pk
HH RCV Buffer Size	45 Kbytes
MHGW RCV Buffer Size	64 Kbytes
Source RCV Buffer Size	4 Kbytes
Modem Speed	28.8 Kbits/sec

Table 2: Parameter Values for Hybrid Network.

Parameter	Value
Source App. Interarrival Rate	0.001secs/pk
Packet Size	9 Kbits
Hybrid Host App. Interarrival Rate	0.5 secs/pk
HH RCV Buffer Size	45 Kbytes
Source RCV Buffer Size	4 Kbytes

Table 3: Parameter Values for Terrestrial Network.

Parameter	Value
Source App. Interarrival Rate	0.5secs/pk
Hybrid Host App. Interarrival Rate	0.5 secs/pk

Table 4: Parameter Values for Low-Data-Rate Traffic.

Parameter	Value
HH RCV Buffer Size	45 Kbytes
MHW RCV Buffer Size	1000 Kbytes
Source RCV Buffer Size	4 Kbytes
Source App. Interarrival Rate	0.001secs/pk
Packet Size	9 Kbits
Hybrid Host App. Interarrival Rate	0.5 secs/pk

Table 5: Parameter Values To Show Buffer Size Effect.

References

- [1] V. Arora, J.S. Baras, Dillon, A.D. Falk, and N. Suphasindhu. "Hybrid Internet Access". *Proceedings of the AIP Conference on NASA Centers for Commercial Development of Space, Albuquerque, NM*, pages 69-74, 1995.
- [2] V. Arora, N. Suphasindhu, J.S. Baras, and D. Dillon. "Asymmetric Internet Access Over Satellite-Terrestrial Networks". *AIAA Proceedings of the 16th International Communications Satellite Systems Conference and Exhibit*, pages 476-482, 1995.
- [3] A. Ballardie. "Core Based Trees (CBT) Multicast Interoperability," *INTERNET DRAFT*. University College London, April 1996.
- [4] A. Ballardie, S. Reeve, and N.Jain. "Core Based Trees (CBT) Multicast Protocol Specification," *INTERNET DRAFT*. University College London, Bay Networks, July 1996.
- [5] A. Ballardie. "Scalable Multicast Key Distribution," *RFC 1949*. University College London, May 1996.
- [6] K. Calvert, E. Zegura, and M. Donahoo. "Core Selection Methods for Multicast Routing," *Technical Report GIT-CC -95/15*. Georgia Institute of Technology, June 1995.
- [7] K. Calvert, E. Zegura, and M. Donahoo. "A Comparison of Two Practical Multicast Routing Schemes," *Technical Report GIT-CC-94/25*. Georgia Institute of Technology, February 1994.
- [8] L. Wei and D. Estrin. "The Trade-off of Multicast Trees and Algorithms," *Technical Report USC-CS-93-560*. University of Southern California, June 1993.
- [9] M. Yajnik, J. Kurose, and D. Towsley. "Packet Loss Correlation in the MBONE Multicast Network". *IEEE Global Internet Conference*, December 1996
- [10] P. Bhagwat, P. Mishra, and S. Tripathi. "Effect of Topology on Performance of Reliable Multicast Communication". *Proceedings of INFO-COM 94*, pages 602-609, June 1994.
- [11] A. Gaid and J.S. Baras. "Simulation of a Hybrid Network in Order to Enhance Performance of Hybrid Internet Service," *Technical Report*. University of Maryland CSHCN, December 1996.
- [12] I. Secka. "IP Multicasting in Hybrid Networks," *Technical Report CSHCN M.S. 97-2/ISR-M.S. 97-5*. M.S. Thesis, Systems Engineering, University of Maryland, Institute for Systems Research and Center for Satellite and Hybrid Communication Networks, May 1997.