Issues on the IP Multicast Service Behaviour over the Next-Generation Satellite-Terrestrial Hybrid Networks

Fethi Filali, Walid Dabbous
INRIA, 2004 Route des Lucioles, BP-93
06902 Sophia-Antipolis Cedex, France
e-mail: {filali,dabbous}@sophia.inria.fr

Abstract

Our main focus here is to study the IP multicast service behaviour over the next-generation of satellite-terrestrial hybrid networks for which we develop a general architecture and we enumerate their main characteristics. We first focus on the IGMP feedback implosion problem and we present the exponential feedback raise that able to provide sufficiently stable expectation values across a large of group sizes. Then, we propose and compare different approaches that can be used to enable the multicast in the on-board satellite switch. We show that the approach based on maintaining a Multicast Beams Table (MBT) containing the list of spot beams concerned by each multicast group is the most suitable approach because it allows an efficient and transparent integration of satellite links in the Internet. In fact, the on-board switch monitors the IGMP reports and queries sent over the satellite link and update the MBT entries.

In the second part, we turn our attention to multicast routing protocols. We present some undiserable behaviour of DVMRP, PIM-DM, and PM-SM. For DVMRP and PIM-DM, we identify some configurations where the satellite receivers may receive duplicated packets and we propose a method to overcome this problem. We then develop a configuration policy of PIM-SM in hybrid networks concerning the choice of the list of Rendezvous Point (RPs) and the switching from the RP-routed tree to the shortest path tree.

Keywords: *satellite-terrestrial hybrid networks, on-board switching, multiple spot-beams, IGMP feedback implosion problem, multicast routing protocols.*

1 Introduction

Multicasting allows us to send data packet to multiple sites at the same time. The key here is the ability to send one message to one or more nodes in a single operation. This provides a tremendous amount of savings in bandwidth when compared to traditional unicast transmission which sends messages to multiple nodes through replication of the message to each node. Besides the performance improvement over unicast transmission, multicast allows the construction of truly distributed applications. Indeed, multicast allows application developers "to add more functionality without significantly impacting the network". Therefore, multicast-based applications and services will play an important role in the future of the Internet as continued multicast deployment encourages their use and development.

Parallelling to the evolution of the Internet services, the Internet infrastructure is integrating several types of wired and wireless communication links. One of the major components of this infrastructure is the satellite links [2]. In fact, even in the past three decades, satellites have played a pivotal role in global telecommunications. It is anticipated that they will play a complementary role in the so-called information superhighway, a term referring to an infrastructure consisting of networks linking homes, business, government, and institutions to a wide range of interactive multimedia services. Potential applications include teleconferencing, tele-learning, high resolution image transfer, home banking and shopping, video on demand, TV/radio/newspaper/data broadcasting. The nature of these services require the adoption of broadband transmission at the T1 rate and beyond. To support these high bandwidth applications, it is anticipated the next-generation satellite communications systems will differ from the traditional systems by including intelligent functions in the on-board satellite, the use of ka-band and V-band, and the use of the spot beams technology. All theses aspects may handle different question in the manner to use the Internet protocols over satellite.

Since the Internet protocols have been designed without taking into account the inherent characteristics of the physical support, the harmonious and the efficient integration of the next-generation satellite systems into the Internet requires the study and the adaptation of these protocols. The dynamic unicast routing problem over unidirectional links, and in particular satellite links, has been investigated in the IETF UDLR Working Group. In fact, they have proposed LLTM (Link Layer Tunneling Mechanism) [7] which performs well by capturing the dependence of the network layer on the bidirectionnality of the communication link. It uses a Dynamic Tunnel Configuration Protocol (DTCP) that provides a means for satellite receivers to dynamically discover the presence of feeds (satellite uplink stations) and to maintain a list of operational tunnel end-points[1]. Feeds periodically announce their

[1] LLTM uses GRE (Generic Routing Encapsulation) [9] to tunnel packets from the satellite receivers to feeds via the terrestrial tunnels config-

tunnel end-point addresses over the unidirectional link. Receivers listen to these announcements and maintain a list of tunnel end-points.

While the dynamic unicast routing over unidirectional links has been resolved, there are several other open issues related to the integration of satellite links in the Internet that remain without solutions or even have not been already addressed such as the scalability of LLTM mechanism, the multicast routing, and the reliable multicast transfer over satellites. In this paper, we study the behaviour of the IP multicast standard model over the next-generation of the hybrid satellite-terrestrial networks. We discuss the problems of the deployment of satellite links in the Internet for multicast delivery and we present some short-term as well as long-term solutions to overcome these problems. These solutions concerning the Internet Group Membership Protocol (IGMP) [5] feedback implosion problem and the construction of an efficient multicast tree that profits from the broadcast capability of the satellite. We present an algorithm that use an exponential feedback raise and which able to provide sufficiently stable expectation values across a large of group sizes. We then propose three approaches that can be used to enable the multicast in the on-board satellite switch and we enumerate the advantages and the drawbacks of each approach. We show that the approach based on maintaining a Multicast Beams Table (MBT) by the on-board satellite switch containing the list of spot beams concerned by each multicast group is the most suitable approach because it allows an efficient and transparent integration of satellite links in the Internet.

The second part of this paper deals with the multicast routing protocols. We examine source-based protocol such as DVMRP [15], and PIM-DM [6] and from those using shared tree we focus on PIM-SM [8]. For each multicast routing protocols we pick up its possible undesirable behaviour in satellite environment and the most suitable tuning method. For DVMRP and PIM-DM, we identify some configuration where the satellite receivers can receive duplicated packets and we propose a method to overcome this problem. For PIM-SM we propose a configuration policy that has two main benefits. First, it builds a delivery tree where members receive data from the broadcast satellite downlink either using a RP-routed tree or a shortest path tree. Second, it minimizes the multicast traffic load in the terrestrial network so as the safe terrestrial bandwidth will be used by unicast applications or, in general, by applications having requirements that can not guaranteed by the satellite link.

The remainder of this paper is organised as follows. The next Section presents the architecture of the next-generation satellite-terrestrial hybrid networks and their main characteristics. In Section 3, we discuss the behaviour of the IGMP in such type of networks where the multicast receivers may have the capability to receive data from the satellite link.

ured by DTCP.

We propose and compare three approaches that can be used to enable the multicast in the on-board satellite switch. We study the behaviour of the second part of the multicast service model that is the multicast routing protocols and in particular DVMRP, PIM-DM, and PIM-SM in Section 4, Section 5, and Section 6, respectively. We conclude this paper in Section 7 and we give some other open issues concerning the deployment of next-generation satellite links in the Internet for multicast delivery.

2 Next-Generation Satellite-Terrestrial Hybrid Networks Architecture

Prior to studying the behaviour of the IP multicast standard model over the satellite-terrestrial hybrid networks, a general network architecture of this type of networks must be specified. We mean by hybrid network a terrestrial network connected to a GEO satellite with *On-Board Switching* (OBS) capability offering a high bandwidth shared by satellite uplinks and a broadcast downlinks toward terrestrial receivers. A satellite uplink is an unidirectional or bidirectional wireless communication link which connect two nodes (router, host, etc.) via the GEO satellite.

In the remainder of this paper we use the following terminology [7]:

- **UniDirectional Link (UDL):** a one way broadcast satellite transmission link.
- **Receiver:** a router or a host that has a receive-only connectivity to an UDL.
- **Send-Only Feed:** a router that has a send-only connectivity to an UDL.
- **Receiver-Capable Feed:** a router that has a send-and-receive connectivity to an UDL.
- **Feed:** a send-only or a receive capable feed.
- **Node:** a receiver or a feed.

The Internet today is an inter-connection of Autonomous Systems (AS) (called also domains). An AS is a connected group of one or more IP prefixes run by one or more network operators which has a **single** and **clearly defined** routing policy [13]. In this section, we describe an architecture for one or more terrestrial-satellite hybrid ASs composed of terrestrial and next-generation satellite communications links.

In the satellite-terrestrial hybrid network architecture shown in Figure 1, we assume that the satellite networks are capable of providing multimedia (e.g. voice, video, data) services to the ground user. Most of such networks in operation today or planned for deployment in the nearest future are characterised by the support of:

- **On-Board Switching (OBS):** OBS provides a method of network switching on the satellite. It can be used to actively control signal routing. To cope with high network transfer rates, the switching of packets from an uplink to a different downlink spot beam needs to be done at a very high speed. There are two type of satellite switched network implementation: (1) Fully Switched where the satellite does all the processing and there is no or very little ground control. (2) IF or RF switching that involves electronically controlled RF/IF switches, which can be configured on a near-real-time basis via Ground Control.

- **Spot-beams technology:** instead of covering the whole footprint of a satellite by a global beam, the beam is divided into a number of spot beams. The benefits of spot beams are twofold: a) the power requirements of user terminals are reduced, thereby permitting the use of smaller antennas in the ground segment and reducing cost; and b) the frequency between beams can be re-used, thereby increasing the capacity of the space segment.

Figure 2 shows a general abstracted topology of a hybrid network. We denote R^i_j to be the satellite receiver number j that belongs to the spot beam number i containing n_i satellite receivers and F_i and SOF_i to be the feed number i and the send-only feed number i, respectively.

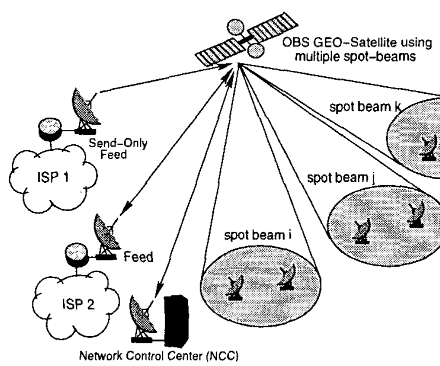

Figure 1: satellite-terrestrial hybrid network architecture characterized by the support of multiple spot-beams and on-board switching technologies

Each spot beam can be emulated to a broadcast unidirectional communication link that connects feeds, send-only feeds, and receivers belonging to this spot beam. Feeds and send-only feeds can send data to all spot beams or to a specified spot beam while every receiver has the capability to receive only the packets forwarded by the on-board satellite processor to its spot beam. Each receiver is also connected to the Internet via a dial-up link (e.g; PPP connection) and so it has two IP addresses, one assigned by the unidirectional Internet access provider (e.g; a satellite-based access provider) and the second one is that of the bidirectional connection assigned by the terrestrial Internet access provider.

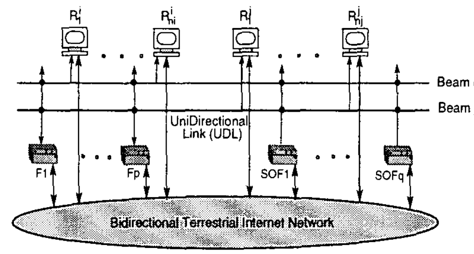

Figure 2: satellite-terrestrial hybrid network abstracted architecture. Each spot beam is emulated by a broadcast communication link. Feeds and send-only feeds can send data to each spot-beam, however receivers can receive only data forwarded to their spot beam.

Considering that the satellite receivers do not have the capability to send data via the UDL, we assume that feeds and receivers implement the Link Layer Tunnelling Mechanism (LLTM) as it is described in [7]. This mechanism allows nodes which are directly connected by a unidirectional link (feeds and receivers) to send datagrams as if they were connected to a bidirectional link.

Thus, users requests (and later acknowledgements as the session progresses) may be sent via the terrestrial link to an Internet host while data may be received on the UDL after it has been forwarded by one of the feeds.

3 IGMP over Satellite

3.1 IGMP Overview

IGMP is an integral part of IP that is used by IP hosts to report their host group memberships to any immediately neighbouring multicast router [5][2]. Hosts inform their local router of their intent to receive transmissions attached to a specific multicast group. The router would then periodically query the LAN to determine if group members are still active. Based on the group membership information learned from IGMP, the router joins a multicast delivery tree for each group that determines routes where multicast traffic has to forwarded over.

If no reports are received for a particular group after some number of queries, the router assumes that there are no group members for that group and prunes itself from the delivery tree of the group. To guarantee that a host will receive multicast traffic if it is the first to join a group, rather than wait for a query, it immediately transmits a report for that group when it joins a new group.

[2]IGMP V2 [10] and IGMP V3 [4] enhance and add additional features to Version 1. They specifically define a procedure for the selection of a multicast querier and add the support of Group-Source report message so that a host can elect to receive traffic format specific source.

In the next section, we present and discuss different approaches that can be used to enable the multicast in the on-board satellite switch.

3.2 Enabling multicast in the on-board satellite switch

Future generation satellites should be able to offer services that are not limited to existing mainstream C- or Ku- Band transparent satellites. As discussed in Section 2, the next generation broadband systems are expected to employ regenerative satellites rather than the current transparent satellites. Such satellites will also operate at higher frequencies (e.g. the 20/30 GHz, Ka-Band) and may be expected to employ spot beams, rather than providing continental coverage. On-board processing (switching) will direct packets to each appropriate downlink spot beams. This will enable multiple uplink terminals (at different locations) to serve as feeds to the multicast downlinks, thus providing (effectively) a space-bourne multicast overlay over the existing terrestrial Internet. This approach is well suited for unicast tranmission, which emulates the use of Internet switching in LANs. However, in order to support efficient multicast, the on-board switch needs to be multicast enabled. We propose three approaches to enable the multicast in the on-board satellite switch:

First, have a hub multicast router to configure a multicast registration table in the down-link satellite switch (e.g. on-board processing switch, either IP, ATM, or DVB). One the appropriate list of multicast address has been configured, the satellite should forward only the multicast packets to the spot beams that have at least one receiver registered in the registration table. This solution is not complex and it is easy to implement without a big added cost and there is no need to activate IGMP in the satellite links. However, it is a static solution and it assumes that the satellite access provider knows in advance the list of multicast groups for each spot beam and so that it is more suitable for configurations where the satellite receivers are themself the end-users.

Secondly, allow the on-board processor to observe group membership, by monitoring IGMP queries and reports and then constructing a list of which spot beams are associated (registered) with which multicast groups. In this scenario, each feed and each receiver enable the IGMP protocol in its satellite interface and IGMP queries and reports will be sent in the broadcast-emulated networks of each spot beam connecting feeds, send-only feeds, and receivers. The IGMP reports sent by a receiver via the UDLR tunnel to its default feed [7], will be broadcasted via the satellite uplink to all receivers belonging to the same spot-beam. When observing an IGMP report, the on-board processor should:

- identify the type of the report (joining a group, joining a source, leaving a group, leaving a source, etc.),
- identify the spot-beam to which belongs the sender of the IGMP report,
- and update the Multicast Beams Table (MBT) which can look like that of Table 1.

Spot beam number	groups list	sources list
1	G_1, G_2	(G_1, S_1) $(G_2, *)$
...	...	...
i	G_k	$(G_k, *)$
...		
n	G_2	(G_2, S_1)

Table 1: Example of the Multicast Beams Table in the on-board switch. The table is indexed by the spot beam numbers and for each spot beam, it maintains the list of active multicast groups and sources.

On the reception of a multicast packet to send to the receivers a lookup to the MBT permits to the on-board switch to identify the list of the spot beam that are interested by this packet.

A third solution consists to enable the routing in the on-board switch and to perform multicast routing. This solution may will be the long-term solution but it is more complex to implement and to manage.

We believe that the short-term efficient solution is the second solution because it permits to next-generation satellites to be multicast-enabled in a dynamic fashion and so that to be integrated in the multicast Internet core. However, we have to resolve the IGMP feedback implosion problem that is not specific to IGMP but to all protocols that need feedback from the receivers. In the next section, we present and discuss solution to this problem in the context of IGMP over satellite.

3.3 IGMP Feedback Implosion Problem

In traditional networks IGMP mechanism rarely pauses a big problem, it is used on a per-router basis with only a few hosts connected to the router's physical links. Even if all these hosts simultaneously answer a feedback query for which one response message would have been sufficient not much harm is done with respect to the network load because the number of hosts directly connected to a router is small.

In the satellite-terrestrial hybrid network shown in Figure 1, one might expect 10^6 or even more hosts to be connected to the same physical link of each spot beam. In Figure 3 we show the behaviour of IGMP over satellite links. The IGMP REPORTS received by the feed are forwarded to the on-board switch and tunnelled to the send-only feeds. We suggest that when receiving the report, the on-board switch forwards the packet **only** to the spot beam to which belongs the sender and update its Multicast Beams Table according to the procedure described in Section 3.2. By this way, the receivers in other

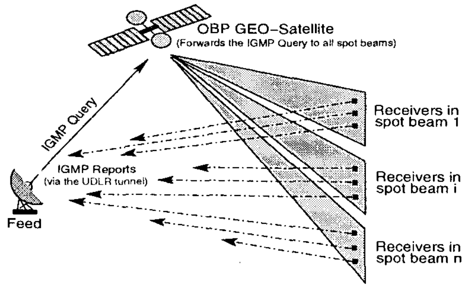

Figure 3: The feed send periodically an IGMP QUERY via the satellite uplink to each spot beam and receivers reply by an IGMP REPORT using the UDLR tunnel [7].

spot-beams will not hear this report and at least one of the receivers interested to the multicast group in each spot beam being reported will also reply to the feed query.

Here feedback algorithms must carefully avoid feedback implosion from each spot beam, i.e. a large avalanche of identical responses to a query that could be answered by any single one of the receivers. Even more, if one of the receivers answers a query other receivers in the same spot beam will not immediately notice that a response was already given. Hence, even more feedback duplicates will be sent to the network. We have to deal with a trade-off between feedback latency *and response duplicates*. With feedback latency we mean the expectation value for the time until the first response is sent.

The authors of [14] have compared the proposed solutions in order to minimise the number of responses and so that to avoid feedback implosion for reliable multicast transfer. They have shown that only one of these algorithms, the exponential feedback raise, is able to provide sufficiently stable expectation values across a large of group sizes.

Let apply this algorithm to the IGMP feedback implosion problem and denote T to be the MAX RESPONSE TIME specified by the querier in the IGMP QUERY packet. Before sending an IGMP query, the router (feed or send-only feed) should:

1. Estimate an upper limit N for the number of satellite receivers that might provide feedback responses. $N = \sum_i N_i$ where N_i is the number of satellite receivers in the spot beam number i.

2. Decide on the desired number R_0 of feedback responses or the desired upper limit for the feedback latency T. Note that these values *cannot* be chosen independently since $T = \tau \log_{R_0} N$ where τ denotes the network latency (round-trip time for responses within the network).

On the reception of an IGMP query, the interested receiver in the spot beam number i should run the feedback-algorithm as follows:

1. Choose a number $x \in [0, 1)$. If $x < 1/N_i$ the responsive receiver immediately sends a feedback-response. Otherwise it sends its response at time $t = T(1 + \log_{N_i} x)$ unless it received a response from another receiver in the same spot beam before that time.

2. If a sufficiently good estimation for the network latency τ can be given the following modification can be applied: The response interval $[0, T]$ is divided into sub-intervals of duration τ. Hosts that would respond within a given sub-interval send their response already at the beginning of the response interval.

To apply the algorithm described above, the feed has to estimate the value of N to compute T and the receivers have to know the value of N_i in order to generate the values of x and t.

In the satellite context, the number N can be approximated if the content provider uses an administrative policy to authorise only receivers which are subscribed in advance to the service to receive data from the multicast group. Otherwise analytic methods as that described in [12] which is based on a binomial(n,p) distribution can be used to estimate this value.

Even if the querier (the feed) has estimated correctly the number of receivers interested by the multicast group being queried, this number is not included in the standard IGMP QUERY [5] so the receivers can not have knowledge about the number of receivers and so that they can not compare the generated value x to $1/N_i$ and generate t. A classic solution consists to include this information in the IGMP query executed in the satellite interface.

4 DVMRP over Satellite

The first protocol developed to support multicast routing is called the Distance Vector Multicast Routing Protocol (DVMRP) [15]. It has been widely used on the MBONE.

DVMRP constructs a different distribution tree for each source and its destination host group. Each distribution tree is a minimum spanning tree from the multicast source at the root of the tree to all the multicast receivers as leaves of the tree. The distribution tree provides a shortest path between the source and each multicast receiver in the group, based on the number of hops in the path, which is the DVMRP metric. A tree is constructed on demand, using a "broadcast and prune" technique, when a source begins to transmit messages to a multicast group.

Construction of a DVMRP spanning tree in a hybrid satellite-terrestrial network is illustrated in Figure 4. A tree is constructed on demand, using a "broadcast and prune" technique, when a source begins to transmit messages to a multicast group. The designated router on the source subnet (R_1 in Figure 4), i.e., the router that has been selected to handle routing for all hosts on its subnet, begins by transmitting a multicast message to all adjacent routers.

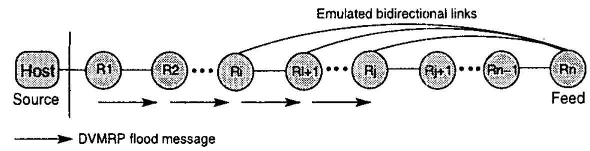

Figure 4: DVMRP behaviour over hybrid satellite-terrestrial networks. The source sends DVMRP messages towards the feed. The satellite receivers in this path send the message to all their interfaces, including the UDLR Tunnel to the feed, except that from which the message was received.

Let us interest to a path from the source that contains satellite feeds or receivers and note by R_i the **first** satellite receiver that belongs to this path and R_n the first router after R_i that has the capability to send data (a feed or a send-only feed) to the unidirectional satellite link.

The shortest path from a router R_k to the designed router R_1 is:

- R_k, ..., R_{k-1}, ..., R_2, R_1 if $k \leq j$ where $j = [\frac{i+n}{2} + 1]$
- R_k, R_{k+1}, ..., R_n, R_i, R_{i-1}, ..., R_2, R_1 otherwise.

Therefore, the routers from R_1 to R_j forward the message to their downstream routers, since they receive the message from their interface used to send data toward the source. In particular, the router R_i sends the message also to the feed R_n via the UDLR tunnel so that the routers from R_{j+1} to R_n receive the message from R_n. In fact, on the reception of a message from R_j, the router R_{j+1} sends a prune message via its interface toward R_j.

The undesirable DVMRP behaviour here is that the routers from R_i to R_j that has a satellite-receiver capability will receive **duplicated copies** of multicast packets. One copy from the terrestrial interface toward the source (it is the shortest path) and another copy from the satellite link since the feed should send data to receivers from R_{j+1} to R_n. It is should be noticed that he multicast copy received from the satellite link will not be delivered to the network layer and so will not disturb the multicast application. However, we argue that it is more efficient that the satellite copy will be used instead of the terrestrial one. For this end, we propose to proceed as follows:

- When the satellite link layer of the routers between R_i and R_j detects that there is multicast packets received via the satellite link which belongs to a group to which the router is already a member on the terrestrial interface, it triggers a DVMRP prune message towards R_i.
- They forward each message received via the satellite interface to the terrestrial interface. By this way each multicast packet received from the satellite link, it is considered by the network layer as that it was received from the terrestrial interface.

Modifications should be added only to the LLTM mechanism to execute the two above tasks.

5 PIM-DM over Satellite

Protocol Independent Multicast-Dense Mode (PIM-DM) [6] is similar to DVMRP. Both protocols employ Reverse Path Multicasting (RPM) to construct source-rooted distribution trees. The difference between DVMRP and PIM-DM is that in DVMRP, prior to forwarding to a certain interface, DVMRP makes sure that the interface leads to a node that will recognize the local node as a node that is in the shortest path between it and the source (poison-reversed route). PIM-DM decides to accept additional overhead in order to simplify the forwarding algorithm. Apart from this, the protocol is very similar to DVMRP and thus, all that has been stated for DVMRP applies to PIM-DM also.

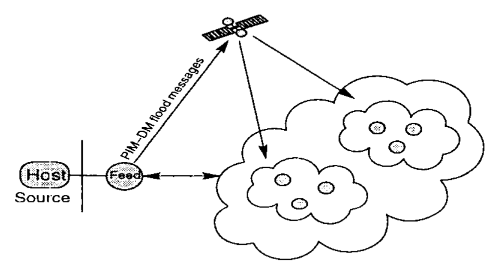

Figure 5: The feed forward the flood message to all the spot beams without checking if the receivers use the satellite link to reach the source.

The use of PIM-DM over satellite has the same problem as that of DVMRP explained in Section 4. Besides this problem, the additional overhead added by PIM-DM in a satellite network can have a big impact of the performance of this protocol. We illustrate this problem via the configuration given in Figure 5.

According to PIM-DM specification, when the feed receives multicast data from the source via one of its terrestrial interfaces, it automatically forwards the packets received to all other interfaces and in particular to the satellite link without checking if there is or no satellite receivers that use the satellite link to send to the source. Or, some of the receivers in the spot-beams may receive unicast data from the source via one of their terrestrial interfaces. The bandwidth consumed by the unused PIM-DM flood messages can be non negligeable and can have a big impact on the resources used by the other established connections.

6 PIM-SM over Satellite

6.1 Overview

The PIM-Sparse Mode (PIM-SM) protocol initially constructs a group-shared tree to support a multicast group. The tree is formed by the senders and receivers both connecting to the rendezvous point, just as a shared tree is constructed around the core with the CBT protocol [3]. After the tree is constructed, a receiver (actually the router closest to this receiver) can opt to change its connection to a particular source to a shortest-path tree. This is accomplished by having this router send a PIM join message to the source. Once the shortest path from source to receiver is created, the extraneous branches through the RP are pruned. Note that different types of trees can be selected for different sources within a single multicast group.

The PIM protocol specifies soft-state mechanisms to periodically refresh system state, adapt to topological changes in the network, and adapt to changes in group membership. While PIM relies on unicast routing tables to adapt to network topology changes, it is independent of the particular unicast routing protocol that is used to construct those tables. Other features of PIM, such as using multiple rendezvous points to eliminate the problem of having a single failure point, are too numerous to describe in this paper.

6.2 PIM-SM configuration policy for hybrid networks

In considering a routing protocol to be used for multicasting in terrestrial-satellite hybrid networks, one has to carefully look at the issues unique to this type of network and make use of the broadcast nature of GEO satellites.

PIM-SM requires routers that are directly attached to downstream members to join a sparse-mode distribution tree by transmitting explicit join messages to the group's primary Rendezvous Point (RP) which acts the root of the tree. PIM-SM creates a shared, RP-routed distribution tree that reaches all group members. PIM-SM provides also a mechanism to switch from a RP-routed tree to a shortest path tree (SPT).

Using PIM-SM effectively in hybrid networks depends on our capacity to: (1) manage the choice of the RP of each multicast group, and (2) configure the policy used by group members to switch from a RP-routed tree to a SPT.

6.2.1 RP Placement

The *bootstrap mechanism* used in PIM-SM employs an algorithmic mapping of multicast group to rendezvous point address, based on a set of available RPs distributed throughout the network by the dynamically-elected BootStrap Router (BSR) from a list of Candidate-BSRs [8]. Routers belonging to the set of Candidate-BSRs or Candidate-RPs should be manually configured in the network [8]. For hybrid networks, we assume that all feeds and send-only feeds are configured as Candidate-BSRs and Candidate-RPs[3].

In order to profit from satellite link broadcast nature, the RP placement policy that we recommend in such type of network depends on satellite uplinks (feeds) positions in the terrestrial network.

For each multicast group we process as follows to select the RP:

- If there is a feed or a send-only feed which is a multicast source of the group, it is chosen as the RP of this group. If there is more than one feed belonging to the group, the feed which is the nearest to the source is chosen.
- If there is no feed belonging to the multicast group, it is recommended to choose the closest feed for multicast sources as the RP of this group. If there is more than one feed in the hybrid network, the feed which has the highest priority for this group will be elected.

This policy of RP placement in a hybrid network can be satisfied by effectively choose the priority value of each RP in the Candidate-RP message sent to the BSR.

6.2.2 Switching from a RP-routed tree to a SPT

PIM-SM specification [8] does not specify a fixed policy to switch from the RP-routed tree to the SPT, but it recommends that the router monitors data packets from sources for which it has no source-specific multicast route entry and initiates such an entry when the data rate exceeds the configured threshold. Let us apply this method to a hybrid network. If at least one satellite receiver that is a member of the multicast group decides to switch to the SPT and when the SPT contains a feed, all satellite receivers will receive multicast packets sent by the source. Multicast receivers that are still using the RP-routed tree will receive a **duplicated packet**: a copy from the terrestrial interface belonging to the RP-routed tree and another copy from the satellite interface.

We propose a switching policy that can be used effectively in a terrestrial-satellite hybrid network. This policy consists as follows:

- If the source is a feed, it makes sense for all satellite receivers members to join source-specific tree and prune the source's packets off the shared RP-centrated tree since it forwards data to members via the satellite link. Or, the RP triggers Register-Stop messages in response to Registers sent by the source only if the RP has no downstream receivers for the group (or for that particular source), or if the RP has already joined the (S,G) tree and it receiving the data packets natively. Then we recommend that all satellite receivers members join directly

[3] [8] recommends that C-BSRs should be equal to C-RPs.

the source. This can be done by properly configured the threshold maintained by the router. For example, we can attribute for each (S,G) a threshold value close to zero to guarantee that each satellite receiver member switch to the SPT when the source is a feed.

- if the source is not a feed, it is not desired to switch from the RP-routed tree to the SPT especially when the RP is a feed. Multicast packets will be sent by the source to the RP via the UDLR tunnel [7]. The threshold maintained by each member for the (S,G) entry should be then greater.

PIM-SM, used with UDLR, grants an efficient use of communication links resources considering that multicast packets will be sent via the satellite downlink. Terrestrial links will then be used effectively by applications that need resources that cannot be offered by the satellite connectivity.

7 Conclusion

In this paper, we have discussed the behaviour of the IP multicast standard model in satellite-terrestrial hybrid networks. We have addressed many problems related to the use of IP-level multicast protocols in such type of networks. We have presented a global architecture of an Internet network including the next-generation of satellite systems that support multiple spot beams and which has on-board switching capability. We then focus on the IGMP feedback implosion problem and we present the exponential feedback raise that able to provide sufficiently stable expectation values across a large of group sizes. To enable the multicast in the on-board satellite switch, we have investigated different approaches. We show that the approach based on maintaining a Multicast Beams Table (MBT) containing the list of spot beams concerned by each multicast group is the most suitable approach that allows an efficient and transparent integration of satellite links in the Internet. In fact, the on-board switch monitors the IGMP reports and queries sent over the satellite link and update the MBT table. In the second part of this paper, we have addressed the problem of multicast dynamic routing protocols over satellite links. For DVMRP and PIM-DM, we have identified some network configurations where the satellite receivers can receive duplicated packets and we propose a method to overcome this problem. We have developed a configuration policy of PIM-SM in hybrid networks concerning the choice of the list of Rendezvous Point (RPs) and the switching from the RP-routed tree to the shortest path tree.

In our future work we will evaluate and compare the different approaches proposed here to enable the multicast in the on-board switch. We will implement the different approaches and solutions proposed here and we evaluate their performance using simulations and real experimentations in our UDLR testbed network.

References

[1] F. Ananaso and F. Priscoli, *The role of Satellite in Personal Communication Services*, IEEE JSAC SAC-13, pp. 180-196, Febrary 1995.

[2] K. C. Almeroth and Y. Zhang, *Using Satellite Links as Delivery Paths in the Multicast Backbone (MBone)*, WOSBIS'98, Dallas, Texas, October 1998.

[3] A. Ballardie, *Core Based Trees (CBT Version 2) Multicast Routing: Protocol Specification*, IETF, RFC 2189, September 1997.

[4] B. Cain, S. Deering, and I. Kouvelas, *Internet Group Management Protocol: version3*, IETF, June 2000.

[5] S. Deering, *Host Extensions to IP multicasting*, IETF, RFC 1112, May 1990.

[6] S. Deering et al., *Protocol independent multicast version 2 dense mode specification*, IETF, draft-ietf-pim-v2-dm*.txt, June 1999.

[7] E. Duros, W. Dabbous, H. Izumiyama, N. Fujii, and Y. Zhang, *A Link Layer Tunneling Mechanism for Unidirectional Links*, IETF, RFC 3077, March 2001.

[8] D. Estrin et al., *Protocol independent multicast sparce-mode (PIM-SM): Protocol specification*, IETF, RFC 2362, June 1998.

[9] D. Farinacci, T. Li, S. Hanks, D. Meyer, and P. Traina, *Generic Routing Encapsulation (GRE)*, RFC 2784, IETF, March 2000.

[10] W. Fenner, *Internet Group Management Protocol, Version 2*, IETF, RFC 1112, November 1997.

[11] F. Filali, W. Dabbous, and F. Kamoun, *Optimization of GEO Satellite Links Deployment in the Internet*, Research Report RR-3925, INRIA Sophia-Antipolis, March 2000.

[12] T. Friedman and D. Towsley, *Multicast session membership size estimation*, In the Proceedings of IEEE Infocom'99, pp. 965-972, March 1999.

[13] J. Hawkinson and T. Bates, *Guidelines for creation, selection, and registration of an Autonomous System* (AS), IETF, RFC 1930, March 1996.

[14] J. Nonnenmacher and E. W. Biersack, *Scalable Feedback for Large Groups*, IEEE/ACM Transactions on Networking, Vol. 7, Issue 3, pp. 375-386, 1999.

[15] D. Waitzman, C. Partridge, and S. Deering, *Distance vector multicast routing protocol (DVMRP)*, IETF, RFC 1075, November 1988.

[16] Y. Zhang, D. De Lucia, B. Ryu, and S. K. Dao, *Satellite Communications in the Global Internet: Issues, Pitfalls, and Potential*, INET'97, 1997.