

IP MULTICAST OVER SATELLITE

Z. Sun, H. Cruickshank and S. Iyengar and M. P. Howarth, CCSR, University of Surrey, GU2 7XH, UK

L. Claverotte, Alcatel Space Industries, Toulouse, France

R. Gomez, Alcatel Espacio, Madrid, Spain

Tel: +44 (0)1483 68 9493, Fax: +44 (0)1483 68 6011, Email: z.sun@surrey.ac.uk

ABSTRACT

Satellites have been very successful for digital broadcast, and it is a natural step to develop IP multicast over satellite to explore the potential of the satellite systems. This paper discusses the issues of IP multicast over satellite, based on IP multicast protocols using different satellite communications systems including DVB-S and DVB-RCS. The paper provides an overview of the networking issues and their interactions. The work is based on the studies of the European IST GEOCAST project within the IST's 5th Framework programme. The paper explains how IP multicast protocols can be implemented over satellite links, looking particularly at the issues involved in implementing IGMP, multicast routing protocols, reliable multicast protocols, and security.

KEYWORDS: *satellite, multicast, Internet, IP, QoS, DVB, IGMP*

1 Introduction

Over the last decade, different satellite systems have been studied for broadband services; these systems include IP over satellite, Digital Video Broadcasting via Satellite (DVB-S) and DVB Interactive Return Channel via Satellite (DVB-RCS). Due to the recent significant increase of applications and services in the Internet, satellites have been used for Internet access and interconnection. Historically, the Internet supports only classical data applications based on best-effort services of the Internet Protocol (IP) such as file transfer (FTP) and world-wide web (WWW). Much work has also been conducted with the objective of supporting real-time multimedia and multicast applications with specific quality of service (QoS) requirements. The success of satellite digital broadcast services and the asymmetric nature of IP traffic have made satellites a potential candidate to deliver IP multicast services.

Research has been carried out in the area of IP multicast over satellites within the project "Multicast Over Geostationary EHF Satellites (GEOCAST)" [1] supported by the EU 5th Framework as part of the EU Information Systems Technologies programme [2]. This project is investigating the potential of satellites for IP multicast applications, and how IP multicast protocols can be implemented over a satellite link, including the issues involved in implementing IGMP, multicast routing protocols, reliable multicast protocols and security.

In Section 2 the paper introduces satellite networking and multicast including network applications and services, IP multicast protocols, addressing, group management, routing and IP packet encapsulation, satellite systems and technologies. Section 3 discusses how the use of satellite systems affects a number of aspects of the behaviour of IP multicast. Section 4 continues this discussion, focussing on satellite security systems. Finally, Section 5 concludes the paper.

2 Satellite IP networking and multicast

Satellite IP networking and multicast issues concern several inter-related areas, including services and applications and their QoS requirements and suitable network protocols. In this Section we begin by briefly considering the applications and services that the network is required to support; we then review IP multicast, and mechanisms for providing network support for QoS. We then consider current satellite architectures and finally discuss how satellite systems fit into IP networks.

2.1 Network applications and services

From the network perspective we can consider two main categories for broadband services: interactive services and distribution services.

The interactive services can be subdivided into three classes of services as follows:

- Conversational services: typical examples are video telephony, video-conferencing, video/audio information transmission, high speed digital information, file and document transfer;
- Message services: typical examples are document mail and video mail; and
- Retrieval services: typical examples are video, high-resolution image, document and data storage (for example in databases).

The distribution services can be subdivided into two classes:

- The class without user individual representation control, such as TV, multimedia video and audio distribution; and
- The class with user individual representation control, such as Pay TV (PTV).

Each of these services may have different QoS requirements, such as delay-sensitive or jitter-sensitive real time data, or loss-sensitive transaction data. In general these services may be either point-to-point (unicast) or point-to-multipoint (multicast) or multipoint-to-multipoint (also multicast); from this it can be seen that multicast needs to be capable of supporting the different QoS requirements.

2.2 IP multicast

We now proceed to review IP multicast technology. Multicast allows a communications network source to send data to multiple destinations simultaneously whilst transmitting only a single copy of the data on to the network. The network then replicates the data and fans it out to recipients as necessary. Multicast can be considered as part of a spectrum of three types of communications:

- Unicast: transmitting data from a single source to a single destination (for example, downloading a web page from a server to a user's browser; or copying a file from one server to another);
- Multicast: transmitting data from a single source to multiple destinations. The definition also encompasses communications where there may be more than one source (i.e. multipoint-to-multipoint). Videoconferences provide an example of this latter, where each participant can be regarded as a single source multicasting to the other participants in the videoconference.
- Broadcast: transmitting data from a single source to all receivers within a domain (for example within a LAN; or from a satellite to all receivers within a satellite spotbeam).

The advantages of multicast are as follows:

- Reduced network bandwidth usage: for example, if data packets are being multicast to 100 recipients the source only sends a single copy of each packet. The network forwards this to the destinations, only making multiple copies of the packet when it needs to send packets on different network links to reach all destinations. Thus only a single copy of each packet is transmitted over any link in the network, and the total network load is reduced compared to 100 separate unicast connections. This is particularly beneficial on satellite systems where resources are limited and expensive.
- Reduced source processing load: the source host does not need to maintain state information about the communications link to each individual recipient.

Multicast can be either best effort or reliable. "Best effort" means that there is no guarantee that the data sent by any multicast source is received by all or any receivers, and is usually implemented by a source transmitting UDP packets on a multicast address (the addressing mechanism is described in further detail below). "Reliable" means that mechanisms are implemented to ensure that all receivers of a multicast transmission receive all the data that is sent by a source: this requires a reliable multicast protocol.

2.2.1 IP multicast addressing

Each terminal or host in the Internet is uniquely identified by its IP address. In IP Version 4 this consists a 32-bit address space, divided into a network number and a host number which respectively identify a network and the terminal attached to the network. A normal unicast IP datagram includes a source address and destination address in the IP packet header; routers use the destination address to route the packet from the source to the destination. Such a mechanism cannot be used for multicast purpose, since the source terminal may not know when, where and which terminals will try to receive the packet.

Consequently, a range of addresses is defined for multicast purposes only. The range of addresses, called Class D addresses, is from 224.0.0.0 to 239.255.255.255. Unlike Classes A, B and C, these addresses are not associated with any physical network number or host number, but instead are associated with a multicast group that is like a radio channel; members of the group receive multicast packets sent to this address, and the address is used by multicast routers to route IP multicast packets to users that register for a multicast group. The mechanism by which a terminal registers for a group, IGMP, is described below.

2.2.2 Multicast group management: IGMP

In order to make efficient use of network resources, the network sends multicast packets only to those networks and subnets that have users belonging to the multicast group. The Internet Group Membership Protocol (IGMP) [3] allows hosts or terminals to declare an interest in receiving a multicast transmission. IGMP supports three main types of message: Report, Query and Leave.

A terminal wishing to receive a multicast transmission issues an IGMP join Report, which is received by the nearest router. This Report specifies the IP multicast class D address of the group being joined. The router then uses a multicast routing protocol (described below) to determine a path to the source. To confirm the state of terminals receiving multicast, a router occasionally issues an IGMP Query to terminals on its network/sub-network. When a terminal receives such a query, it sets a separate timer for each of its (potentially many) group memberships. When each timer expires, the terminal issues an IGMP Report to confirm that it still wishes to receive the multicast transmission. However, in order to suppress duplicate reports for the same Class D group address, if a terminal has already heard a report for that group from another terminal it stops its timer and does not send a Report. This has the benefit of avoiding flooding the subnetwork with IGMP Reports.

When a terminal wishes to finish receiving the multicast transmission it issues an IGMP Leave request¹. If all the members of a group in a subnet have left, the router does not forward any more multicast packets to that subnet.

2.2.3 IP multicast routing

In a normal IP router used for unicast, the routing table contains information that specifies paths that lead to a given IP destination addresses. However, this routing table is not useful for IP multicast since multicast packets do not contain information about the location of the packet's destinations. Therefore different routing protocols and routing tables have to be used. Multicast routing protocols address the issue of identifying a route for data to be transmitted across a network from a source to all its destinations, while minimising the total network resources required for this.

In IP multicast, the routing table is effectively built from destinations to the sources rather than from sources to destinations, since only the source address in the IP datagram corresponds to a single physical location. Tunnelling techniques may also

¹ The Leave message is supported in IGMP Version 2. In Version 1, a host / terminal quietly changes its state to non-member, and no message is sent to the router.

be used to support multicast over routers that do not have multicast capabilities.

A number of multicast routing protocols have been developed by the IETF. These include Multicast Extensions to OSPF (M-OSPF) [4], Distance Vector Multicast Routing Protocol (DVMRP) [5], Protocol-Independent Multicast - Sparse Mode (PIM-SM) [6] and PIM Dense Mode (PIM-DM) [7], and Core-Based Tree (CBT) [8]. An overview of these multicast routing protocols is available in [9].

Here we briefly review the underlying principle of operation of two protocols. DVMRP and PIM-DM are "flood and prune" algorithms: in these protocols, when a source starts sending data, the protocols flood the network with the data. All routers that have no multicast recipients attached send a prune message back towards the source (they know they have no receivers because they have received no IGMP join Reports). These protocols have the disadvantage that a "prune" state is required in all routers (i.e. "I have pruned on this multicast address"), including those routers with no multicast recipients downstream.

Flood and prune protocols use Reverse Path Forwarding (RPF) to forward multicast packets from a source to the recipients: the RPF interface for any packet is the interface that the router would use to send unicast packets to the packet source. If a packet arrives on the RPF interface it is flooded to all other interfaces (unless they have been pruned), but if the packet arrives on any other interface it is silently discarded. This ensures efficient flooding and prevents packet looping.

DVMRP uses its own routing table to compute the best path to the source, whereas PIM-DM uses an underlying unicast routing protocol.

2.2.4 IP multicast scope

Scoping is the mechanism that controls the geographical scale of a multicast transmission, by making use of the time to live (TTL) field in the IP header. It tells the network how far (in terms of router hops) any IP packet is allowed to propagate, allowing IP multicast sources to specify whether packets should be sent only to the local sub-network, or to larger domains or the whole Internet. This is achieved by each router reducing the TTL by 1 when forwarding the packet to the next hop, and discarding the packet if the TTL is 0. Each subnet may additionally have filters or firewall to discard some packets according its security policy that may be beyond the control of the multicast source.

It can be seen in a satellite network that even with a small TTL value, IP multicast packets can reach a very large number of members of a multicast group scattered in a very large geographical area.

2.2.5 Address mapping and configuration

Different network technologies may use different addressing scheme for assigning addresses, also called physical addresses, to devices. For example, an IEEE 802 LAN uses a 48-bit address for each attached device, and ISDN uses the ITU-T E.164 address scheme. Similarly, in a satellite network each ground earth station or gateway station has a physical address to be used for circuit connections or packet transmissions. However, the routers that are interconnected by the satellite network know only the IP addresses of the other routers. Therefore, address mapping between each IP address and its related physical address is required, so that packet exchanges between the routers can be carried out through the satellite network using the physical addresses. The precise details of this mapping depend on the underlying data link layer protocols used over the satellite.

2.3 Network QoS

Internet network layer protocols (IP) provide a “best effort” support to services and applications across different data link layer technologies, such as LAN, MAN, WAN and satellite links. IP datagrams have to be encapsulated and transported across these different data link layer technologies. In a satellite environment, the data link layer may be traditional transparent data links, or it may be broadband links based for example on DVB-S or DVB-RCS.

To improve on the best effort service provided by the IP network layer protocol, new mechanisms such as Differentiated Services (“DiffServ”) [10], and Integrated Services (“IntServ”) [11], have been developed to support QoS. In the DiffServ architecture, services are given different priorities and resource allocations so that various types of QoS can be supported. In the IntServ architecture, resources have to be reserved for individual services. However, resource reservation for individual services does not scale well in large networks, since a large number of services have to be supported, each maintaining its own state information in the network’s routers. Flow-based techniques such as Multi-Protocol Label Switching (MPLS) [12] have also been developed to combine layer 2 and layer 3 functions to support QoS requirements.

At the transport layer the Internet community has developed mechanisms that provide a basic level of QoS: for unicast these are TCP, used to provide guaranteed and ordered data delivery for reliable connection-oriented services, and UDP for best effort connectionless services. For multicast, a number of reliable multicast protocols have been developed for research and experimental use; UDP is again used for best effort connectionless multicast traffic.

Important network QoS parameters include end-to-end delay, delay variation and packet loss. These have to be measured in an end-to-end reference path, where the propagation delay of satellite links has to be taken into account properly.

2.3.1 IP packet encapsulation

Different network technologies may also use different frame formats or frame sizes for transporting IP datagrams. IP packet encapsulation is a process that puts the packet into the payload of a data link layer frame for transmission over the network: for example, Ethernet and Token Ring LANs have their own standard frame formats. Similarly, in a satellite link, IP packets have to be encapsulated into frames. In DVB-S, IP packets including multicast are encapsulated in an Ethernet-style header using a standard called Multi-Protocol Encapsulation (MPE).

Due to the differences of framing format, different encapsulation techniques may be used. Sometimes, an IP packet may be too large to fit into the frame payload. In such a case, the IP packet has to be broken up into smaller segments (fragmented) so that the IP packet can be carried over several frames. In this case, additional overhead is added to the segments so that on arriving at the destination, the original IP packet can be reassembled from the segments. It can be seen that the encapsulation process may have a significant impact on network performance due to the additional processing and overhead.

2.4 Satellite architectures

2.4.1 Satellite network roles

Each of services can be provided either over terrestrial networks or satellite networks. The particular benefits provided by satellites include their geographically extended coverage (including land and sea), their efficient delivery to a large number of users on a large scale, and the low marginal cost of adding additional users. A satellite can play several different roles in a network:

- Last mile connections: end-users are directly connected to the satellite, which provides direct forward and return links. Traffic sources connect to the satellite feeder or hub stations through the Internet, tunnelling or dial-up links. The end-user terminals are the destination for IP datagrams, and can be considered to consume the IP packets.
- First mile connections: the satellite provides forward and return link connections directly to a large number of ISPs’ routers or gateways, which in turn forward IP datagrams on to the end-users in accordance with IP routing

protocols. As with the last mile connections, traffic sources connect to the satellite feeder or hub stations through the Internet, tunnelling or dial-up links, but the here ISP does not consume the IP packets.

- Transit connections: the satellite provides connections between Internet gateways or ISPs' gateways. The traffic is routed through the satellite links according to specified routing protocols and defined link metrics in the networks so as to minimise connection costs and to meet required QoS constraints for the given traffic sources.

2.4.2 Routing and satellites

The Internet consists of a collection of subnetworks, or autonomous systems (ASs) or domains. There is no overall structure or architecture, but several backbone networks exist to which the ASs may be connected. Within each AS, routers (called Interior Gateway Routers, IGRs) normally use an interior gateway routing protocol. Between ASs, routers use a border gateway routing protocol; these routers are either called Border Gateway Routers (BGRs) or sometimes Exterior Gateway Routers (EGRs). Figure 1 gives an example how satellites can be used to support each of the types of connections, with IGRs used to route traffic within each domain and BGRs used between different domains.

Each AS is operated by a single organisation, with all routers using the same routing protocols and network metrics. These routers exchange information about links and networks resources and can dynamically update their routing tables with shortest path information, allowing them to route IP packets efficiently. Different ASs may use different routing protocols and network metrics to update their routing tables. Routing between ASs is provided by the BGRs, which obtain routing information to other ASs. There are also policies set in the BGRs to control whether IP packets are allowed to be routed through a particular AS or whether an IP packets may be accepted from other named ASs.

Although a satellite network can be used for connections within an AS or between ASs, it is often practical to have the satellite network form its own AS. This is because the satellite network is often operated and managed separately from the terrestrial networks, and it has different characteristics in term of bandwidth, connection cost and coverage. In such a case, IGR protocols can be used within the satellite network.

The original unicast IGR protocol was a distance vector protocol, called Routing Information Protocol (RIP). Due to the count-to-infinity and slow convergence problems of the protocol, it was replaced by a link state protocol in May 1979. In

1988, the IETF began to work on a successor, called Open Shortest Path First (OSPF) that became an Internet standard in 1990.

In the case of multicast routing protocols, the set listed in Section 2.2.3 are multicast IGR protocols. To establish multicast communication across ASs, BGRs use a multicast border gateway routing protocol, specifically BGMP [13].

2.4.3 Satellite network technologies

Satellite systems and technologies concern two aspects: the ground segment and the space segment. In the ground segment, there are several constraints such as physical size, access, and trade-offs between transmission power, data rate and mobility. In the space segment, various types of technology can be used for the data link layer, such as transparent (bent-pipe), on-board processing, on-board packet switching and, recently, on-board DVB switching [14], [15], [16]. In addition, various hybrid options can also be implemented, such as IP over DVB.

Transparent satellite links provide data link layer connections, where on-board signalling and control are minimal, but they may not provide optimised resource utilisation. Conversely, on-board processing and switching satellites can provide optimised resource utilisation, but at the cost of complexity of on-board signalling and control.

Future satellites with on-board DVB switching will be able to integrate broadcast and interactive services by combining DVB-S and DVB-RCS standards. A DVB regenerative payload can multiplex information from diverse sources into a standard downlink DVB-S stream [17]. Another example of the use of DVB on-board switching is to interconnect LANs using IP over MPEG-2 encapsulation, via a regenerative satellite payload [18].

The GEOCAST project includes designs for both a transparent satellite system and a satellite system with on-board processing and switching capabilities [1]. The transparent satellite system provides a simple solution to support a star topology for networking. For the OBP option, the forward link is based on DVB-S/MPEG-2 and the return link on DVB-RCS. The OBP satellite system has additional functionality, being capable of supporting multiple spotbeam "star" (point-to-multipoint centred on a gateway earth station) and "mesh" (multipoint-to-multipoint) topologies, Figure 2. The OBP system is also more flexible, with better utilisation of satellite bandwidth resources.

None of the above data link layer technologies were originally designed to support IP multicast over satellites, but they are now being adapted to support

this capability, using for example the packet encapsulation techniques.

If networks evolve towards an all-IP solution, a further option needs to be investigated: an all-IP satellite with on-board router. Such an option will need a significant amount of new system design, and will need to convince industrial players of the benefits of developing and deploying satellite payloads based on this new router technology instead of existing technologies. The benefit of an IP-router-in-the-sky approach is that the routing algorithms can be used to integrate the satellite links in an IP multicast routing tree at the source, trunk or end branch, as first mile connections, transit connections or last mile connections.

3 IP multicast over satellite issues

3.1 IGMP behaviour in satellite environments

In a satellite environment, multicast group management together with the scoping mechanism may provide an efficient solution to support IP multicast with large numbers of users distributed over a large area. However, IGMP over satellites raises interoperability issues, as we now describe.

In a conventional terrestrial LAN, an IGMP Report is heard by other multicast receivers on the LAN, and this prevents flooding of the LAN with multiple reports. In a satellite system, individual ground stations cannot hear each other; given the large number of multicast receivers that are expected in satellite systems (potentially of the order of 10^5 or 10^6) multiple IGMP Reports could cause significant flooding of the satellite network with IGMP traffic. One of a number of adaptations of IGMP and multicast must therefore be implemented. Two options are as follows, illustrated with an example of multicast from an uplink gateway earth station out to multiple end-user terminals each with a router (Figure 3):

- Multicast channels can be statically configured to be transmitted across the satellite link to each downlink router, with IGMP traffic only operating between a router and the end-user terminal (Figure 3a). There is no transmission of IGMP traffic across the air interface in this case. This is a simple option, but potentially wastes scarce satellite channel capacity if there are no listeners on a particular multicast channel within any spotbeam;
- Multicast channels are (as in conventional terrestrial networks) only transmitted across the satellite link if there is one or more listening end-user. IGMP messages are transmitted across the air interface. When the uplink router receives an IGMP Report from a terminal following an IGMP Query, either the router

must retransmit the IGMP Report via the satellite to all ground stations to avoid flooding, or else other receivers will also transmit IGMP Reports resulting in flooding (Figure 3b).

In architectures that have no router on the downlink side, IGMP “proxying” can be used to forward multicast traffic to group members while avoiding transmission of IGMP traffic over the air interface.

3.2 Multicast routing protocols in a satellite environment

We illustrate the issues in transmitting multicast routing protocols across a satellite with two examples based on multicast interior gateway routing protocols.

In the first example, we consider a flood and prune algorithm (such as is used in DVMRP or PIM-DM). When a source starts to transmit, the data is flooded across the network. In Figure 4a the underlying data link layer supports a point-to-multipoint connection, and the data from the source is correctly flooded out from router R4 to routers R1, R2 and R3. This requires a point-to-multipoint circuit from every such source on the multicast group; this could be expensive in the case of a large multicast group dynamically configured so that every satellite terminal can potentially transmit data from a data source. On the other hand, in Figure 4b the source transmits through router R4 to the uplink gateway router R1. This router then has to flood the data back out through its RPF interface in order to multicast to routers R2 and R3. This is in contravention of the normal RPF algorithm, and requires modification of the routing algorithm.

In the second example, we consider the CBT multicast routing protocol. This protocol creates a tree that joins receiving members of the multicast group. When a source transmits to the group, the data is forwarded by all network routers until it reaches either the core of the tree or a router on the multicast tree. The tree then propagates the data both out to its downstream leaves and back up into the core. In general therefore, the tree carries multicast traffic in both directions, depending on where the data from the source first reaches the tree. However, satellite links with terrestrial return paths have different forward and return path routes, and so are not suitable for such bi-directional multicast routing protocols.

3.3 Reliable multicast protocols over satellites

Reliable multicast protocols address the issue of ensuring that data is multicast from a source to all the multicast recipients and that each packet sent by the source is successfully received by all recipients. Reliable multicast protocols usually also ensure ordered and non-duplicated delivery of packets.

Since they provide an end-to-end service they are conventionally regarded as transport layer protocols in the context of the OSI Reference Model.

A wide range of reliable multicast protocols has been developed and described in the literature. One reason for this is that efficient multicast is a much more complex problem than efficient unicast, and consequently many multicast protocols have been developed for specific classes of application. Two examples of different application classes are real-time applications (requiring low delay with moderate packet loss acceptable) and multicast file transfer (requiring zero packet loss, but delay-insensitive), each of which has its own specific multicast requirement.

Two of the principal drawbacks associated with satellite links are their error characteristics and the round-trip delay, particularly in geostationary links. We first consider the error characteristics. Historically, satellite links have had high bit error rates, and in addition the channel coding typically used on satellites to maintain a low bit error rate means that errors tend to occur in bursts [19]. The consequent corruption of data means that when there are a large number of multicast end-users there is a significant probability that one or more recipients will not receive the data; this has implications for the design of reliable multicast network protocols [20]. Broader issues of the behaviour of reliable multicast protocols over satellite were considered in [21]; the performance of reliable multicast protocols over satellite has been modelled [22], [23] and measured on satellite emulators [24].

The high round-trip delay times, especially of satellite in geostationary orbits, is well-known to have an adverse impact on two-way real-time communications (for example, telephone conversations or videoconferences), and also effects the behaviour of network protocols such as TCP. A number of mitigating techniques for TCP traffic have been developed [25]. However, no corresponding standard mechanisms have yet been developed for reliable multicast protocols.

In summary, developing reliable multicast protocols and optimising them, particularly for scalability, throughput, flow control and congestion control, is an on-going research issue both for terrestrial networks and for networks that include satellite links.

4 Satellite IP security for multicast

The challenge of security in satellite environments is considered to be one of the main obstacles to the widespread deployment both of satellite IP multicast and of satellite multimedia applications in

general [26]. The main problem is that eavesdropping and active intrusion is much easier than in terrestrial fixed or mobile networks because of the broadcast nature of satellites. In addition, the long delays and high bit error rates experienced on satellite systems may cause loss of security synchronisation. This demands a careful evaluation of encryption systems to prevent Quality of Service (QoS) degradation because of security processing. A further issue, specific to multicast, is that the number of members in a multicast group can be very large and can change very dynamically.

Encryption mechanisms depend on the data link layer protocols used. In this Section we describe in detail one particular security system, designed specifically for satellite systems, namely that used in DVB-RCS, which is link level security. We then briefly consider IPSEC and multicast IP security mechanisms, which is IP network level security and can be used to provide end-to-end security.

4.1 DVB-RCS security

Security in general is intended to protect the end-user identity (including their exact location), data traffic to and from the user, signalling traffic and also protect the network operator against use of the network without appropriate authority and subscription. In DVB, two levels of security can be applied:

- DVB common scrambling which is used for conditional access for satellite broadcasts in DVB-S system. This type of scrambling is not suitable for individual and multicast IP traffic.
- Individual user scrambling in the forward and return link.

Although the user/service provider could use their own security systems above the data link layer, it is usually desirable to provide a security system at the data link layer so that the satellite link is secure without recourse to additional measures. Link level security is particularly desired by satellite access network operators in order to secure satellite links and provide their clients (such as ISPs) with data confidentiality.

For DVB, the satellite interactive network is based on the DVB/MPEG-TS Standard. The security concept is shown in Figure 5, which is taken from [27]. There are two stages in the security procedures in DVB-RCS: First the satellite terminal will be authenticated by the satellite network, by using a secret key stored in the set top box called Cookie. The Cookie can be update periodically. The second stage is data privacy (encryption/decryption). Data privacy can be performed either in the MPEG-TS level or it can be performed in the layer above the MPEG-TS and

usually within the Multi Protocol Encapsulation (MPE) that carry the IP packets.

4.2 IP multicast security

In contrast to link level (such as DVB-RCS) IP network level security has the potential to provide a better integration of satellite and IP networks and providing end-to-end security. The security architecture of the Internet Protocol known as IP security (IPSEC) is the most advanced effort in the standardization of Internet security. It consists of an authentication protocol: Authentication Header (AH), a confidentiality protocol: Encapsulated Security Payload (ESP) and it also includes an Internet Security Association Establishment and Key Management Protocol (ISAKMP). These security protocols are designed for both IP version 4 (IPv4) and IP version 6 (IPv6) environments.

In secure IP multicast, one of the principal issues is that of ensuring that the key used to encrypt traffic is known to all the member of the group, and only to those members: this is the issue of key management and distribution. The size and dynamics of the multicast group have a great impact on the key management distribution system, especially for large groups. There are several architectures for key management that are currently the subject of research [28], [29]. Another area of significant research effort is that of ensuring that key management is scalable to the large groups that are expected in satellite multicast; one of the most promising such mechanisms is the logical key hierarchy [30] and its derivatives. These keys could then be used in a security architecture such as IPSEC. This research is being conducted independently of any satellite considerations, but the results are expected to be applicable to secure IP multicast satellite systems.

5 Conclusions

This paper has reviewed the technical challenges that faces designers of satellite systems capable of carrying IP multicast traffic efficiently. The main applications of IP multicast over satellite are data distribution and multimedia streaming at a large scale and with global coverage. We have shown how multicast networking technologies have to be modified to take into account the use of a satellite link: these technologies include IGMP, multicast routing protocols, and reliable multicast protocols. Security systems also need to be carefully designed for satellite IP multicast, both to ensure confidentiality of traffic (and authentication of terminal assets) and to ensure efficient use of expensive satellite resources by implementing effective key management in multicast groups. Overall these technical challenges need to be brought together and integrated with the chosen link

layer standards so as to provide the required quality of service for the network applications and services.

6 Acknowledgements

The authors are pleased to acknowledge the support from the European Union IST Programme and the GEOCAST (IST 1999 11754) and ICEBERGS (IST 2000 31110) projects.

7 References

- [1] GEOCAST project home page, <http://www.geocast-satellite.com/>.
- [2] Information Society Technologies Programme, <http://www.cordis.lu/ist/>.
- [3] Fenner W. Internet group management protocol, version 2. *IETF RFC2236*, Nov 1997.
- [4] Moy J. Multicast extensions to OSPF. *IETF RFC1584*, Mar 1994.
- [5] Waitzman D, Partridge C, Deering S. Distance vector multicast routing protocol. *IETF RFC1075*, Nov 1988.
- [6] Estrin D et al. Protocol independent multicast - sparse mode (PIM-SM): protocol specification. *IETF RFC2362*, Jun 1998.
- [7] Deering S, Estrin DL et al. The PIM architecture for wide-area multicast routing. *IEEE Trans. Networking* 1996; **4**(2):153-162.
- [8] Ballardie A. Core based trees (CBT version 2) multicast routing. *IETF RFC2189*, Sep 1997.
- [9] Sahasrabuddhe LH, Mukherjee B. Multicast routing algorithms and protocols: a tutorial. *IEEE Network* 2000; **14**(1):90-102.
- [10] Blake S et al. An architecture for Differentiated Services. *IETF RFC2475*, Dec. 1998.
- [11] Braden R, Clark D, Shenker S. Integrated Services in the Internet architecture: an overview. *IETF RFC1633*, Jun 1994.
- [12] Rosen E, Viswanathan A, Callon R. Multiprotocol label switching architecture. *IETF RFC3031*, Jan 2001.
- [13] Thaler D. Border Gateway Multicast Protocol (BGMP): protocol specification. *IETF Draft, work-in-progress, draft-ietf-bgmp-spec-03.txt*, 30 Jun. 2002, expires Dec. 2002.
- [14] Akyildiz IF et al. Satellite ATM networks: a survey. *IEEE Communications* 1997; **35**(7):30-43.
- [15] Bem DJ, Wiekowski TW et al. Broadband satellite systems. *IEEE Communications* 2000; **3**(1):2-14.
- [16] Yegenoglu F, Alexander R, Gokhale D. An IP transport and routing architecture for

- next-generation satellite networks. *IEEE Network* 2000; **14**(5):32-38.
- [17] Lamarca M, Prat J et al. DVB-Forward: a digital television / Internet payload. *Proc. 19th AIAA Int. Comms. Sat. Systems Conference*, Toulouse, France, 2001.
- [18] Claverotte L et al. DILAN: DVB PROC for the interconnection of LANs using the STENTOR satellite. *Proc. 19th AIAA Int. Comms. Sat. Systems Conference*, Toulouse, France, 2001.
- [19] Heissler JR, Barsoum YA, Condello R. An analysis of the Viterbi decoder error statistics for ATM and TCP/IP over satellite communication. *Proc. Milcomm 1999*, pp.359-363.
- [20] Howarth MP, Cruickshank H, Sun Z. Unicast and multicast IP error performance over an ATM satellite link. *IEEE Comms. Letters* 2001; **5**(8):340-342.
- [21] Koyabe M, Fairhurst G. Reliable multicast via satellite: a comparison survey and taxonomy. *Int. J. Sat. Comms.* 2001; **19**(1):3-28.
- [22] Donner A, Bovelli S, Shabdanov S. Reliable multicast based on DVB-RCS. *Proc. 20th AIAA Int. Comms. Sat. Systems Conference*, Montreal Canada, May 2002, paper AIAA 2002-1879.
- [23] Howarth MP, Sun Z. Performance of an IP reliable multicast protocol over a GEO satellite ATM link. *Proc. 20th AIAA Int. Comm. Sat. Systems Conf.*, Montreal Canada, May 2002, paper AIAA-2002-1876.
- [24] Linder H, Miloucheva I, Clausen HD. A forward error correction based multicast transport protocol for multimedia applications in satellite environments. *Proc IEEE Intl. Perf. Comp. and Comms. Conf.*, Feb 1997, pp.419-425.
- [25] Allman M, Glover D, Sanchez L. Enhancing TCP over satellite channels using standard mechanisms. *IETF RFC2488*, Jan. 1999.
- [26] Cruickshank H et al. Securing multimedia services over satellite ATM networks. *Int J. Sat. Comms.* 1998; **16**(4):183-195.
- [27] ETSI EN 301790. Digital Video Broadcasting (DVB) Interaction Channel for Satellite Distribution Systems. 2002-11.
- [28] Harney H, Schuett A, Colegrove A. GSAKMP Light. *IETF Internet Draft, work-in-progress, draft-ietf-msec-gsakmp-light-sec-01.txt*, Jul 2002, expires Dec 2002.
- [29] Baugher M et al. The group domain of interpretation. *IETF Draft, work-in-progress, draft-ietf-msec-gdoi-06.txt*, Oct. 2002, expires Apr. 2003.
- [30] Wallner D, Harder E, Agee R. Key management for multicast: issues and architectures. *IETF RFC2627*, June 1999.

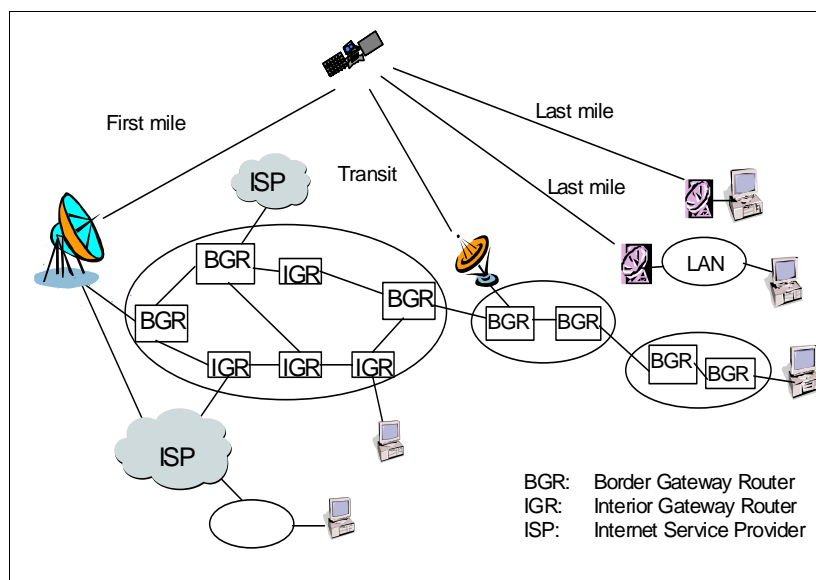


Figure 1. Satellite IP multicast connections




Figure 2. The GEOCAST system, illustrating star and mesh topologies

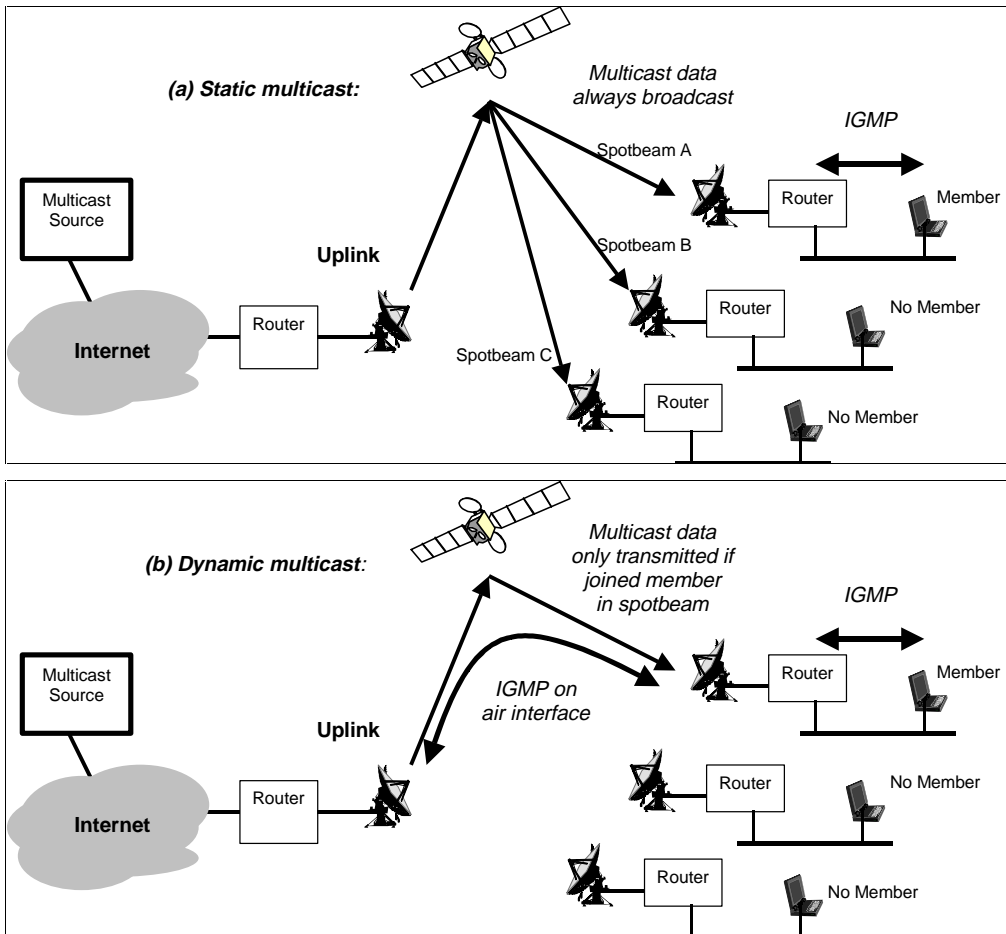


Figure 3. IGMP over satellite: (a) static and (b) dynamic multicast




Figure 4. Multicast routing flooding: two approaches

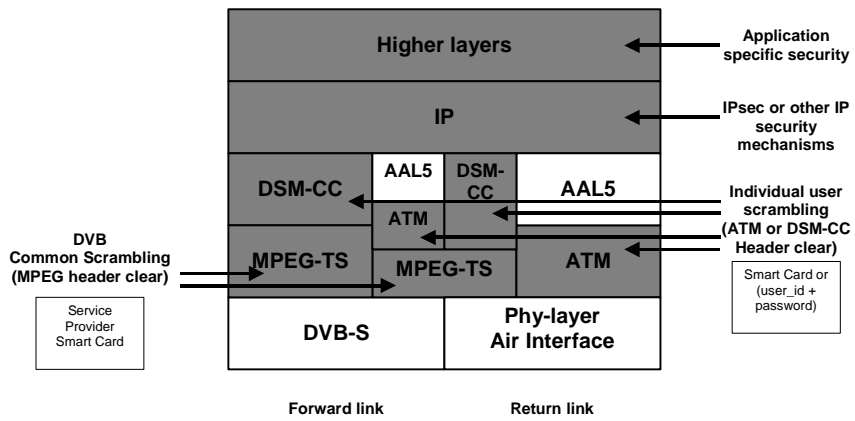


Figure 5. IP stack and security in DVB-RCS