IP ROUTING ISSUES IN SATELLITE CONSTELLATION NETWORKS

L. WOOD,[*1] A. CLERGET,[2] I. ANDRIKOPOULOS,[1] G. PAVLOU[1] AND W. DABBOUS.[2]
[1]Centre for Communication Systems Research (CCSR), University of Surrey, Guildford, United Kingdom.
[2]RODEO group, INRIA Sophia-Antipolis, France.

This is a preprint of an article accepted for publication in the International Journal of Satellite Communications Special Issue on Internet Protocols over Satellite, vol. 18 no.6, Nov/Dec 2000. Copyright © 1999, 2000 John Wiley and Sons.

SUMMARY

The growth in use of Internet-based applications in recent years has led to telecommunication networks transporting an increasingly large amount of Internet Protocol (IP)-based traffic. Proposed broadband satellite constellation networks, currently under development, will be required to transport IP traffic. A case can be made for implementing IP routing directly within the constellation network, in order to transport IP traffic well and to provide good support for IP multicast and for emerging IP-based Quality of Service (QoS) guarantees. This paper examines strategies for implementing and operating IP routing effectively within satellite constellation networks, given known constraints on the constellation resulting from satellite mobility, global visibility, routing and addressing.

KEY WORDS: satellite constellation networks; Internet Protocol (IP); routing; tunnelling; Multi-Protocol Label Switching (MPLS); Border Gateway Protocol (BGP); Quality of Service (QoS); multicast.

1 INTRODUCTION

Satellite-based networking has developed in complexity over the years, rising up and building upon established work at the various networking layers as described by the ISO OSI Reference Model.[1] Networking using satellites began by using simple transparent bent-pipe repeaters in geostationary orbit, where uplinked signals were amplified, frequency-shifted and broadcast down to a large ground area. Sharing of this broadcast physical and data-link layer capacity led to the introduction of increasingly complex media-access control (MAC) schemes to use capacity effectively, most notably with slotted aloha and its variants for use with very small aperture terminal (VSAT) networks.[2]

The development of multiple spotbeams per satellite led to on-board switching and MAC, with control of capacity allocated via circuits and a Logical Link Control (LLC) sublayer.

Development of direct radio or laser inter-satellite links (ISLs) between satellites and the design of constellations utilising ISLs, such as *Iridium*, *Teledesic* and *Spaceway*, has led to consideration of dynamic adaptive routing algorithms for communication across a toroidal mesh of ISLs between multiple satellites; the space segment has now reached the network layer, and satellites in such constellations must support onboard routing as well as onboard switching. In this case, the satellite constellation itself is a true network; in conjunction with its terrestrial gateway stations it forms an *autonomous system* (AS).

Over the same period that satellite-based networking has been developing, the Internet has been developing and growing in size. The TCP/IP protocol suite used on the Internet has become established as the most popular method for computer network communication in the world. In this paper, we examine how IP routing can be implemented in a complex satellite constellation network, in order to support TCP/IP-based services well.

Section 2 provides overviews of the history of TCP/IP over satellite, of constellation networks, and of logical approaches to routing in mobile constellation networks. Section 3 gives reasons for considering IP routing in constellation networks, discusses common objections to implementing IP routing onboard satellites, and examines the approaches taken by the proposed commercial constellations. Section 4 provides an overview of methods with which the constellation network can route Internet traffic. Section 5 describes how IP routing for Internet traffic can coexist together with other routing of non-IP traffic on an ATM-based infrastructure, via the use of Multi-Protocol Label Switching (MPLS). Section 6 concludes the paper.

* Correspondence to Lloyd Wood, CCSR, University of Surrey. Email L.Wood@surrey.ac.uk

2 OVERVIEW

2.1 A brief history of TCP/IP over satellite

The Transmission Control Protocol/Internet Protocol (TCP/IP) suite has been carried transparently over satellite ever since experiments were first conducted with SATNET in the 1970s,[3] and TCP/IP implementations have been shown to work well over satellite links. However, the long propagation delay to satellites in geostationary earth orbit (GEO) has imposed limitations on interactive applications and on existing TCP implementations.

Work on large windows[4] and selective acknowledgements[5] has been designed to overcome TCP's problems with paths that exhibit high bandwidth-delay products, such as links over geostationary satellites.

Current TCP congestion control algorithms[6] can mistake bursty satellite channel errors (which can be the result of how data-link layer coding choices perform under poor conditions) for network congestion. This leads to sub-optimal use of available satellite link capacity when recovering from errors, due to congestion avoidance decreasing TCP's sending rate dramatically, followed by a slow return to the previous transmission rate. This is something that really results from a lack of explicit congestion notification to distinguish between network congestion and link errors. Tweaking congestion control algorithms to improve performance in the satellite environment cannot compensate for this lack of information on the real cause of the problem.

TCP/IP works over even geostationary satellites, has done so in the past, and will continue to work over satellite in the future. However, there is more to the IP family of protocols than simply being able to support end-to-end communication using TCP; for example, interactive applications may use the connectionless User Datagram Protocol (UDP), the Real-time Transport Protocol (RTP) or communicate between groups of endhosts using IP multicast.

The trend towards complex switching and routing onboard satellite, and the network topologies created by an orbiting constellation of broadband satellites with ISLs, have created the need for constellation networks to be able to route traffic internally over multiple satellites between sources and destinations on the ground. Although unicast transmissions, such as those for TCP virtual circuits, can be supported end-to-end across any proprietary network by tunnelling, implementing support for other protocols in the TCP/IP suite, particularly multicast, is less straightforward, requiring routing support in the new constellations that are described below.

2.2 Constellation network topology

Low-earth orbiting (LEO) satellite constellations have been proposed, using orbits much lower than the geostationary orbit, in order to give global coverage, more frequency reuse of the limited earth-space communication frequencies available, and higher system capacity as a result of this frequency reuse; the decrease in propagation delay when compared to GEO is an added bonus, but can be insignificant or not quantifiable for a number of applications, such as automated file transfer.

Use of non-geostationary orbits results in the need for satellite-to-satellite handover even for fixed earth stations. Use of ISLs in the constellation leads to a complex orbiting mesh network topology, where permanent ISLs are established between satellites following each other in the same circular orbital plane. Semi-permanent ISLs can be set up between satellites in neighbouring planes, but must be broken and reestablished at the highest latitudes of each orbit as the planes cross (an example topology is shown in Figure 1).

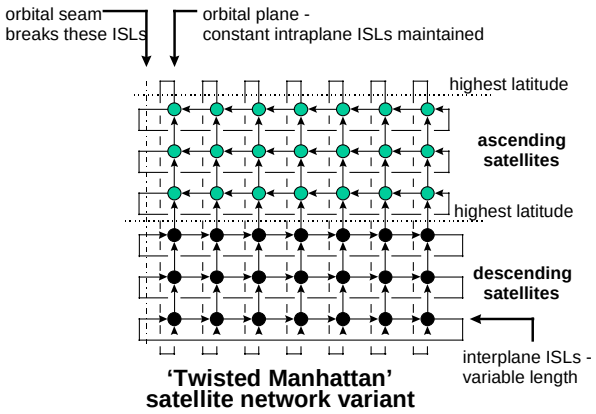

ISLs have added direction to illustrate crossing of orbital planes at highest latitudes, where neighbours swap places.

Figure 1: ISL topology of LEO constellation

Use of a polar (Walker star) constellation geometry[7] leads to a cylindrical mesh network with an orbital 'seam' between counter-rotating planes where cross-plane ISLs may not be viable due to the satellites' high passing speeds; use of a rosette (Ballard) geometry[8] leads to a fully-toroidal mesh network[9] that is a variant of the class of networks known as the Manhattan Network[10]. In both cases the network is a mesh, providing multiple redundant ways in which the ISL links can be traversed from one satellite to another.

Routing algorithms are needed to determine the best way to traverse the mesh; a flexible packet-based, rather than static circuit-based, routing approach can take advantage of the redundancy inherent in this mesh.

2.3 Satellite mobility and routing issues

Due to their low altitude (typically 700 to 1400 km), LEO satellites move at rapid speeds relative to the ground terminals. Speeds at over 25,000 km/hour, with satellite visibility times of only a few minutes before handover occurs to another satellite, are the norm.

This high mobility leads to a rapidly and regularly-changing network topology, and raises numerous issues for the networking layer with respect to routing.

Terrestrial Internet routing protocols, such as Open Shortest Path First (OSPF)[11] and Routing Information Protocol (RIP),[12] rely on exchanging topology information when network connections are established or changed. In LEO constellations, this topology information quickly becomes obsolete and must constantly be refreshed with new information. The overhead of regularly providing this information is an obstacle to considering satellites as conventional Internet routers.

However, the topology of these constellations exhibits interesting and useful properties:

- Predictability;
- Periodicity in the space segment;
- Regularity;
- Constant number of satellite network nodes.

To perform routing in this highly dynamic but tractable context, several strategies have been proposed:

2.3.1 Path maintenance via Virtual Topology Routing

The idea behind Dynamic Virtual Topology Routing[13] is to exploit the periodic and predictable nature of the constellation topology.

Time intervals $[t_0=0,t_1]$, $[t_1,t_2]$, ... , $[t_{n-1},t_n=T]$, where T is the period, are chosen so that:

- Over an interval $[t_i,t_{i+1}]$, the topology can be modelled as a constant graph G_i, i.e. link activation and deactivation take place only at discrete times t_0, t_1, ..., or t_n.
- The interval $[t_i, t_{i+1}]$ is small enough to consider the costs of individual ISLs as constant over this time interval. The costs of these links could be computed from a function of inputs such as distance between the satellites, duration before link deactivation, geographic position, or other factors – assigning higher cost to high-latitude ISLs with a short time before deactivation, for example.

Over these time intervals, the 'instantaneous' topology, G_i, is fixed. Optimal shortest paths and alternate paths can be established across the network graph between all pairs of satellites, using well-known methods such as the Dijkstra shortest-path algorithm. These optimal paths can be calculated in advance for the topology on the ground and then uploaded to all satellites via broadcast command.

It is also possible to add an optimization procedure to choose among alternate paths between two satellites in order to minimize the number of satellite-to-satellite handovers required over the period.[14]

This path-based approach makes it possible to attempt to hide the mobility of satellites from standard connection-oriented network protocols, such as ATM, that may be running over the constellation, simplifying their view of the constellation and thus their routing.

2.3.2 The virtual node concept

The virtual node concept[15] aims to exploit the regularity of the constellation's topology. Again, the goal is to hide the mobility of satellites from routing protocols running over the constellation.

In this scheme, information concerning terrestrial constellation users, and how to communicate with them, is state that relates to a region of the Earth and is maintained in a fixed

position relative to the surface of the Earth. Constellation users communicate with the virtual entity containing state pertaining to them: this is the virtual node. This virtual node is embodied at any given time by a satellite, and a virtual network of these nodes is embodied at any time by the satellite constellation.

As the satellites move and as users perform handovers, state, such as routing table entries or channel allocation information, is transferred continuously from one satellite to another. Routing is performed in the fixed virtual network, by using a common routing protocol.

Since we are considering carrying IP traffic in the constellation and the use of connectionless packet-based routing protocols, we will consider this path-independent strategy further.

2.3.3 Strategies dependent on topology

These strategies use proprietary routing protocols that have explicit knowledge of the constellation topology and the satellite mobility. Such protocols require that there is always a path between two communicating ground hosts, and that routing is loop-free.

Each proprietary protocol will be very specific to the design of a certain type of constellation. The Footprint Handover Routing Protocol[16] is a simple example of such a protocol for polar Walker star constellations.

3 IP ROUTING ONBOARD SATELLITE

3.1 Reasons for considering IP routing in constellations

Assuming that the constellation network has ISLs, and that it will be expected to carry IP traffic, there are a number of compelling reasons for wanting to support IP routing of that same IP traffic in the constellation network's space segment:

3.1.1 IP multicast

One-to-one end-to-end connectivity across an internetwork can be accomplished using a virtual circuit. For group applications, where more than two users simultaneously communicate with each other, the number of circuits required increases rapidly with the number of users in the group: ½[n(n-1)] bidirectional virtual circuits for n users.

To prevent applications from needing to know about all users in the group or being responsible for maintaining all these virtual circuits, and to decrease network load, we require multicast.

Multicast allows a source to simultaneously send data to all users on the internetwork interested in receiving the data, but in a more efficient manner than establishing multiple virtual circuits or simply flooding the entire internetwork with unnecessary broadcast packets.

Support for IP routing permits straightforward support for IP multicast, which allows management of multi-way IP communications using the Internet Group Management Protocol (IGMP)[17] and related multicast protocols.

The set of all end hosts interested in a multicast communication forms a multicast group. Management of membership of local end hosts belonging to this group on a subnetwork is the responsibility of the local multicast router, using IGMP. End hosts not in the group do not see the data (perhaps because there is no need for the data to be sent across their subnetwork, as there are no group members in that subnetwork) or, if they do see it, discard it.

To communicate the data efficiently to all users across the Internet belonging to the multicast group, the internetwork must set up a spanning tree connecting networks where multicast routers have indicated interested users exist. Messages can be sent along this spanning tree and replicated at tree branches. How this tree is established depends upon the type of information being communicated, and the expected scale of the group. There are two basic sets of assumptions in IP multicast protocols:

- *Source-based tree* multicast protocols are data-driven, in that construction of the multicast spanning tree begins top-down from the source outward, and data on the state of the tree is flooded to all routers before being pruned back by routers on subnetworks with no interested members requesting the tree no longer reach them. The multicast trees constructed allow data to travel in one direction, from source to group, emulating broadcast; the Distance Vector Multicast Routing Protocol is an example of this mode.[18] The initial flooding assumes that potential group members are densely distributed throughout the internetwork, i.e. that many subnetworks contain at least one

group member and will be interested in receiving the communication, and that internetwork capacity is plentiful.

- *Core-based tree* multicast protocols have receiver-initiated multicast spanning trees, where a router becomes involved in a branch of a multicast distribution tree only when one of the hosts on its subnetwork requests membership by issuing a join message. There may be one or more central core routers which receive join and leave messages.[19] The lack of any initial flooding and the assumption of constrained capacity and fewer interested members, sparsely distributed, means that this scales better for internetworks.

An ISL-using satellite constellation network communicating with the terrestrial Internet possesses considerable capacity in the space segment due to its broadband microwave or laser ISLs, and there will be considerable capacity in the fibre-based terrestrial Internet ground segment. The throughput constraint lies in the limited capacity of the earth-space air interface between the two.

Being able to duplicate IP multicast packets in the ISL network for redistribution to all communicating ground parties involved at remote terrestrial terminals makes best use of the earth-space interface, as no unnecessary packet duplication or repetition in multiple virtual circuits needs to occur across the limited earth-space interface.

Without IP routing and native support for IP multicast, implementing any support for IP multicast becomes increasingly problematical, as the IP multicast group and tree will need to be projected with difficulty onto other network layers and routing paradigms, where IP multicast routing functionality must be duplicated.

3.1.2 Supporting IP QoS

The traditional network service on the Internet is best-effort datagram transmission. IP packets are sent from a source to a destination without any guarantee that the packet will be delivered.

For traditional two-way data applications which are elastic in nature in that they tolerate packet delays and packet losses, this best-effort model is satisfactory, and any necessary reliability can be implemented without redundancy at the end-points e.g. via acknowledgements in TCP.[20]

However, the emerging real-time applications have very different characteristics and requirements than data applications. They are less elastic, less tolerant of delay variation and need specific network conditions in order to perform well. The Internet protocol architecture is being extended to provide support for real-time services by adding Quality of Service (QoS) models to meet these application requirements.

There are two architectures that are being defined in this context: *Integrated Services* and *Differentiated Services*. Support for IP routing within the satellite constellation would make it possible to support IP QoS via one of these architectures.

3.1.2.1 Integrated Services

The primary goal of the Integrated Services architecture and QoS model is to provide IP applications with end-to-end 'hard' QoS guarantees, where the application may explicitly specify its QoS requirements and these will be guaranteed and met by the network.[21]

For this to be accomplished, the Resource Reservation Protocol (RSVP) is used to signal the resource requirements of the application to the routers situated on the transit path between the source and destination hosts.[22]

The Integrated Services architecture supports two new classes of service, in addition to the existing best-effort class:

- The *Guaranteed Service* guarantees both delay bounds and bandwidth availability, setting a maximum queuing delay.
- The *Controlled Load Service* approximates the end-to-end behaviour provided by best-effort service under unloaded conditions. The network ensures that adequate bandwidth and packet processing resources are available to handle the requested level of traffic.

The major drawback of Integrated Services is that the amount of state information, which is required to be maintained per node, is proportional to the number of application flows, and does not scale. As resource requirements must be negotiated over a set path or a set spanning tree, where the routers in the path or tree maintain soft state pertaining to the flows

passing through them, support for the regularly-changing paths resulting from satellite mobility would act to invalidate RSVP guarantees, as the path between endpoints for which an RSVP guarantee is set up does not remain constant. RSVP renegotiation when this happens would be extremely undesirable.

RSVP is not suited for deployment on high-bandwidth backbones or on transit networks due to its reliance on per-flow state and per-flow processing. Aggregation of flows will be necessary to make large-scale RSVP tractable.

3.1.2.2 Differentiated Services

The Differentiated Services (DS or diff-serv) architecture has been proposed to overcome perceived limitations of Integrated Services.[23] DS allows IP traffic to be classified into a finite number of priority and/or delay classes. Traffic classified as having a higher priority and/or delay class receives some form of preferential treatment over traffic classified into a lower class.

The Differentiated Services architecture does not attempt to give explicit 'hard' end-to-end guarantees. Instead, at congested routers, the aggregate of traffic flows with a higher class of priority has a higher probability of getting through. Traffic with a marked delay priority is scheduled for transmission before traffic that is less delay-sensitive.

The information needed to perform actual differentiation in the network elements is carried in reserved bits in the Type of Service (TOS) field of the IPv4 packet headers, or the Traffic Class field of the IPv6 packet headers. This is referred to as the differentiated services codepoint (DSCP).[24]

Since the information required by the buffer management and scheduling mechanisms is carried within the packet, no complex per-flow signalling protocols are required. As a result, the amount of state information, which is required to be maintained per node, is only proportional to the small overall number of service classes and is not proportional to the large number of application flows.

The Differentiated Services architecture is composed of a number of functional elements, namely packet classifiers, traffic conditioners and per-hop forwarding behaviours (PHB). A PHB describes the externally-observable forwarding behaviour of a differentiated services node, as it is applied to a collection of packets with the same DSCP that are traversing a link in a particular direction.

Each service class is associated with a PHB. PHBs are defined in terms of behaviour characteristics relevant to service provisioning policies, and not in terms of particular implementations. PHBs may also be specified in terms of their resource priority relative to other PHBs, or in terms of their relative observable traffic characteristics. These PHBs are normally specified as group PHBs and are implemented by means of buffer management and packet scheduling mechanisms.

According to the basic differentiated services architecture definition, these elements are normally placed in ingress and egress boundary nodes of a differentiated services domain and in interior DS-compliant nodes. However, it is not necessary for all the elements to be present in all the DS-compliant nodes; that depends on the functionality required at each node.

At each differentiated services user/provider boundary, the service provided is defined by means of a Service Level Specification (SLS). The SLS specifies the overall performance and features which can be expected by a customer or another network.

3.1.2.3 IP QoS in the constellation

As we are considering a satellite constellation that is effectively a high-bandwidth mobile backbone, the Integrated Services model simply does not scale for use within the constellation. Aggregation of RSVP-specified flows would be extremely complex to implement, and mobility would be difficult to overcome even with use of virtual nodes.

We should consider supporting RSVP at the edges in the border routers at the gateway earth stations, and the more scaleable differentiated-services model within the constellation.

From a differentiated-services viewpoint, as the satellites in a constellation can be expected to be mass-manufactured and identical, with identical routing functionalities, we can consider the constellation as both a single Differentiated Services (DS) Domain and as a single DS Region, where a common, single, set of per-hop behaviour (PHB) groups is implemented within the routers in every satellite and in every ground station. PHBs specific to the capabilities of the satellite constellation can be defined.

Marking the DSCP in order to select an appropriate PHB will be carried out at the border gateway routers, which are also DS ingress and egress boundary nodes that will undertake any necessary traffic conditioning.

As a single region, domain and PHB behaviour set, there should be no need for additional PHB mapping within the constellation between satellites. This will greatly simplify meeting service-level specifications in comparison to other, less homogeneous, diff-serv-capable networks.

3.2 Problems with IP routing onboard satellite

This section describes reasons often stated for not considering IP routing onboard satellite, and argues that these reasons are not insurmountable.

3.2.1 Variable-size IP packets

A common misconception is that, as the satellite air interface must allocate channel capacity in some predefined manner via FDMA/TDMA, fixed datagram sizes are needed to fit neatly into the frame structures for the allocated slots in the wireless channel. As IP is known to have variable-length packets, an objection to the use of IP is raised.

It is possible to fit IP packets into any fixed-length frame structure by the use of either explicit IP-level fragmentation, where the packet is broken into sections, each with a fragmentation identifier (ID), small enough to transmit across the fixed-length interface as IP packets themselves, or by implicit lower-level fragmentation, where IP is broken up in order to be carried by a MAC-level or a tunnelling protocol, discussed further in section 4.1. Padding can be used where appropriate to fill up frame structures not completely filled with all or part of an IP packet payload.

IP-level fragmentation is generally undesirable, and its occurrence can be minimised by use of path message transfer unit (path MTU) discovery[25] and the setting of a maximum MTU size for IP packets. Although use of path MTU discovery with IPv4 has caveats beyond the scope of this paper, its use is mandatory with IPv6.[26]

3.2.2 Routing table management

Routing table size and complexity is often cited as an obstacle to performing IP routing onboard satellite.

For performance equivalent to terrestrial equipment, space-qualified computing hardware is generally considerably more difficult and expensive to produce, and satellite computing performance (and thus routing performance) can be expected to lag behind equivalent terrestrial performance at any point in time.

Once a satellite is launched it cannot be upgraded for the duration of its expected lifetime, meaning that the satellite performance can be expected to fall increasingly behind terrestrial performance, and must be designed with a margin to meet expected needs at the end of the satellite lifetime. This is a clear argument for placing as much as possible of the complexity of the satellite constellation network in the ground segment in order to future-proof it – even to the extent of limiting the space segment to nothing above the data-link layer and rejecting use of ISLs.

The space environment, with radiation and temperature variations, is harsh on processors. This limits available processing capacity in comparison with equivalent terrestrial Internet routers. The available power budget is also limited.

Any single LEO satellite is unlikely to be able to hold information on how to route towards the physical addresses of all of the necessary connected networks in the world, much less keep constantly handing its entire routing table for the Internet over to the next satellite as the satellites move and assume new virtual nodes or management of earth-fixed cells. This creates a scalability problem for on-board routing tables and processing, and is the problem commonly referred to by people with an understanding of terrestrial Internet routing protocols. The constellation would be overwhelmed with information about the terrestrial Internet. The satellites can be seen to function better if they do not need to know about terrestrial Internet addresses or about terrestrial routing.

The movement of the onboard IP routers in the non-geostationary satellite constellation can also create a similar scaling problem if information concerning their continuous motion is propagated continuously as routing updates within the terrestrial Internet. The terrestrial

Internet functions well without knowing about the motion of satellites in the space segment.

In both cases, the problem is not with IP routing, but with full integration and the rate at which large amounts of routing information must be updated for both the terrestrial Internet and the constellation network. Keeping routing updates from propagating from one to the other is a way to prevent these problems.

There is also a secondary problem with the growth of backbone routing tables as the Internet increases in size. This continual growth in routing table size means that the tables might eventually outstrip the availability of the satellites' onboard routers to hold them, leading to loss of service to parts of the terrestrial Internet. Field upgrades to add more table memory are not feasible for satellites; again, excess table capacity must be designed in for the satellites' lifetimes, in order to meet the expected size of routing tables at the end of the satellites' expected lifetimes before replacement. This excess capacity is undesirable from an engineering viewpoint, as it relies on matching the expected satellite lifetimes with uncertain projections of routing table sizes. The satellite network itself is relatively fixed in size, requiring a fixed-size table; removing Internet-related information removes this routing problem.

In all cases, the limited earth-space capacity also benefits if routing updates do not need to be unnecessarily propagated across it, and the amount of routing information and state held in the satellites, and resulting size of the routing tables, needs to be minimized.

These goals can be achieved by separating and isolating satellite and Internet routing updates to their respective routing realms, using one of the methods discussed later in section 4.

3.2.3 Speed of routing vs. switching

IP routing involves examining packet headers for a global destination address and then executing a table lookup for the correct forwarding action to take for the packet, rather than simply transmitting the packet from one switched interface to another. As a result, IP routing is generally perceived as requiring more processing power than, and being slower in operation than, simpler switches that do not need to consider anything beyond the local interfaces and that do not need to examine global addresses in headers. As on-board processing capabilities are constrained by the environment, this is cited as a reason why IP routing is not suitable onboard satellite.

However, continual advances in processing power, better lookup algorithms and the move from simple bus-based routers to crossbar and shared-memory switching fabric designs have acted to improve IP routing performance.[27] This raises the bar on what is feasible for onboard processing.

Work on 'shortcut' IP switching techniques such as MPLS demonstrates increased throughput and routing performance with shorter queuing delays and fewer local state overheads.[28]

3.3 The commercial constellations

Since IP multicast and IP QoS are still being defined at present and are not yet widely implemented in internetworks, while a number of commercial broadband constellations are finishing their design stages and nearing production, it is safe to assume that implementing support for IP multicast and for IP QoS does not figure strongly in the designs of these commercial constellations.

Of the most visible of the proposed broadband constellations, Hughes' *Spaceway* and Lockheed Martin's *Astrolink* are fixed GEO constellations adopting ATM-based switched communication across ISLs and in the earth/space interface, although they will be using custom ATM MAC/LLC layers and custom signalling.

Of the LEO constellations, *Teledesic* is understood to have designed its own custom-designed protocols over ISLs and in the earth/space interface, while Alcatel's *Skybridge* is taking a ground-based ATM approach without the use of either onboard routing or ISLs.

These commercial constellation networks can be expected to support end-to-end communication of unicast IP traffic via tunnelling, as described in section 4, but support for IP beyond that is an open question.

Given the use of either a custom protocol or ATM with AAL encapsulation, implementation of IP QoS, IP multicast, or of future enhancements to the IP architecture in the proposed commercial constellations looks to be problematic.[29]

4 IMPLEMENTING ROUTING IN THE CONSTELLATION NETWORK

There are several different approaches to implementing routing in the constellation network.

If IP routing is considered desirable, it is also necessary to avoid propagating IP routing updates between the separate routing realms of the constellation network and the Internet.

The routing approaches discussed here are:

- *tunnelling* over another network protocol or over IP in the constellation network. End-to-end communications are supported well, and the clear separation of the tunnelling network layer and the external network prevents the need to communicate routing updates between the two;
- *network address translation* (NAT) using twice NAT, which separates internal and external IP addressing and routing;
- *exterior routing protocols*, where border routing is managed at the edges of the network using the Border Gateway protocol (BGP), in order to use different internal routing protocols or network layers while controlling internal and external propagation of routing updates.

4.1 Tunnelling

4.1.1 What is tunnelling?

Tunnelling is often used to route packets in one network through an intermediate network that belongs to a different routing realm and that can have a differing network layer. Packet formats, addressing space, and routing paradigms in the two networks may be *entirely* different.

When IP is tunnelled, a virtual IP hop (a tunnel) is created between the two IP-capable routers at the borders of the intermediate network. When entering the tunnel, the IP packet, including the IP header, is sent as payload data to the other side of the tunnel using the network layer of the intermediate network. Steps in tunnelling, shown in Figure 2, include:

- Possible implicit fragmentation of an IP packet to fit in the payloads of packets used in the intermediate network.
- Building local packets or cells suitable for routing or switching across the intermediate network using its routers or switches by the addition of appropriate header information.
- Setting the destination address in these packets or cells to the internal network address of the IP-capable router at the other end of the tunnel.

When leaving the tunnels, the IP packets are reassembled as required and forwarded along the next IP hop.

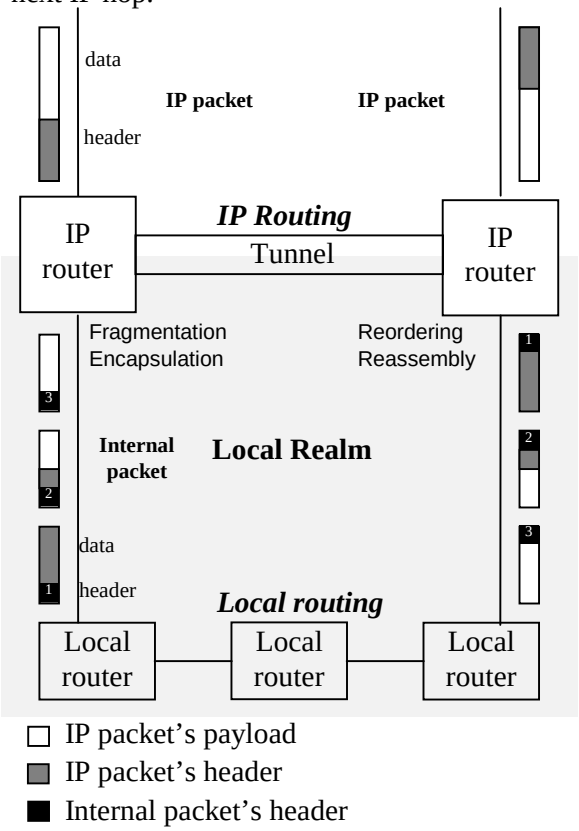

Figure 2: Tunnelling

Tunnelling is often used as a transitional measure to support new functionality added to parts of today's Internet without requiring the entire Internet to be upgraded at once to support new features for those who wish to adopt them.

One of the best examples of tunnelling in action in this way is the MBone,[30] a virtual network across the Internet interconnecting subnetworks where multicast routing is supported. Multicast-capable routers (or end hosts running software to act as routers) forward packets to neighboring multicast routers within the MBone topology. A tunnel is configured between nearby routers that are separated by non-multicast-capable routers, and multicast packets are sent within these tunnels.

Another example of transitional tunnelling, using IP-in-IP tunnelling, is the 6Bone, a testbed virtual network that is assisting in the deployment of IPv6 over IPv4 as part of the Internet's gradual move to IPv6.

4.1.2 Tunnelling over different protocols

4.1.2.1 IP over ATM

As this paper is written, many satellite constellation operators and manufacturers have focused on ATM as the network protocol for the constellation, but with use of proprietary ATM signalling protocols and MAC layers. Support for ATM and interworking with ATM networks is a commercial goal.

ATM virtual path connections (VPCs) can be maintained between all pairs of satellites, using for example the Dynamic Virtual Topology Routing concept. All the ATM virtual channel connections (VCCs) that share the same pair of satellite entry and exit points can be aggregated into the same VPC. Switching is done onboard only according to the VPC label.

If an ATM service is provided to interconnect two constellation users, IP packets can be tunnelled and carried by ATM cells, using for example classical IP-over-ATM encapsulation.[31]

4.1.2.2 IP over a proprietary protocol

A proprietary network layer and routing protocol can be specifically optimized for the constellation. Such a protocol can avoid transmitting unnecessary routing information while propagating other useful network-specific information such as internal delay, expected traffic load or instantaneous traffic load.

This appears to be the approach adopted by the proposed *Teledesic* constellation network.

4.1.2.3 IP over IP

It may sound curious to suggest the tunnelling of IP over IP,[32] but this approach does allow separate addressing, separate routing realms and avoids propagation of routing information between the constellation network and the terrestrial Internet. This approach has the advantage of using existing IP routing protocols, with the possibility of relatively straightforward support for IP routing features such as IP multicast or IP QoS, unlike tunnelling over non-IP protocols.

However, IP-in-IP encapsulation imposes a header overhead. One of the few advantages of NAT, discussed in section 4.2, is that NAT avoids this header overhead

A second disadvantage of supporting only IP routing in the constellation is the tunnelling of non-IP communications over the IP layer. This is unacceptable e.g. for ATM, and an alternative approach is necessary if the constellation is to support more than just IP.

4.1.3 Tunnelling in the satellite constellation

To isolate the constellation network's routing realm from external networks or routing realms, or to send IP traffic across a constellation network where IP routing is not implemented or is not supported in the network layer, we can create tunnels across the network that link IP-capable entities on the ground, namely:

- Isolated ground hosts using the constellation for connectivity to the terrestrial Internet or to other isolated ground hosts or networks.
- Small routers using the constellation network to interconnect a local area network (LAN) to the terrestrial Internet or to other ground hosts or networks.
- Large border gateways interconnected with the rest of the Internet, through which traffic from the previously-listed entities would travel to reach the terrestrial Internet.

Seen from the IP level, the network topology is as illustrated (Figure 3). The border gateways, routers, and ground terminals, shown in black, are where tunnelling would occur.

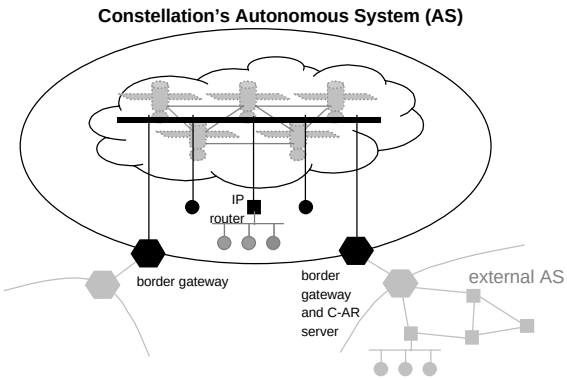

Figure 3: Tunnelling in the constellation[33]

As tunnels must be established between all pairs of hosts, these IP ground entities create a fully

connected graph, thus creating a virtual network across the satellite network.

For destinations outside the constellation, small routers and border gateways can reach the egress border gateway in one hop over the constellation, whereas isolated ground constellation users need to send their packets to an IP routing entity (e.g. a border gateway), that will tunnel the packet to the egress point.

One significant difference between the use of tunnelling across the constellation network and the use of tunnelling described earlier for the MBone and 6Bone is that for the constellation network the tunnelling is a permanent, rather than a transitional, measure with no future benefits.

4.1.4 Constellation Address Resolution Servers

To send a packet through a tunnel from one edge router to another, it is necessary to know the constellation address (the address in the constellation realm) of the communicating peer on the other side of the tunnel. This address could be pre-configured in the tunnelling entity, but since the number of possible peers is potentially very large and the virtual network is a fully-connected graph, this approach does not scale well with size or adapt well over time as new constellation user networks join the constellation. An on-demand strategy for retrieving these constellation addresses appears more reasonable.

On an Ethernet local network, an IP entity that wants to send a packet to another local IP entity first needs to retrieve its Ethernet address using an Address Resolution Protocol (ARP).[34] Here, we can use a similar strategy, except that we do not resolve the address by broadcasting a request, but by interrogating a Constellation Address Resolution Server (C-ARS), in a similar fashion to ATM-ARP. All that a ground IP edge router needs in order to communicate with other IP hosts at the edges of the constellation's autonomous system is the destination IP address, and the constellation address of a C-ARS; this would be situated in a gateway station that ideally also controls the satellites and has detailed knowledge of the satellite constellation.

To avoid concentrating all the address resolution traffic around a single C-ARS and to provide redundancy, we may have multiple C-ARS servers situated in other ground stations that communicate network updates to one another.

Having the C-ARS servers all linked over the satellite network may result in dangerous failure modes. A possible option would then be to have a core of primary C-ARS synchronized by sharing state over dedicated redundant terrestrial links. Additional secondary C-ARS could be added in remote regions and synchronized by communicating over the constellation itself with a primary C-ARS. A host whose Address Resolution Request failed would switch to a safer primary C-ARS.

Allocation of IP ground hosts to C-ARS servers can be dynamic, based on geographical position, or static, which leads to inefficiency in the case of mobile ground hosts.

4.1.5 The constellation realm

As discussed in section 2.3, a specific routing scheme can hide the high mobility of the constellation from the ground users and from the rest of the Internet.

In particular, we want to avoid generating huge amount of routing traffic between the separate routing realms of the Internet and the constellation, while still propagating necessary updates concerning route changes internally.

In the constellation realm, we can expect to exchange very little information on that dynamic topology, since the topology changes are mostly predictable.

The necessary information for routing, i.e. the position of a node within that topology, can be deduced from its constellation address.

For the satellite interfaces:

- Satellite ID, ISL Interface ID
- or Virtual node and interface ID

For the ground host interfaces:

- Fixed geographical position
  - Earth-fixed cell ID, related to latitude and longitude,
  - or Virtual Node ID.
  - or Moving geographical position
    - Moving cell ID, current satellite and downlink interface ID.
- MAC Address (code, time slot, and/or frequency) or Host ID

A simple mapping between address and position within the topology allows us to use routing

protocols that use very little network capacity to exchange topology information.

4.1.6 Advantages of tunnelling

- Tunnelling can allow us to adapt the tunnelling network layer and routing protocols inside the constellation network to the needs and constraints of that network.
- Tunnelling is a simple solution to separate routing updates and addressing in the constellation network from routing updates and addressing in the Internet.

4.1.7 Disadvantages of tunnelling

- Tunnelling imposes some processing overhead – dealing with headers, encapsulation and fragmentation.
- Tunnelling can give a false picture of the number of hops between two points. All tunnels appear the same length – a single virtual IP hop – to the packets being tunnelled through them and to routers exchanging routing information over them. This can result in the constellation network becoming the preferred route over alternative shorter terrestrial routes to the same destination, as the true number of hops within the constellation is not visible. The time-to-live (TTL) hop-count field in the IP packet header is not decremented in the tunnel because the header is encapsulated as data.
- Mapping IP QoS and IP multicast onto the tunnelling network and supporting them in that network is a non-trivial problem.
- Events in the constellation network are not visible to the terrestrial Internet, making it difficult to notify the Internet about network state that may have an impact on QoS. Notification to the Internet of congestion in the constellation network via Explicit Congestion Notification (ECN)[35] would be difficult without complex handling of notification events between the tunnelling layer and the terrestrial Internet.

4.2 Network Address Translation

Network Address Translation (NAT) is the term for techniques used in private IP networks to manage internal address space by separating it from the global Internet address space. NAT translates the internal-realm addresses in every IP packet to new addresses suitable for use in the external realm.

NAT has been used to avoid having to renumber a private network when topology outside the network changes, for firewalls, and for a variety of other reasons. As an increasingly popular technique, its use is being documented by the IETF NAT working group.[36]

4.2.1 NAT in the constellation network

The satellite constellation network can be viewed as a private network with a single operator controlling both the space segment and the terrestrial gateways interfacing to the terrestrial Internet. The external topology of the terrestrial Internet changes from moment to moment as far as non-geostationary satellites in the network are concerned, thanks to their orbital motion; if the satellites support IP routing, routing updates become a problem, as previously discussed.

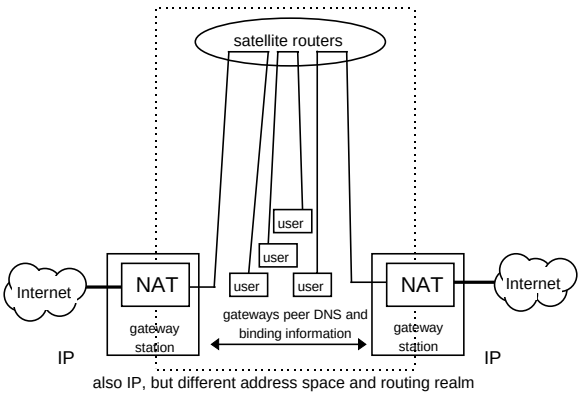

Figure 4: NAT in the constellation network

NAT can be considered as a useful way of separating constellation network addressing from global Internet addressing to provide separate address realms. It can remove the need for propagation of routing updates between the constellation and the terrestrial Internet.

NAT would be implemented in the terrestrial gateway stations interconnecting the satellite constellation network with the terrestrial Internet. These gateways translate between internally-visible addresses of constellation network users and externally-visible addresses associated with that gateway. The gateway is viewed as the endpoint of all communications

for routers and users internal and external to the constellation (Figure 4). The gateway provides transparent routing by straddling and having knowledge of both addressing realms.

It would be necessary to propagate translations and address bindings between all the gateways in the terrestrial segment of the network, in order to handle e.g. gateway failures or link outages.

4.2.2 Types of NAT

There are a number of established varieties of NAT.

Traditional or Outbound NAT hides the private address space from global visibility by translating private ports and addresses seen in the headers of outbound packets. It removes the need for internal routing updates to propagate outside the private network to the global Internet. However, routing updates from the global Internet are still propagated within the private network, which is undesirable for our scenario. Since translation is only carried out on the way out, only hosts within the private network can initiate sessions.

Bidirectional or two-way NAT adds DNS support via an Application-Level Gateway (ALG) and address binding to allow sessions to be initiated externally as well; routing updates from outside are still propagated internally.

Twice NAT translates addresses in both directions, rewriting internal addresses in packet headers to addresses associated with the gateway externally, and rewriting external addresses to addresses associated with the gateway internally. This means that the external Internet and internal network only need to know how to route to the appropriate gateway at the edges between the networks; it is no longer necessary to propagate routing updates into what are now separate routing realms. This is useful for the satellite network to decrease the routing state held onboard satellite to that of only the private satellite network. A DNS ALG, where names bind to different addresses depending on whether you're inside or outside the private network, is also necessary.

Twice NAT offers the ability to abstract from a global *physical network address* to a *logical network address* that is used only within the constellation to identify and route to the translating earth station gateway. As IP multicast communication is already abstracted to a logical group address, it is not necessary to translate multicast packet headers, and multicast can be handled as it is for terrestrial networks.
Routing table lookup within the satellite network can then become as simple as masking the destination address to determine in which block of addresses it lies, and with which destination gateway or constellation user network that block of addresses is associated.

The translation of packet headers takes place in the gateway, where global terrestrial routing tables are held. This moves complexity from the space segment into the ground segment, where more processing power is available.

4.2.3 Implementation problems with NAT

NAT is often regarded as undesirable as it affects the implementation of existing applications and security services. As well as rewriting IP packet headers, it becomes necessary to rewrite in-band information in packet payloads that duplicates or relies on the header address or port information, using specific ALGs for each protocol to do so. This can add considerable implementation difficulty and processing overhead. NAT affects security using IPSec[37] as a result of this.

NAT also breaks explicit IP fragmentation, since only the first fragment of a packet possesses information identifying the protocol and the source and destination port used by the applications, while the remaining fragments are assigned fragmentation IDs that are not unique. This makes tracking of multiple simultaneous connections from the same end host difficult. Use of Path MTU discovery can discourage fragmentation.

4.2.4 NAT with QoS

From a QoS point of view, NAT interacts badly with RSVP, the resource reservation signalling protocol associated with IP Integrated Services (int-serv), because the use of NAT invalidates RSVP Integrity Objects among other issues, due to the use of in-band information in those objects.[38]

This does not entirely prevent the use of RSVP with NAT. Since a broadband constellation network acts as a high-capacity backbone or transit network for its customers, RSVP integrated services would in any case experience

the scalability problems described earlier; RSVP-capable NAT gateways would need to map to a roughly-equivalent differentiated-services PHB within the constellation network.

4.2.5 Other NAT problems

NAT for multihomed constellation network users – where a user's small terrestrial network has a satellite link for redundancy in case their terrestrial connection to the Internet fails – also poses implementation problems.

As the user's terrestrial network should not have to dynamically renumber itself into the constellation's realm if the terrestrial connection fails – avoiding renumbering is a motivation for implementing NAT in subnetworks on the Internet – one might assign the user's network addresses within the constellation's realm and require NAT at the router on the user's outbound terrestrial link. This NAT gateway could peer with and exchange bindings with constellation NAT gateways across the constellation network. However, this means that renumbering into the constellation's realm is necessary upon addition of the satellite link, making adding satellite redundancy to an existing network difficult.

4.2.6 NAT in summary

NAT is an attractive way of decreasing routing table overhead by gaining address space separation into separate realms at the IP level without the need for tunnelling. However, NAT's considerable technical implementation problems mean that its use must be carefully evaluated.

From a non-technical viewpoint, these problems can arguably be considered a feature as far as the constellation network operator is concerned: IP-level support for services can be enabled and disabled on a per-ALG basis, and being seen to discourage implementation of IPSec can be viewed as helpful in meeting the security concerns of international governments.

Implementing protocol support at the edges of the constellation network in the ground segment, which is easier to manage and change, is an advantage of NAT – but this is an advantage shared by tunnelling and by exterior routing, which do not have NAT's protocol-specific implementation problems.

4.3 External routing for constellation networks

4.3.1 Exterior routing protocols

4.3.1.1 Autonomous systems

For administrative purposes, today's Internet is divided into many different autonomous systems (AS). An AS is a collection of networks under a common administration using a consistent routing protocol. It is identified by a unique 16-bit number assigned by the Network Information Center (NIC).

Splitting the Internet into ASs makes it possible for groups of networks using different routing strategies to cohabit. When packets travel between ASs, they must cross a pair of connected border gateways. The constellation network can be seen as an AS.

4.3.1.2 The BGP protocol

Autonomous systems (AS) must communicate and exchange routing information to make global routing possible. Border gateways run an exterior routing protocol that enables them to determine routes to other AS. These routes are then propagated in the autonomous system through the internal routing protocol.

The Border Gateway Protocol (BGP)[39,40] is an example of an exterior routing protocol widely deployed in the Internet, having mostly supplanted the older Exterior Gateway Protocol (EGP). BGP connections (established over TCP) are established between the border gateways:

- *External BGP connections*, between the border gateways of neighbor autonomous systems, are used to advertise routes to networks of the autonomous system, and routes to other autonomous systems' networks. The border gateway should only advertise routes that it itself uses, but it is possible to restrict these exported routes for political reasons.
- *Internal BGP connections*, between all the border gateways of the same autonomous system, are used to exchange routes learned from external connections. They then decide on an egress point for networks outside the AS by minimizing the *external metric*, evaluated locally by the border routers using criteria such as the length of the AS path. Multiple gateways connected to a neighbouring AS may be chosen between for communication using the *border gateway*

preference value (awkwardly called the Inter-AS metric in BGP) that is advertised by this neighbouring AS.

Routes to external networks via border gateways are then imported to all the routers of the autonomous system, using the internal routing protocol to indicate the choice of border gateway, whatever that routing protocol may be.[41]

4.3.2 BGP in the constellation network

4.3.2.1 BGP traffic

The border gateways of an autonomous system must be connected with BGP connections internal to the AS to form a completely connected graph. Routing updates received at one of the border gateways must be propagated to the rest of the border gateways, in a similar fashion as discussed for NAT. This generates a lot of traffic, as the updates concern information on routes to all networks in the Internet external to the AS. This traffic will best be handled in the constellation network by dedicated terrestrial links, to avoid large amounts of routing table state update traffic passing through the satellite constellation.

Networks that use the same AS route can be aggregated so that fewer updates are sent to internal peers.

4.3.2.2 Choosing ingress and egress points

Routing from or to an external network is done hierarchically: the ingress or egress point is chosen by the border gateways, and the shortest route or most appropriate route to that point is determined by the internal routing protocol. External metric is therefore given more importance than internal metric.

In terrestrial networks, this sounds reasonable, since the external metric represents the cost of sending a packet through other ASs over a large geographic distance, whereas the internal metric measures the lesser cost of the transmission inside a single AS.

However, in the constellation network, the length of the internal path may easily be as large as or larger than that of the external one, since this autonomous system itself covers the entire surface of the Earth. This discrimination between external and internal metric may therefore be inverted.

It is possible to force incoming traffic for a given constellation user network to come through one of the nearest border gateways by advertising a route to this network only from these gateways. For outgoing traffic, having all the border gateways import their best route inside the AS mapping the external metric on the internal one enables us to optimize the path based on a combination of external and internal metrics.

Note that this may lead to asymmetrical routing, and that this asymmetrical routing may not work with NAT, where incoming and outgoing traffic must cross the same NAT gateways to avoid propagation of per-flow state between peered NAT gateways.

5 CO-EXISTING WITH IP ROUTING

5.1 IP routing on ATM with MPLS

Support for IP routing is extremely useful for handling IP traffic well, but is not useful for routing non-IP traffic, such as ATM or frame relay.

Given the large amount of work on wireless and satellite ATM links, it is likely that ATM will provide an underlying link-layer protocol over which IP traffic will be carried within the satellite constellation. A satellite-specific MAC layer must be defined for the ISLs and for the earth-space air interface, much as a MAC layer would be needed for IP.

With the use of ATM, the interworking of IP and ATM poses a number of interesting problems, particularly with respect to routing and support for IP multicast and QoS.

One solution for IP multicast over ATM is the Multicast Address Resolution Server (MARS), which maps IP multicast addresses into ATM server addresses.[42] However, the MARS family does not cope with mobility, does not scale easily to multiple servers that must share state, and is difficult to implement because ATM's routing paradigm and resulting multicast model differ considerably from those of IP.

Support for IP QoS over ATM is a non-trivial problem due to difficulty in mapping the IP QoS models accurately to available ATM service classes.

Multi-Protocol Label Switching (MPLS) is a developing technology which warrants serious consideration for the IP-over-ATM scenario.[43] It appears likely to be standardised as the IP-over-ATM transport method of choice by the ITU.

MPLS uses a label-swapping paradigm to integrate the flexibility and efficiency of Layer-3 IP routing with the high-speed Layer-2 switching of ATM. Fixed-size labels are assigned to the IP packets according to their respective egress destination nodes in the satellite network. Labelled packets that are destined for the same egress traverse a label-switched path (LSP) that is bound to an equivalent Layer-3 route. Every MPLS-enabled ATM switch on the LSP checks the label and rapidly forwards the packet to the appropriate output interface with little lookup overhead. Labels are exchanged by means of a label distribution protocol (LDP).[44] LSPs can be pre-established for reserved traffic, or created when required after initial layer-3 IP routing of a flow, and may be updated according to the IP routing tables.

Relative to the interconnection of IP edge routers tunnelled over an ATM core, MPLS improves the scalability of routing, due to the reduced number of immediate peers and elimination of the '*n*-squared' logical links between the *n* IP routers at the edge of the ATM core that are operating IP routing protocols.

A major feature of MPLS is that all ATM software above the ATM adaptation layer (AAL), including signalling, does not need to be involved in the routing of IP traffic, and does not have to be present or even defined. As the MPLS flows depend upon the IP routing tables, IP routing has full control of IP traffic. It is possible to use MPLS control plane to provide IP routing of IP traffic in parallel with native ATM-Forum control plane for ATM traffic without interference (so-called 'ships in the night' operation).

MPLS provides the following major advantages for IP over ATM, which are of benefit to IP traffic in constellation networks:

- Layer-3 IP routing can be used as is, without any need for interoperability with ATM switching or for explicit circuit setup;
- Hierarchical MPLS can be used to hide external routing information from internal nodes. Multiple layers of tunnelling are possible via label stacks, allowing e.g. BGP information to be distributed through nodes that remain unaware of routing outside the constellation AS;
- MPLS forwarding can be accomplished with little computational overhead – of benefit to on-board processing;
- IP multicast spanning trees can be mapped in a straightforward manner onto the ATM network by mapping branches of the multicast trees directly to the relevant LSPs;
- IP QoS, whether it be integrated or differentiated services,[45] can be supported using the available IP routing and traffic engineering functionality. Mapping differentiated services behaviour aggregates to available label values is possible;[46]
- downstream merge to a node supports multipath communications across a mesh network well. Layer 3 forwarding and hashing functions on source and destination addresses can be used to prevent out-of-order packets in individual source/destination flows, or individual label stacks can be used;
- MPLS can be configured to use explicit routing controlled by egress switches, in order to divert traffic from a congested part of the network, for example, or to allow the egress ground-based gateways and user terminals to explicitly control routing across the constellation network based on the visible satellites and distributed topology information.

5.2 Constraint-based routing

QoS routing can be important for LEO constellations, due to their redundant mesh topologies and choice of available paths to a destination. Given a QoS request for a flow, QoS routing could return the route that is most appropriate to the QoS requirements.

Constraint-based routing considers not only the topology of the network and the QoS requirements of the flow, but also resource availability of the links, and possibly other information specified by the network administrators – such as assigned link costs from Virtual Topology Routing. By taking all these factors into consideration, constraint-based routing may find a longer and lightly-loaded path better than the heavily-loaded minimum-delay or -hop path, distributing network traffic more evenly, avoiding congestion and improving network utilization.

The primary components of a constraint-based routing scheme are the advertisement of link state information and the selection of metrics and route computation algorithms. This is applicable to MPLS-based architectures. MPLS LSPs allow constraint-based routing with per-LSP statistics, and LDP can be extended to provide constraint-based routed label switched paths (CRLSPs).[47] Given precise information on how traffic flows through the network, constraint-based-routing can determine how to dynamically configure LSPs for explicit routing to carry the traffic through the network more efficiently and provide effective QoS.

6 CONCLUSIONS

Despite common objections, IP routing is not impossible to use in satellite constellation networks. It can be implemented in a manageable fashion by a combination of border routing protocols, tunnelling, network address translation and MPLS.

Use of IP routing offers benefits to the IP traffic that is routed, allowing straightforward implementation of support for IP multicast and for IP quality of service.

MPLS appears to be a realistic method for implementing support of advanced IP routing functionality on an ATM-based backbone likely to be found in the context of a satellite constellation network, while allowing good support for non-IP protocols. Use of MPLS in this scenario is worth investigating in further research.

REFERENCES

1. ISO (1984). 'Basic Reference Model for Open Systems Interconnection', **ISO7498**, 1984. Recommendation **X.200** of the International Telecommunications Union.
2. G. Maral, *VSAT Networks*, J. Wiley & Sons, 1995.
3. V. G. Cerf, 'Packet satellite technology reference sources,' **RFC829**, DARPA, November 1982.
4. V. Jacobson, R. Braden and D. Borman, 'TCP Extensions for High Performance', **RFC1323**, proposed standard, May 1992.
5. M. Mathis, J. Mahdavi, S. Floyd and A. Romanow, TCP Selective Acknowledgement options', RFC2018, proposed standard, October 1996.
6. M. Allman, V. Paxson and W. Stevens, 'TCP Congestion Control', **RFC2581**, proposed standard, April 1999.
7. J. Walker, 'Some circular orbit patterns providing continuous whole earth coverage', *Journal of the British Interplanetary Society*, vol. 24, pp. 369-384, 1971.
8. A. Ballard, 'Rosette constellations of earth satellites', *IEEE Transactions on Aerospace and Electronic Systems*, vol. 16, no. 5, September 1980, pp. 656-672.
9. L. Wood, 'Network performance of non-geostationary constellations equipped with intersatellite links', Masters thesis for University of Surrey, Rapport 95-9, ENST Toulouse, November 1995.
10. N. Maxemchuk, 'Routing in the Manhattan street network', pp. 503-512, *IEEE Transactions on Communications*, vol. 35 no. 5, May 1987.
11. J. Moy, 'OSPF version 2', **RFC2328**, Internet standard **STD0054**, April 1998.
12. G. Malkin, 'RIP version 2', **RFC2453**, Internet standard **STD0056**, November 1998.
13. M. Werner, 'A dynamic routing concept for ATM-based satellite personal communication Networks', *IEEE Journal on Selected Areas in Communications*, vol. 15, no. 8, pp. 1636-1648, October 1997.
14. M. Werner, C. Delucchi, H-J. Vogel, G. Maral and J-J.de Ridder, 'ATM-based routing in LEO/MEO satellite networks with intersatellite links', *IEEE Journal on Selected Areas in Communications*, vol. 15, no. 1, pp. 69-82, January 1997.
15. R. Mauger and C. Rosenberg, 'QoS Guarantees for Multimedia Services on a TDMA-Based Satellite Network', *IEEE Communications Magazine*, July 1997.
16. H. Uzunaliogly, W. Yen, I. Akyildiz, 'A Connection Handover Protocol for LEO Satellite ATM Networks', *Proceedings of ACM Mobicom '97*, pp. 204-214, October 1997.
17. W. Fenner, 'Internet Group Management Protocol, Version 2', **RFC2236**, proposed standard, November 1997.
18. D. Waitzman, C. Partridge and S. Deering, 'Distance Vector Multicast Routing Protocol', **RFC1075**, experimental, November 1988.
19. C. Shields, J. J. Garcia-Luna-Aceves, 'The Ordered Core-based Tree Protocol', *INFOCOM '97*, 1997.
20. J. Saltzer, D. Reed and D. Clark, 'End-to-End Arguments in System Design', pp277-288, *ACM Transactions in Computer Systems*, November 1984.
21. R. Braden, D. Clark and S. Shenker, 'Integrated Services in the Internet Architecture: an Overview', **RFC1633**, informational, June 1994.
22. R. Braden, L. Zhang, S. Berson, S. Herzog and S. Jamin, 'Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification', **RFC2205**, proposed standard, September 1997.
23. S. Blake, D. Black et al., 'An Architecture for Differentiated Services', **RFC2475**, informational, December 1998.
24. K. Nichols, S. Blake, F. Baker and D. Black 'Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers', **RFC2474**, informational, December 1998.
25. J. Mogul and S. Deering, 'Path MTU Discovery', **RFC1191**, draft standard, November 1990.
26. J. McCann, S. Deering and J. Mogul, 'Path MTU Discovery for IP version 6' **RFC1981**, proposed standard, August 1996.
27. S. Keshav and R. Sharma, 'Issues and Trends in Router Design', *IEEE Communications Magazine*, vol. 36 no. 5, pp. 144-151, May 1998.

28. E. C. Rosen, A. Viswanathan and R. Callon, 'Multiprotocol Label Switching Architecture', work in progress as internet-draft, IETF MPLS working group, August 1999, approved as proposed standard by the IESG, 16 September 1999.
29. L. Wood, H. Cruickshank, Z. Sun, 'Supporting group applications via satellite constellations with multicast' pp. 190-194, *Proceedings of the Sixth IEE Conference on Telecommunications (ICT '98)*, March 1998.
30. H. Eriksson, 'MBone: The Multicast Backbone', *Communications of the ACM*, vol. 37, pp. 54-60, August 1994.
31. M. Laubach and J. Halpern, 'Classical IP and ARP over ATM', **RFC2225**, proposed standard, April 1998.
32. W. Simpson, 'IP in IP Tunnelling', **RFC1853**, informational, October 1995.
33. P. Narvaez, A. Clerget, and W. Dabbous, 'Internet Routing over LEO Satellite Constellations', *Third ACM/IEEE International Workshop on Satellite-Based Information Services (WOSBIS '98)*, October 1998.
34. D. Plummer, 'An Ethernet address resolution protocol, or converting network protocol addresses to 48-bit Ethernet address for transmission on Ethernet hardware', **RFC826**, standard, November 1982.
35. K. Ramakrishnan and S. Floyd, 'A Proposal to add Explicit Congestion Notification (ECN) to IP', **RFC2481**, experimental, January 1999.
36. P. Srisuresh and M. Holdrege, 'IP Network Address Translator (NAT) Terminology and Considerations', **RFC2663**, informational, August 1999.
37. S. Kent and R. Atkinson, 'Security Architecture for the Internet Protocol', **RFC2401**, proposed standard, November 1998.
38. P. Srisuresh and M. Holdrege, 'Protocol complications with the IP Network Address Translator (NAT)', work in progress as internet-draft, IETF NAT working group, June 1999.
39. Y. Rekhter and T. Li, 'A Border Gateway Protocol 4 (BGP-4)', **RFC1771**, draft standard, March 1995.
40. Y. Rekhter, P. Gross, 'Application of the Border Gateway Protocol in the Internet', **RFC1772**, draft standard, March 1995.
41. K. Varadhan, 'BGP OSPF Interaction', **RFC1403**, proposed standard, January 1993.
42. R. Talpade, M. Ammar, 'Multicast Server Architectures for MARS-based ATM multicasting' **RFC2149**, informational, May 1997.
43. R. Callon et al, 'A Framework for Multiprotocol Label Switching', work in progress as internet-draft, IETF MPLS working group, July 1999.
44. L. Andersson et al, 'LDP Specification', work in progress as internet-draft, IETF MPLS working group, June 1999.
45. I. Andrikopoulos and G. Pavlou, 'Providing Differentiated Services to MPLS Networks', *Proceedings of the 7th IEEE/IFIP International Workshop on Quality of Service (IWQoS '99)*, pp. 207-215, June 1999.
46. L. Wu et al, 'MPLS support of differentiated services by ATM LSRs and frame relay LSRs', work in progress as internet-draft, IETF MPLS working group, June 1999.
47. B. Jamoussi et al, 'Constraint-based LSP setup using LDP', work in progress as internet-draft, IETF MPLS working group, August 1999.

AUTHOR BIOGRAPHIES

Lloyd Wood received his masters in electronic engineering from Loughborough University in 1994. He gained an MSc in satellite communication engineering from the University of Surrey in 1995, after studying constellation topology at ENST Toulouse. Lloyd then joined what is now the Centre for Communication Systems Research (CCSR), where he has participated in a number of ACTS and Esprit projects. He maintains the *Lloyd's satellite constellations* web resource while pursuing his PhD in the area of Internet protocols and delay in satellite constellation networks. His interests include multicast, routing, and orbital geometry.

Antoine Clerget graduated from the Ecole Polytechnique in 1995, where he followed graduate studies in a number of subjects, including mathematics and computer science. He then undertook further studies in telecommunications and computer science at ENST, before joining the RODEO team at INRIA Sophia-Antipolis in September 1997 to begin his doctoral work. Antoine's main areas of interest are in unicast and multicast routing, congestion control, asymmetrical and unidirectional network topologies and transport in mobile networks, such as low-earth-orbiting satellite constellations.

Ilias Andrikopoulos gained a Diploma in Physics from the University of Athens in 1993, and an MSc in Information Technology from University College London in 1994. He joined the University of Surrey in 1995, where he is a research fellow within the CCSR. He has been involved in a number of European ACTS and UK projects related to broadband and Internet networking while pursuing his PhD in the area of QoS provisioning in IP networks. His interests also include TCP/IP over ATM, congestion control, network and traffic management, scheduling and buffer management for high-speed networks. He is a student member of the IEEE and an associate member of the IEE.

George Pavlou gained a Diploma in Electrical Engineering from the National Technical University of Athens. He obtained MSc and PhD degrees in Computer Science from University College London, where he acted as a senior research fellow and lecturer. He is currently professor of information networking at the Centre for Communication Systems Research in the University of Surrey, where he leads the activities of the networks research group. Over the last ten years he has led a number of collaborative research projects in networking, with emphasis on management and control. Routing and quality of service are among his key research areas of interest. George has contributed to ISO and ITU-T standardisation work.

Walid Dabbous graduated from the Department of Electrical Engineering at the Lebanese University in Beirut in 1986. He obtained his DEA and Doctorat d'Université from the University of Paris XI in 1987 and 1991 respectively. Walid joined the RODEO Team at INRIA in 1987. He has been a staff researcher at INRIA Sophia-Antipolis since 1991, and has led the RODEO team since 1996. Walid's main research interests span the areas of high-performance communication protocols, congestion control, reliable multicast protocols, audio and video conferencing over the Internet, efficient and flexible protocol architecture design and the integration of new transmission media, such as satellite links, into the Internet.