

SUPPORTING GROUP APPLICATIONS VIA SATELLITE CONSTELLATIONS WITH MULTICAST

L. Wood, H. Cruickshank, Z. Sun

Centre for Communication Systems Research, University of Surrey, United Kingdom

Construction of low-earth-orbiting satellite constellations for global wireless communication to small user terminals has begun with the launch of a number of Motorola's *Iridium* satellites in 1997 and the building of a number of Qualcomm's *Globalstar* satellites for launch in 1998. These early constellations, and their competitors, aim to provide low-bandwidth satellite telephony world-wide, including paging, faxing and 2400bps modem services.

Future satellite constellations have been proposed for high-bandwidth wireless data transmission. These include *Teledesic* (1), Alcatel's *SkyBridge*, formerly known as *Sativod* (2), and Motorola's *Celestri*, whose components were formerly known as *M-Star* and *Millennium* (3).

In providing data networking services, the ability to internetwork with pre-existing ground networks will be an important requirement for these constellations.

Proponents of these schemes have stressed their advantages over existing geostationary satellite networking solutions of higher overall capacity, due to frequency reuse, and of decreased latency in communication due to the decrease in propagation delay in going from high-altitude geostationary orbit (GEO) to low-altitude low earth orbit (LEO).

However, support for low-latency group applications, such as multiparty videoconferencing, is problematical. Here, the networking aspects of the broadband satellite constellations are discussed, and the suitability of the constellations for multicast is assessed.

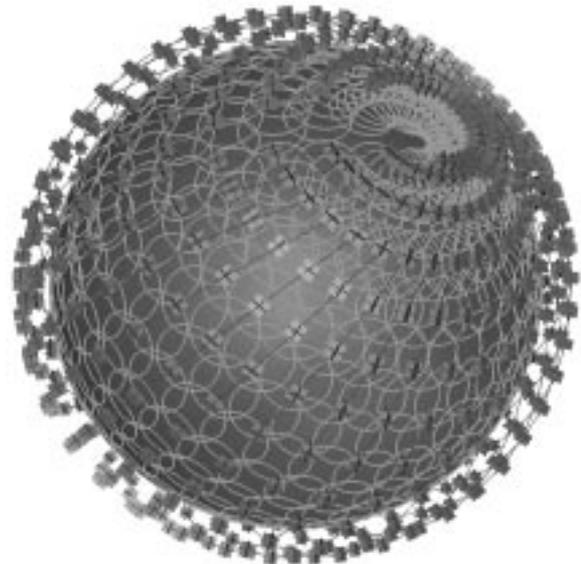
1.0 DATA CONSTELLATION OVERVIEW

1.1 Teledesic

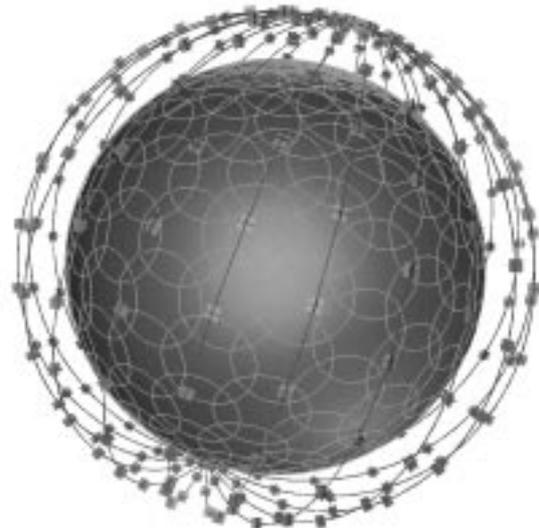
Announced in 1994, *Teledesic* is the first and most ambitious of the LEO broadband constellations. Although originally planned as 21 near-polar orbital planes of 40 active satellites and 4 in-orbit spares per plane at an altitude of 700km, in early 1997 it was scaled back to 12 near-polar planes of 24 active satellites and 3 in-orbit spares at an altitude of 1400km.

Each *Teledesic* satellite has eight intersatellite links, in the 60GHz band, to its two nearest neighbours in each of the four cardinal directions, forming a complex

cylindrical mesh network, a topology described as 'geodesic'. Constellation topology is discussed in more detail in (4).



a. Original design



b. 288-satellite redesign

Figure 1 - *Teledesic* satellites and ground coverage rendered by SaVi (5)

Teledesic ground terminals will provide data rates from 16kbps to 2Mbps in 16kbps increments, while gateway 'GigaLink' terminals between *Teledesic* and ground networks have data rates from 155Mbps (OC-3)

to 1.2Gbps (OC-24). Each satellite acts as a fast packet switch, using a proprietary fixed-length 512-bit packet format with a proprietary connectionless adaptive-routing protocol, rather than the increasingly popular Asynchronous Transfer Mode (ATM).

Ka-band frequencies were allocated to *Teledesic* at the 1995 World Radio Conference (WRC), and the US Federal Communications Commission (FCC) granted a licence in March 1997. Construction has not yet begun, but service is scheduled to begin in 2002.

1.2 SkyBridge

In February 1997 Alcatel announced its *SkyBridge* constellation, a redesign of its earlier *Sativod* proposal, and filed an application for frequency with the FCC.

SkyBridge is intended to share Ku-band with existing GEO satellites. It uses two overlapping offset constellations of 32 satellites each (each being four planes of eight active satellites per plane, all at an altitude of 1457km) to give a choice of satellites to ground stations. For each satellite pair, the satellite furthest from the centre of a GEO footprint that satellite footprints overlap with is used for minimum interference with the GEO satellite.

SkyBridge does not use intersatellite links. Its satellites are planned to act as in-orbit 'bent-pipe' transponders, or amplifying repeaters, forwarding traffic to local gateways. This can therefore be thought of as a ground network with a space-based component, rather than as a space-based network, as all the switching fabric is ground-based.

User terminals offer data rates from 16kbps to over 60Mbps, in increments of 16kbps, using an ATM cell-switching interface.

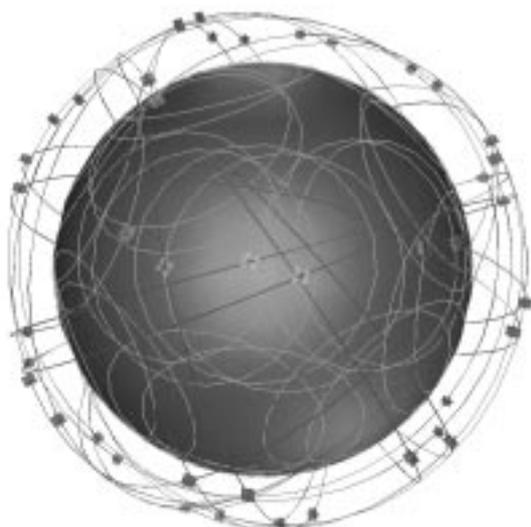


Figure 2 - *SkyBridge* satellites and ground coverage rendered by SaVi (5)

Alcatel plans to deploy one 32-satellite constellation by 2001, and the second by 2002.

1.3 Celestri

In July 1997 Motorola announced *Celestri* and filed an application for a Ka-band licence with the FCC.

Celestri combines a LEO constellation (formerly developed at Motorola under the name *M-Star*) with a GEO constellation (developed under the name *Millennium*). The LEO constellation has nine planes, each containing seven active satellites and one spare, at an altitude of 1400km.

Each satellite has six optical intersatellite links, forming a toroidal mesh network around the Earth, and an onboard ATM switch with throughput of 17.5Gbps.

Communication rates to terminals are either from 64kbps to 2Mbps in increments of 64kbps, or at the fixed rates of 10Mbps, 16Mbps, 51Mbps (OC-1) or 155Mbps (OC-3). Protocol adapters for transparently carrying various network protocols over the *Celestri* network are to be included in terminals, presumably to tunnel these protocols over the underlying ATM infrastructure.

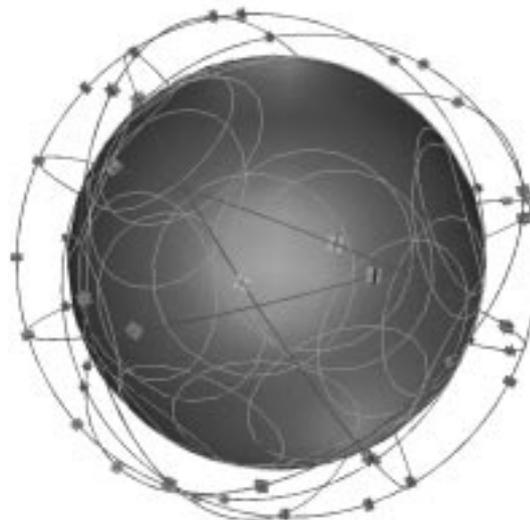


Figure 3 - *Celestri* LEO satellites and ground coverage rendered by SaVi (5)

Motorola plans to launch *Celestri* from 2001, and to have it fully operational by the end of 2002.

2.0 MULTICAST

2.1 What is multicast?

One-to-one end-to-end connectivity across an internetwork can be accomplished by circuit-switching

or by a virtual circuit. For group applications, where more than two users simultaneously exchange messages and maintain state, the number of circuits required would increase rapidly as the number of users increases. To prevent applications from needing to know about all users in the group or needing to be responsible for maintaining all these circuits, and to decrease network load, we require multicast, which is the efficient emulation of broadcast within the constraints of a network environment.

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 simply flooding the entire internetwork with unnecessary broadcast packets.

The set of all interested hosts forms a multicast group, and group management is an internetwork function. Users not in the group do not see the data (perhaps because there is no need for the data to be sent across their network) or, if they do see it, discard it.

To communicate the data efficiently to all users in the group, the internetwork must set up a spanning tree, along which the messages can be sent and replicated at tree branches.

2.2 Multicast in existing network stacks

The Internet Protocol (IP) implements multicast groups as part of the network address space, using reserved Class D addresses, where there is a one-to-one relation between a multicast group and a D address at any given time. Considerable work has been carried out on the implementation of a variety of multicast protocols across the Internet. IP multicasting protocols rely on the Internet Group Management Protocol (IGMP) to manage multicast groups (6). There is a wide choice of available multicast routing protocols, such as the Distance Vector Multicast Routing Protocol, DVMRP (7) and Multicast Open Shortest Path First, MOSPF (8).

These protocols can be split into two main groups, according to the assumptions that they make about group distribution and network traffic.

Dense-mode multicast protocols assume that multicast group members are densely distributed throughout the internetwork, i.e. that many subnetworks contain at least one group member, and that internetwork capacity is plentiful. These are data-driven, in that construction of the multicast tree begins top-down from the source outward, and data on the state of the tree is propagated to all routers. The multicast trees constructed allow data to travel in one direction, from source to group, emulating broadcast.

For full group-to-group communication, a multicast tree must be set up from every user; this scales badly for large groups and imposes joining and leaving

overhead and delays. DVMRP and MOSPF are dense-mode.

Sparse-mode multicast protocols assume that group members are sparsely distributed and that capacity is constrained. In sparse mode, the multicast tree is receiver-initiated, i.e. a router only becomes involved in the construction of a multicast distribution tree only when one of the hosts on its subnet requests membership. Core-Based Trees (CBT), with a single central router from which the tree branches out in all directions, have been suggested for groups where there are many active senders within the group, allowing multi-way communication over a single tree (9). This effectively minimises the amount of multicast state routing information that needs to be stored and the number of routers involved. This single-tree approach differs from the (source, group) pairings of DVMRP and MOSPF, and scales better for large networks, with minimal joining and leaving overheads. Protocol-Independent Multicast - Sparse Mode (PIM-SM) constructs a multicast tree around a chosen router, called a rendezvous point, similar to the core in CBT (10).

Given the constraints of a satellite constellation network with intersatellite links, a sparse-mode multicast protocol would appear to make best use of limited, expensive, capacity. However, as the satellites move in their orbits it would become necessary to periodically move the core of the tree, and this would also require support for native multicast routing in the satellite on-board switches.

As we saw in section 1.0, proposed commercial satellite constellations intend to implement ATM-based networks. With its circuit-switching focus, ATM did not originally have the concept of multicast, and the closest it came to multicast was the concept of a multicast server, where virtual circuits from all receivers would connect to a single source. This scales extremely badly in large networks with many users, as the links adjacent to the server soon run out of capacity and available Virtual Path Identifiers (VPIs) to allocate to connections.

Work is now taking place to add multicast support to ATM, e.g. SEAM, which uses a single core-based tree and combines ATM Virtual Circuits to minimise VC switching overhead (11). However, IP multicast is a considerably more mature area. It is likely that much of the traffic over ATM networks, including satellite ATM networks, will be generated by applications using IP, and tunnelled through ATM.

2.3 Internetworking with multicast

When networks on an internetwork are not multicast-aware, multicast messages can be *tunnelled* through them by *encapsulating* (and, if necessary, fragmenting) the multicast messages within network messages native to that network.

This is the strategy adopted by the MBone, or Multicast Backbone, where isolated IP-multicast-aware networks are internetworked together across multicast-unaware networks by tunnelling IP multicast messages through ATM circuits between the multicast islands. This forms a virtual network layered on top of the physical network.

However, tunnelling gives a false picture of the latency of a connection as measured by a time-to-live (TTL) or hop counter, since the header containing the count value is encapsulated and is not decremented in the tunnel, no matter how long the tunnel is. All tunnels appear the same length to the packets being tunnelled through them.

Tunnelling does not make optimal use of the intervening network. The tunnel must often be set up manually, and the tunnel endpoints cannot be moved easily.

Adapting the TCP/IP and ATM network stacks to each other to increase throughput when IP is layered over ATM has been an area of considerable interest for some time, and adaptations such as TCP Boston, where the implementations of both protocols are modified slightly in order to increase tunnelling performance, are starting to appear. (12)

As an extension of this, protocol federation of ATM multicast protocols with IP multicast protocols is one possible, albeit highly unlikely, solution to allow true IP/ATM internetwork multicasting.

2.4 Multicast over satellite

Given the maturity and widespread adoption of IP multicast protocols, we can expect a demand for IP-multicast-aware group applications to be used in wireless environments, including satellites.

Satellite-based videoconferencing could be accomplished by tunnelling the IP-multicast messages through the satellite gateways, but this would require the setting up of multiple tunnelled virtual circuits between geographically-separate users, making group management difficult and using more satellite capacity than would be necessary if the satellites' on-board switches were to support IP multicast directly. None of the proposed commercial schemes do this.

SkyBridge treats each of its satellites as a simple transparent repeater, resulting in each satellite connection acting as a simple short one-hop ATM tunnel that is transparent to IP multicast, and the problem of internetwork multicast is moved away from the satellites into the ground networks that are utilising *SkyBridge* satellites for connectivity.

Celestri is effectively a standalone ATM network utilising ATM switching in the ground/air interface

and between satellites. Tunnelling of IP multicast packets over a changing-geometry ATM switching fabric will not permit multicasting within the constellation. The inclusion of a complementary GEO component in *Celestri* for broadcast applications appears to tacitly acknowledge and attempt to side-step the networking limitations of the LEO component. However, the increase in propagation delay for the GEO component is offset by the decrease in the amount of switching required.

Teledesic utilises a purpose-designed 512-bit fixed-length packet, and is likely to encapsulate or tunnel even ATM across its network. This protocol tunnelling across the *Teledesic* geodesic mesh is likely to make implementation of IP multicast within the constellation network impossible.

3.0 CONCLUSIONS

Efficient use of satellite constellations for group applications requires that the satellites' on-board switches include native support for multicast. This appears unlikely in the commercially-proposed schemes.

Leaving implementation of multicast solely to the IP-routing ground networks, rather than making it a problem for both ground and mesh satellite networks, would appear to make the problem of implementing efficient internetwork multicast with a satellite component more tractable.

SkyBridge can be considered the least ambitious of the proposed schemes when viewed from a satellite networking viewpoint, as the satellites are simple transponders rather than complex routing switches with intersatellite links. However, by being the most transparent from a tunnelling viewpoint *SkyBridge* paradoxically appears to offer the most at present for the implementation of IP multicast applications – at least until the appearance of native IP multicast routing onboard satellites.

4.0 REFERENCES

1. Sturza, M. A., "Architecture of the *Teledesic* satellite system", Proceedings of the International Mobile Satellite Conference 1995, 212-218.
2. Spector, P. L., Olson, J. H. et al, FCC filing of February 28, 1997, "Application of SkyBridge LLC for authority to launch and operate the SkyBridge system", SkyBridge LLC.
3. Kennedy, M. D., Lambergman, B. et al, FCC filing of June 1997, "Application to construct, launch and operate the Celestri multimedia LEO system", Motorola Global Communications, Inc.

4. Wood, L., "Network performance of non-geostationary constellations equipped with intersatellite links", MSc thesis, University of Surrey, November 1995.
5. Worfolk, P. A. and Thurman, R. E., "SaVi - software for the visualization and analysis of satellite constellations", developed at The Geometry Center, University of Minnesota.
6. Deering, S., "Host extensions for IP multicasting", RFC1112, August 1989.
7. Partridge, C., Deering, S. et al, "Distance vector multicast routing protocol", RFC1075, November 1988.
8. Moy, J, "OSPF version 2", RFC2178, July 1997.
9. Ballardie, A. J., Francis, P. F. et al, "Core-based trees", Proceedings of ACM SIGCOMM, 1993.
10. Estrin, D., Farinacci, D. et al, 'Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification', RFC 2117, June 1997
11. Grossglauser, M. and Ramakrishnan, K. K., "SEAM: Scalable and Efficient ATM Multicast", Proceedings of Infocom '97.
12. Bestavros, A. and Kim, G., "TCP Boston: a fragmentation-tolerant TCP protocol for ATM networks", Proceedings of Infocom '97.

Supporting group applications via satellite constellations with multicast

pp. 190-194, *Proceedings of the Sixth IEE Conference on Telecommunications, Edinburgh, 29 March – 1 April 1998*, IEE Conference Publication no. 451, ISBN 0-85296-700, ISSN 0537-9989

Lloyd Wood, Haitham Cruickshank, Zhili Sun

Networks Research Group,
Centre for Communication Systems Research,
University of Surrey, Guildford GU2 5XH

L.Wood@surrey.ac.uk

<http://www.ee.surrey.ac.uk/CCSR/Networks/>

Permission for web reuse is granted by the
Institution of Electrical Engineers as
copyright holders.
