SLOT CLOUDS: GETTING MORE FROM ORBITAL SLOTS WITH NETWORKING

L. Wood*, A. da Silva Curiel+, J. Anzalchi+, D. Cooke+, C. Jackson+

* Global Defense and Space Group, Cisco Systems.

+ Surrey Satellite Technology Ltd, University of Surrey, Guildford.

IAC-03-U.4.07 Corresponding authors: **lwood@cisco.com a.da-silva-curiel@sstl.co.uk**

Abstract

Operators of traditional geostationary satellites are allocated well-defined orbital slots for their satellites. By breaking the uplink/downlink dependency, adapting onboard processing in the transponders to decode to baseband, co-locating multiple satellites supporting a variety of different uplink and downlink frequencies, and enabling interfaces with networked communication between the satellites using intersatellite links, the variety of services that the satellites can support can be increased. This permits more flexible use of all available satellite capacity. We call this concept the 'slot cloud'. The co-located satellites in the orbital slot together form a network and, particularly when using and communicating with the Internet Protocol, can be viewed as a network 'cloud' that provides functionality in a flexible manner.

Introduction

With a traditional 'bent-pipe' geostationary satellite, the satellite link is treated as just that: a single link in each direction between ground terminals. Although this link consists of an uplink followed by amplification, frequency downshifting and a downlink returning the signal content to the ground, the single satellite link budget includes all of these steps combined. There is a strong relationship – a codependency – between a signal's uplink and its downlink.

Often, even when demodulating or decoding a signal to baseband onboard the satellite, the relationship between the design of the uplink and the downlink remains very strong. This codependency can make for clarity of design and engineering optimization when the satellite is used for its intended purpose. This coupling between uplink and downlink can also permit flexibility in use of the single established channel through both the uplink and downlink that results, e.g. in allowing ground terminals to use turbo coding across links using satellites deployed before turbo coding had been developed, without requiring changes to the satellites.

However, this codependency can also limit the flexibility of link use, terminal design, and the range of networking services that can be offered by available satellite capacity as a whole.

Moving towards on-board processing (OBP) can decrease this uplink/downlink codependency. Increased onboard processing and switching capabilities on computationally 'smarter' satellites can introduce bridging and then networking functionality within and between satellites in an evolutionary fashion. Breaking the link dependency entirely can increase the flexibility of use of each satellite's uplink, downlink and payloads in various ways not envisaged by the original link designers. Breaking the link dependency allows us to evolve the middle of the link: the satellite.

Some trends in satellite payload design

An evolutionary approach to satellite networking can be articulated by first noting some existing trends in satellite payload design:

1. The tendency towards use of targeted shaped or multiple spotbeams to get as much useful coverage and capacity as possible from geostationary satellites. The number of separate payloads that can be carried and the number of separate transponders that can be carried to support this coverage increases as payload size shrinks and satellite size increases. (Though whether satellite size will continue to increase is itself an interesting question.[1])
2. The tendency towards increased on-board processing (OBP). Existing OBP is often

digital signal processing (DSP), used to improve signal gain by cleaning up the signal to varying degrees. Here the received signal can be demodulated to baseband and decoded to generate a 'clean' signal, before the carrier signal is regenerated and modulated at the required downlink frequency. OBP is also a natural complement to managing traffic across multiple similar spotbeams, where multiplexing and a greater or lesser degree of smart switching between feeder links and multiple spotbeams is needed onboard the satellite to utilize the spotbeams effectively if efficient direct communication between ground terminals in different spotbeams is desired. (This has traditionally not been the case; for satellite telephony, most calls are terminated in the terrestrial network, so efficient one-hop handling of satellite telephone to satellite telephone calls has not been a priority, and is often treated as a 'double hop' to and from Network Operation Centres on Earth.)

3. The tendency towards shared satellite buses, with payloads from different manufacturers and operators, with different owners and purposes, that share that bus, and that need to share a common bus standard for power and on-board housekeeping functions, at the very least. We will return to this point later.

An evolutionary approach to satellite networking

The trend towards use of OBP for DSP is of interest. When a baseband signal becomes available at some stage within the satellite payload due to signal processing, it becomes possible to consider accomplishing the minor incremental steps of:

a. copying or streaming that baseband signal elsewhere, or

b. fully decoding and then parsing the streamed contents of the baseband signal to see frames and their contents.

Replicating a decoded baseband transmission to a second transponder to have it regenerated and repeated at an entirely different downlink frequency can be done.

A further incremental step is to copy all (if switching out and streaming entire baseband streams) or parts (if switching on frames) of the baseband transmission out of the satellite through an intersatellite link (ISL).

Intersatellite links are known to work, and their feasibility has been repeatedly proven *en masse* by NASA's TDRSS network supporting the shuttle, International Space Station and other spacecraft,[2] by the commercially unsuccessful Iridium network,[3] and by the SILEX Artemis laser experiments.[4] In not having to pass through the atmosphere, intersatellite links are not plagued with the power limitations, useful frequency range limitations, or regulatory impediments that plague uplink and downlink allocations.

A single, seemingly-redundant, ISL repeater can be used to form a layer-2 'bridge' to another, similar payload with similar ISL terminal that is later launched on another satellite, allowing streams of data to be repeated or redirected between transponders at different frequencies on different satellites once both satellites are launched and the ISL bridge is established. This is shown in Figure 1.

Doing this permits some flexibility and reconfiguration as required in choice of uplinks and downlinks for data streams.

Adding switching on the parsed contents of that decoded baseband transmission means that bridging and networking capabilities can then selectively duplicate and copy frames or packets to different transponders, as desired. (This switching functionality need only be present at the one, later-launched, end to still be useful.) Decisions on where to switch each frame to can be made on each of those frames, leading to bridging of link-specific frames and, once frame contents are examined, even to networking functionality at the IP or other network layer.

Examining and selectively switching on frame contents permits multiplexing of baseband traffic from different uplinks and selection and prioritization of traffic, and can be used to increase the flexibility of transponder use.

It becomes useful to interconnect all similar payloads. A convention for streaming MPEG television transport streams across intersatellite links might emerge and become a standard, and this would encourage all payloads with television transponders to follow that design in order to interoperate with each other, so that one uplinked television signal could be relayed worldwide from multiple interconnected satellites.

Another step is to have the satellite payloads migrate from local bridging to considering network-layer traffic, and letting the switching decision for frame contents be guided by onboard routing functionality that is driven by:

- knowledge of available connections to transponders,
- topology of interconnected payloads and their transponders,
- and policy and administrative decisions concerning what traffic is best sent where.

It then becomes possible to connect together very dissimilar payloads at the network level and have them interoperate. As an example, an MPEG transport stream passing through a broadcast transponder bears little resemblance to an HDLC serial stream on a point-to-point link at first glance. The bridging done for these separate families of payload and designs is entirely different. Yet both HDLC and MPEG streams can contain IP packets, and once the IP packets are parsed from the decoded streams it becomes possible to switch IP packets into and out of each stream as required.

(In a world where television viewers are coming to expect on-demand interactive content and watch television using TiVo or other computing devices that are also fully Internet-enabled network devices each capable of requesting a unique mix of traffic, being able to multiplex IP streams within MPEG streams becomes more important. Being able to uplink a broadcast MPEG channel once and have it downlinked worldwide from interconnected networked satellites frees up uplink frequencies for yet more channels, for terrestrial reuse, or for other purposes such as IP-in-MPEG trunking.)

Once this network-level consensus point has been reached, interoperability between payloads and transponders of different families for different purposes becomes possible. The functionality and flexibility of all payloads and transponders becomes available across all interconnected payloads in the geostationary satellites holding position in their orbital slots, subject to interoperability and peering administrative agreements at the policy level.

At the network layer, these interconnected satellites across the same or multiple geostationary orbital slots can be perceived as a cloud of functionality that exists to allow network end devices using satellites in the cloud to communicate anywhere across that cloud. We call this concept the 'slot cloud'.

Benefits from the steps in the approach

There are a number of benefits that result from implementing the various steps that have been described above.

1. Flexibility of uplink and downlink choice. Satellite service operators are vulnerable to uplink frequency jamming; a broadcast uplink ground station can be a single point of failure. Deliberate jamming of signals for political reasons is not unknown; for example, all U.S. satellite feeds intended for Iran are believed to have been jammed by transmitters based in Cuba.[5] By being able to change uplink stations and uplink transponders, reconfigure downlink transponders to receive traffic from elsewhere, and redirect uplinked traffic to the desired downlink transponder, satellite operators become less vulnerable to physical denial-of-service attacks, because the uplink is no longer a single chokepoint. This does not require packet or frame-level processing; the ability to switch at the stream level, for e.g. broadcast television operators, will

suffice. Similarly, if a downlink transmitter fails, traffic streams can be easily rerouted from the existing uplink to another downlink. Breaking the uplink/downlink dependency means that an interfering signal does not get a 'free ride' to the downlink, and on-board processing and switching also pave the way to enabling authentication of ground transmitters and controlling access to downlink capacity.

2. Removing the 'double hop' problem between communicating ground terminals. Having to go from Earth to satellite and back, and then again, to deliver data end-to-end is undesirable. This can happen for ground terminals using different satellites, where the overall delay and number of hops would be decreased by a direct intersatellite link. This problem also occurs for ground terminals using the same satellite, where the satellite is not smart enough to be able to switch traffic directly between the terminals, and must relay the traffic to a ground network operations centre, which switches traffic from the link supporting one ground terminal to the link supporting the other ground terminal. This 'double hop' problem is a particular problem for active IP multicast group members, where anything sent by a group member behind a ground terminal can only be rebroadcast to other group members also using satellite ground terminals by the ground network operations centre which controls the broadcast facility. Onboard processing of IP packets can allow the multicast replication of packets and rebroadcasting to all interested group members behind ground terminals to be done onboard, decreasing both uplink/downlink capacity use and decreasing latency of communications between group members. This is particularly important for realtime group applications. Onboard switching and packet replication also allows a given feeder uplink to support a larger fan-out of spotbeams and ground terminals than it would otherwise.
3. The ability to add satellite functionality in an incremental fashion. By allowing interconnections directly between satellites, the available functionality in an existing satellite cloud is enhanced by adding a new satellite. This moves away from the existing all-or-nothing replace-asset-at-end-of-lifetime model. It may, over time, even subvert the need for extremely large do-everything satellite payloads, and permit the launch of smaller satellites whose payloads are chosen to incrementally enhance the capability of an existing 'slot cloud', or cluster of interconnected geostationary satellites. This concept has been articulated by Takats.[6]

Once a slot cloud is operational for a specific set of similar payloads with similar purposes, the network effects of being connected to a local cloud should be compelling and outweigh other considerations for other payloads on individual satellites. Similar network effects make it compelling to eventually interconnect individual slot clouds that have grown around their own established bridging conventions to meet the needs of their own protocols and markets, using an established shared network layer.

This gradual interconnection of local clouds leads to the geostationary ring network that has been consistently articulated by Pelton,[7] although the result of these small incremental steps will be administratively messy and redundant in the manner of the terrestrial Internet, lacking the elegant simplicity of either Clarke's original simple geostationary network,[8] or the one-great-leap-for-dotcomkind single-operator mass-manufactured homogenous large-scale constellation network proposals of the 1990s.

An extremely large network 'cloud' that is only one hop from, and capable of being connected to, all satellites is the terrestrial Internet cloud on Earth. The benefits of interoperability with this cloud are likely to dictate that a number of individual slot clouds settle upon the IP protocol for some degree of common network interoperability, as the benefits of a merged ground/space infrastructure and leveraging existing IP functionality are compelling, and a large amount of terrestrial experience is available.

The satellite as a network and AS boundary

It is not enough to view each satellite as a single bridge or router in space, with layer-2 or layer-3 connections to ground terminals and to other satellites.

Rather, a satellite platform and its multiple onboard payloads can constitute a network in itself; the payloads can themselves have different purposes, owners and operators. This has long been the case to some degree for shared scientific missions, but there are more recent examples in the commercial arena of a single satellite hosting multiple payloads for different purposes and owners. Recent examples of these payloads on shared missions include:

- the shared Echo Star 9/Telstar 13 platform, where EchoStar has the Ku- and Ka-band payloads and transponders and the C-band transponders are operated by Intelsat.[9]
- the shared Optus and Defence C1 satellite, which is an example of military and commercial convergence and cooperation, with Ku-band commercial services and military transponders at other frequencies.[10]
- the shared Galaxy 13/Horizons 1 satellite, where Galaxy is the C-band payload and Horizons is the Ku-band payload owned by PanAmSat and JSAT.[11]
- Inmarsat supplying its Regional Broadband Global Area Network (BGAN) initially via leased capacity on a Thuraya satellite before Inmarsat I-4 satellites launch to support it.
- Surrey Satellite Technology (SSTL) microsatellites, where there can be four or more on-board computers of different designs and manufacture (for TT&C or for mission purposes such as solid-state data recording of remote sensing images). Computers onboard recent SSTL satellites all have IP stacks, used for communicating internally and with ground stations. The UK-DMC satellite, the British contribution to the international Disaster Monitoring Constellation, includes an onboard Cisco mobile access router capable of talking to these IP devices.

If the trend in increasing size of satellite platforms continues, and as component sizes shrink and there is pressure to spread satellite construction and launch costs across as many willing paying customers as possible, we will see increasing use of a shared platform by an increasing different communication payloads.

In such cases, the payloads can form part of different autonomous systems (ASs) that are administrated by their individual operators. The boundary between these network autonomous systems will not exist in space in intersatellite links between dissimilar satellites, but internally within a satellite between dissimilar payloads that must share a common bus, as shown in Figure 2. (The satellite platform and bus provides necessary supporting infrastructure, just as in terrestrial networks utility companies provide electricity for air-conditioned racks of routers.)

It should be easier to launch similar payloads at different times on different satellite platforms, and get them interoperating at a bridging level via a shared intersatellite link whose unique design is common to and meets the local needs of those similar payloads, than it will be to design two separate satellites in entirety and mandate that those satellites adhere to a specified external intersatellite link design.

However, having payloads also communicate internally within a satellite over a shared common bus design suitable for the satellite platform makes best use of that satellite platform. The 'network cloud' can be within the satellite at the administrative and policy boundaries between payloads with different owners, just as each satellite can be part of a slot cloud and larger-scale network. The SSTL onboard computers talking to their experimental payloads, such as the onboard Cisco Mobile Access Router, can be considered a very early example of trialling such an approach.

SSTL is currently developing a low-cost, rapid-response, commercial geostationary minisatellite platform (GMP) named Gemini.[12] This development is supported by the British National Space Centre, via the same Mosaic small satellite programme that supported the UK-DMC satellite which hosts the Cisco onboard router.

GMP satellites are intended to support autonomous orbit determination and intersatellite ranging capabilities using GPS. That will allow them to co-locate with each other, enabling the formation of a GEO cluster of satellites in a single slot. The low-cost IP-based GMP platform is intended for the establishment of new and niche services, and will be suitable for development of a local satellite slot cloud interconnecting GMP satellites.

Advantages of the autonomous GEO cluster approach include:

- Improved capacity – a large number of satellites can share a single slot to give a much larger single payload compared to that which can be put on even the largest single satellites.
- System robustness – failure of a single small satellite only results in small diminishes in overall system capability
- Low operations overhead – an autonomously operated interconnected cluster significantly reduces the costs of the ground segment infrastructure needed to interconnect separate satellites.

Considering revenue

It can be argued that satellite operators earn revenue from the services that they deliver to the ground via the downlink when they complete delivery of the signal. The necessary uplink and completing circuits/calls terrestrially simply forms a necessary part of the cost for the satellite operator of doing business, and it's a cost that operators would be interested in decreasing. As the downlink is often tied to the design of a large number of deployed terminals pointing in a specific direction, there is less scope for introducing flexibility there.

By being able to choose between uplinks to a specific downlink because cross-connectivity is available, operators can select an appropriate uplink to minimise connection costs, to decrease overall end-to-end latency, to avoid interference problems (jamming/denial of service attacks) or to substitute for outages, and to peer effectively with other networks feeding and being fed traffic. This helps minimise the terrestrial transit costs for network traffic.

We speculate that optimistic economists might even postulate the emergence of a market in competing uplinks, helping to drive down costs further for those who must deliver a service on a particular downlink. We are not economists ourselves; we simply note that predicting the emergence of free markets in areas previously neither free nor markets is something that economists can be relied upon to do, and brokering of satellite transponder capacity already happens at e.g. the London Satellite Exchange and SatCap.

In reality, satellite operators do not earn money from the downlink and delivering a signal to earth, but from simply providing a service that is useful to their customers. As customers becoming more network-aware and demanding, satellite operators are moving up the network stack in the functionality and services that they provide. Rather than simply providing a satellite link, operators must now provide a range of increasingly sophisticated services to end users – services that just happen to transit satellite links. Many satellite operators are making the transition to becoming specialist Internet Service Providers (ISPs), and must pay more attention to terrestrial internetworking and Internet Protocol-related issues than they have hitherto.

Broadcast and IP multicast services are important to satellite ISPs, because these services can effectively increase utilization of available capacity by enabling the operator to resell the same reused link capacity more than once to multiple consumers. Multicast packet replication where necessary onboard the satellite across multiple downlink transponders/spotbeams frees up otherwise redundantly-used uplink/feeder capacity.

Having IP functionality in the satellite as well as on the ground can help address ISP network design tasks such as supporting network IP Quality of Service (diffserv or intserv) by handling IP QoS consistently throughout the operator network and supporting QoS semantics correctly at lower satellite link layers.[13]

Conclusion

We have presented arguments for introducing increased use of networking in the design of geostationary satellites, and have attempted to describe an approach that, if followed, can enable this networking. We have attempted to articulate the long-term benefits in following such an approach.

The approach described is ambitious, but we have broken it down into a series of less ambitious, sequential steps. These steps can be followed to adopt an evolutionary approach to getting satellite communications through its own metamorphosis to a world of different interconnected payload systems exchanging communications to locally-agreed and eventually common overlay standards, just as the terrestrial Internet evolved from and overlaid a number of separate standalone communications networks.

As a very first unambitious step, if we can just stop thinking of the 'link budget' and start talking of the separate 'uplink budget' and 'downlink budget', and consider how DSP is a form of onboard processing, and how onboard processing permits frame and packet processing, we'll be getting somewhere.

Acknowledgements

Discussions with David Hodgson and Peter Syverson helped with framing some of the arguments presented here; we thank them for their time.

References

1. M. Caceres, 'Industry Insights: Far-term forecast for satellite demand', *Aerospace America*, February 2003.
2. D. Brandel, W. Watson, & A. Weinberg, 'NASA's advanced Tracking and Data Relay Satellite System for the years 2000 and beyond', *Proceedings of the IEEE*, vol. 78, p. 1141-1151, July 1990
3. R. J. Leopold, A. Miller and J. L. Grubb, 'The Iridium system: a new paradigm in personal communications', *Applied Microwave and Wireless*, vol. 5 pt 4, pp. 68-78, 1993.
4. T. Tolker-Nielsen and G. Oppenhaeuser, 'In-orbit test results of the optical intersatellite link, SILEX. A milestone in satellite communication', IAC-02-M.2.01, IAF Symposium, Houston, October 2002.
5. R. Windrem, 'U.S satellite feeds to Iran jammed', MSNBC News, 11 July 2003.
6. P. Takats, 'A new paradigm: geostationary satellite clusters with intersatellite links', AIAA 2003-2322, 21st International Communication Satellite System Conference, Yokohama, April 2003.
7. J. N. Pelton et al., '*Satellite Communications 2001: The transition to mass-consumer markets, technologies, and systems*', International Engineering Consortium, ISBN 0-933217-85-4.
8 A. C. Clarke, 'Extra-terrestrial relays', *Wireless World*, pp. 305-308, October 1945.
9. J. Banke, 'Shared communications satellite lofted into orbit by Sea Launch', Space.com Spaceflight, 8 August 2003.
10. 'France launches Australian MilSat half owned by Singtel', Space Daily, 11 June 2003.
11. M. Long, 'Boeing to build satellite for PanAmSat, JSAT joint venture', Electronic News e-inSite, 5 September 2001.
12. D. Liddle, P. Davies, S. Jason, J. Paffett, C. Underwood and M. Sweeting, 'A low-cost geostationary minisatellite platform', IAC-03-M.5.01, Proceedings of the 54th International Astronautical Congress, Bremen, September-October 2003.
13. L. Wood, A. Clerget, I. Andrikopoulos, G. Pavlou and W. Dabbous, 'IP routing issues in satellite constellation networks', *International Journal of Satellite Communications*, special issue on IP, vol. 19 no. 1, pp. 69-92, January/February 2001.

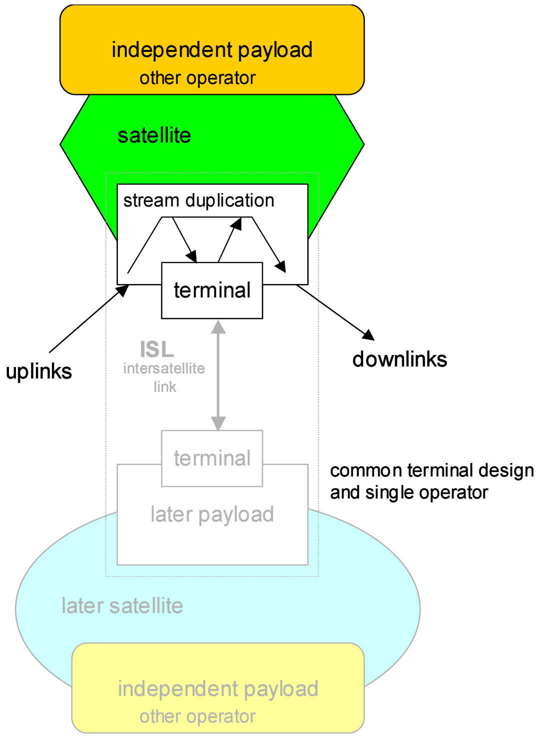

Figure 1: first satellite payload capable of streaming uplinked traffic to/from ISL

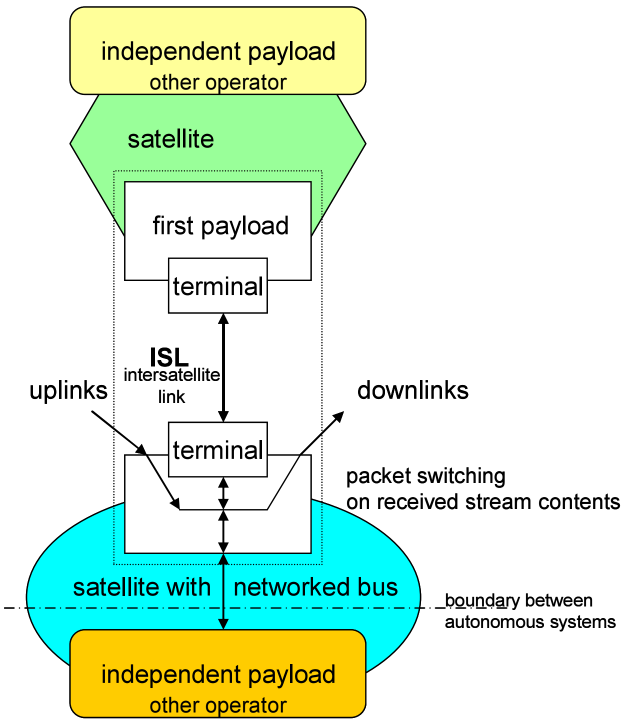

Figure 2: more advanced payload capable of sending stream, of switching on frames in received traffic, and of communicating with other payloads onboard.