Internet to Orbit

Exploring the Potential of IP Networking Technologies in Space

By Daniel Floreani and Lloyd Wood

On October 4, 2004, Burt Rutan's SpaceShipOne rocketed into history, becoming the world's first private manned spacecraft to reach the edge of space by flying to an altitude of over 100 km (climbing to just over 70 miles up). The builders of SpaceShipOne, backed by Microsoft cofounder Paul G. Allen, were awarded the US$10 million Ansari X Prize by the X Prize Foundation (xprizefoundation.com), a not-for-profit organization that uses competitions to encourage innovative breakthroughs in private space flight. The X Prize Foundation intends to jumpstart the personal spaceflight industry through competition between the most talented rocketry experts and entrepreneurs in the world.

Once a top-down venture driven by governments and based on highly customized, costly components and systems, space exploration is finally becoming a free-market endeavor. With governments no longer able to afford or justify unrestricted budgets, and with a new generation of space enthusiasts mastering the principles of space technology and rocket propulsion, the commercial market has stepped up its involvement in space activity.

This new generation of deep-pocketed space entrepreneurs includes Richard Branson of the Virgin Group, who has teamed with the prize-winning flyers of SpaceShipOne, Paul Allen's Mojave Aerospace Ventures, to create Virgin Galactic, which will build five spaceliners to bring ordinary (though wealthy) citizens into space. Elon Musk, a cofounder of PayPal, the electronic payment system, has established SpaceX to develop a family of launch vehicles intended to reduce the cost and increase the reliability of access to space. The first two launch vehicle designs from SpaceX, Falcon I and Falcon V, are rockets capable of placing approximately 670 kg and 6020 kg, respectively, into low Earth orbit (LEO). Falcon I will be priced at US$5.9 million per launch, plus payload-specific costs and range-related fees, making this vehicle a cost-effective ride to orbit.

To help lower the costs of his launch rockets, SpaceX's Musk has turned to an old standby in the world of standards-based networking. In a 2003 *Wired News* article, Musk said, "In launch vehicles, for communications you typically have bundles of serial cables that are as thick as someone's arm. We thought that made no sense, so we put in an Ethernet system."

It's no coincidence that Musk is on the board of the X Prize Foundation.

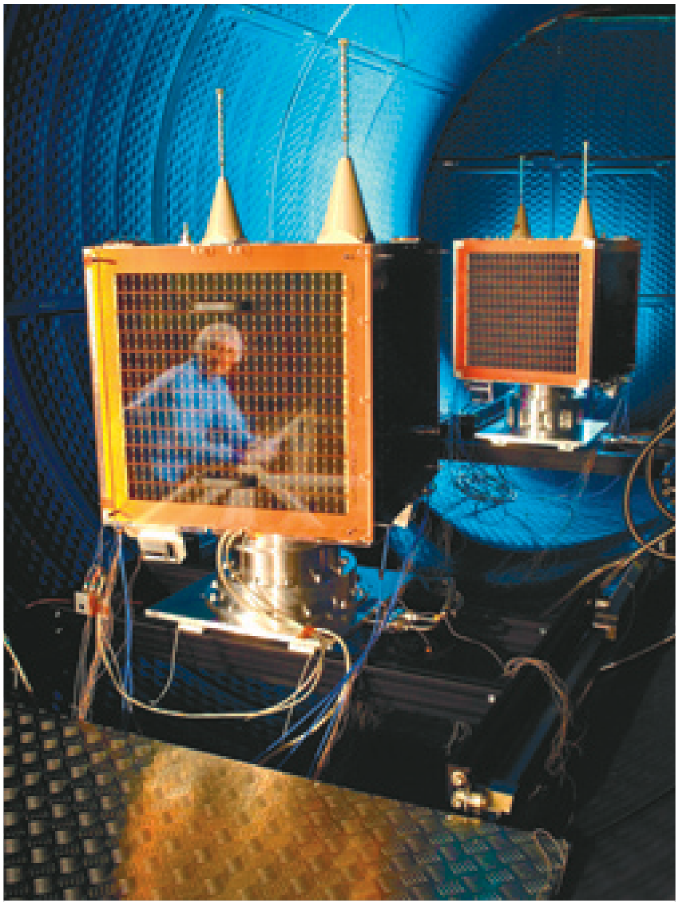

Surrey Satellite Technology Ltd.

UK-DMC The UK-DMC and NigeriaSAT-1 Disaster Monitoring Consortium satellites enter vacuum chamber testing at Rutherford Appleton Laboratory.

Cisco and Space

Cisco has been eyeing space for some time, and in September 2003 a Cisco 3251 Mobile Access Router was launched into low Earth orbit as a secondary experimental payload onboard the UK-DMC (Disaster Monitoring Consortium) satellite built and operated by Surrey Satellite Technology Ltd. (SSTL) of Guildford, England. The UK-DMC is part of a "constellation" of five LEO satellites capable of providing large multispectral images of the Earth's environment for international disaster monitoring and a variety of civil and commercial uses. The satellites and

FIGURE 1 In the "bent-pipe" architecture, banks of transponders receive an RF signal, amplify or regenerate it, then send a lower-frequency signal back to Earth.

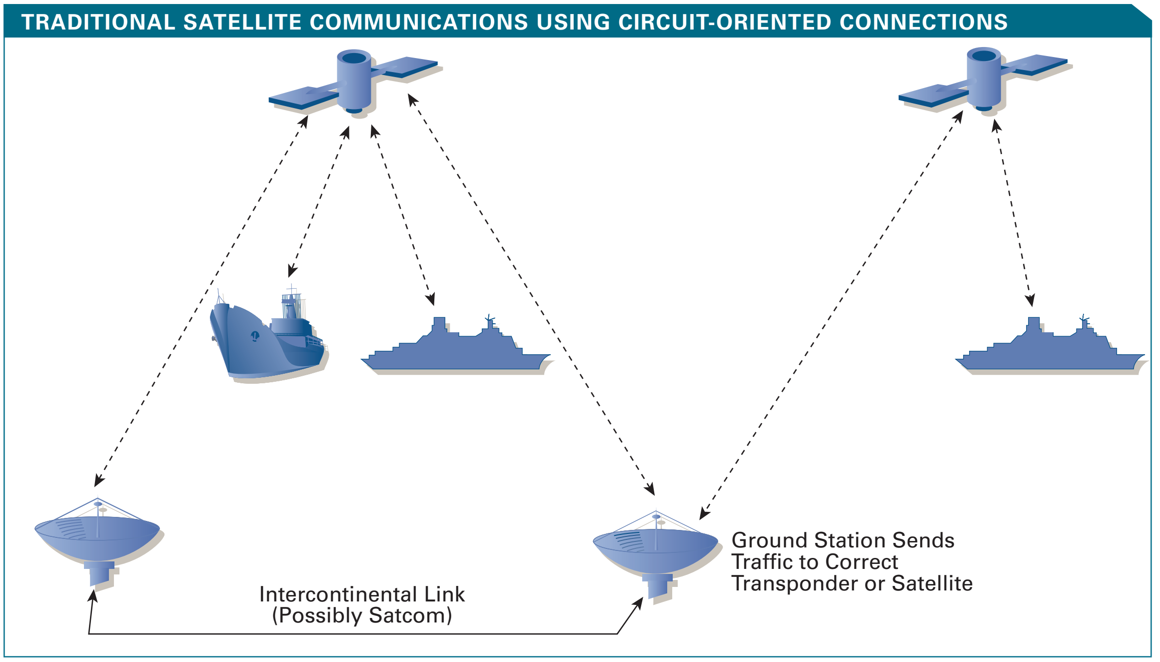

ground stations together form a worldwide IP network that extends into space.

This project marked the first time that commercial Cisco hardware was taken to and tested in space. According to Rick Sanford, director of the Cisco Global Space Initiatives group, the venture's ultimate aim is to help lower the cost of building satellites while improving their communications capabilities with already-established data networks on earth, particularly those using Internet-based communications.

"Typically, satellites are a ground-up effort," says Sanford. "We want to see if we can reduce the cost of satellite communications by applying the same open standards and commercially-produced equipment that have been used to build the Internet to satellites."

In order for this new generation of satellite manufacturers and launch companies to drive down costs, shrink project timetables, and increase the flexibility and capabilities of their hardware, it is necessary to adopt open standards of increasing capability and intelligence. Many of these new satellites will use best-of-breed commercial-off-the-shelf (COTS) technologies and subsystems; this will help enable more rapid on-demand deployment of satellite infrastructure.

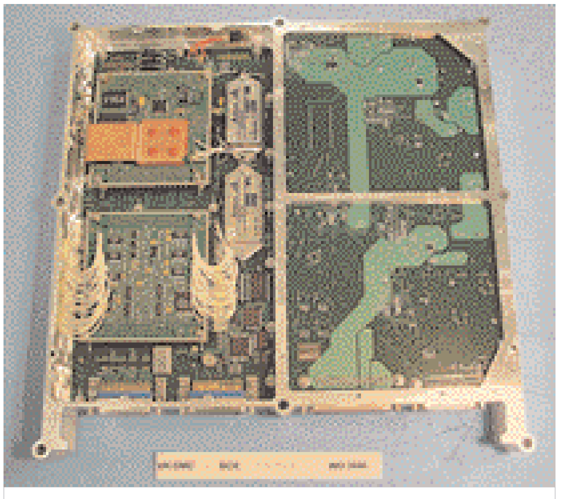

Surrey Satellite Technology Ltd.

CLEO The Cisco router in Low Earth Orbit assembly sits in the left half of its payload tray.

DANIEL FLOREANI is a network architect for the space team in Cisco's Global Defense, Space and Security group. He holds a PhD in modeling tactical packet radio networking from the Institute for Telecommunications Research at the University of South Australia and can be reached at danielf@cisco.com.

LLOYD WOOD is a space initiatives manager in the Global Defense, Space and Security group, where he is responsible for Cisco's first router in orbit. He holds a PhD in internetworking satellite constellations from the Centre for Communication Systems Research at the University of Surrey in the UK. Lloyd Wood can be reached at lwood@cisco.com.

Standards-Based Communications

Standards-based communications is an important element in the drive to expand the market for satellite services. With IP as a standard protocol for communication onboard satellites and between satellites and ground stations, space-based communications will be able to migrate towards the characteristics of ground-based communications. Unlike today, where satellite service providers maintain their own services and links, and each link can be dedicated to a single service—paging, landline phone, portable phone, or television—IP enables convergence. Service providers will be able to use one IP-based connection to carry multiple services and support many different applications. For example, a single space-based Internet connection could simultaneously support audio streaming, telephony, and data services. This could replace the need for separate dedicated hardware for satellite radio, TV, and telephone services.

An IP-based network infrastructure would also simplify the ability to use the service, dramatically expanding the market. Today, choosing a satellite service provider is a lengthy, complex process. The user, typically a large company or a government or military organization, must consider the following in its selection process:

- Coverage area—This can range from an area as small as a state in the US (or smaller) based on spot beams, to a coverage area that can extend to almost one-third of the Earth. Typically, the larger the coverage area, the smaller the available shared link capacity per user.
- Service requirements—Services are typically fixed or mobile. Mobile services require portable or mobile dishes often less than 1 meter in size. Fixed ground stations can be larger, with dishes from 2.5 meters to 8-meter dishes, and often cost millions to install and operate. Smaller dishes have been known for decades as very small aperture terminals (VSATs). The other service consideration is the radio frequency used, including C, X, Ka, Ku, and S-bands, to name but a few. Each band has different physical characteristics, such as susceptibility to rain fade and ability to penetrate clouds and foliage.
- Available bandwidth—This is determined by terminal sizes and the radio frequencies they use. Truly mobile terminals (less than 30 cm in size) use the newest technology and can typically provide 64-kbit/s bidirectional speeds. Portable VSAT terminals typically provide up to 512 kbit/s, but more often provide between 64 and 256 kbit/s for cost reasons. Trailer-based terminals (2.5m dishes) can provide 1 to 2 Mbit/s, and larger static dishes can provide 8 to 60 Mbit/s and higher.
- Radio frequencies and medium-access mechanisms—Today, medium-access mechanisms and frequency-reuse options that get the most from a shared satellite link are largely proprietary, with little ability to reuse equipment between two providers, or have equipment from multiple providers share a satellite link. As users of satellite services are increasingly requiring symmetrical communications to handle returning data or video, as opposed to asymmetric broadcast services today, more sophisticated allocation mechanisms for capacity reuse are required to support these demands and get the most out of the satellite link.

Frequency Reuse and Multicast with Spot Beams

Satellites are used extensively today for television or for satellite radio, via XM or Sirius Radio. These broadcast services are generally engineered to cover large footprints to reach many users. The challenge is that if one large footprint is also used to receive return channel traffic, as becomes increasingly likely with more users transmitting data, it creates contention for shared channel access, limiting the performance of VSAT satellite networks using many terminals.

One solution to this problem that engineers have created is called using spot beams. Similar in concept to cellular systems on Earth, spot beams are smaller areas (up to the size of a US state or smaller) that can reuse available frequencies also used in other spotbeams for distant areas and thus support more users and more return traffic per beam. Introducing spotbeams changes the nature of broadcast, and moves broadcast from the satellite to each broadcast beam.

Using beam capacity efficiently for multicast becomes challenging. With IP multicast, multicast should only be sent to the spot beams that have active clients present who have registered for a specific IP multicast group. With IP routing functionality onboard the satellite, it becomes possible to replicate packets onboard the satellite for individual spotbeams as required, and switch traffic directly between spotbeams so that clients sharing a group can communicate. This is preferable to replicating packets on the ground to send up the feeder link.

Routing functionality enabling multicast onboard can enable a service providers to make more revenue from finite link capacity, because packet replication for the spotbeams can be done on the satellite, rather than on the ground. This saves feeder uplink capacity and also decreases the delay in many-to-many multicast communications, because there is no need to return to the ground to switch traffic between group members. However, enabling this flexibility for IP multicast requires integration of IP and Layer 2 switching onboard the satellite, combining the routing and switching functionality with the spot beams.

Setting up a satellite link is clearly not straightforward. Compare this process with that of using Internet-based connectivity, where the user chooses a suitable connection speed from a local Internet service provider (ISP). Key challenges must be overcome before satellite communications will be as easy to use as networking on Earth. The most significant barrier is the lack of common standards for medium access and frequencies. Companies must work together to build standards-based technologies that allow multiple systems to work together seamlessly. Once this standardization has occurred, satellite users should not have to worry about how a ground terminal at the other end of the link has been configured for access, or where or how their traffic will come back to Earth. IP packets could even be routed in space, as well as on the ground, to reach their correct destinations.

FIGURE 2 With routing functionality on every satellite and interconnections between satellites, each ground node needs only one last-hop communications link for connectivity to the network "cloud."

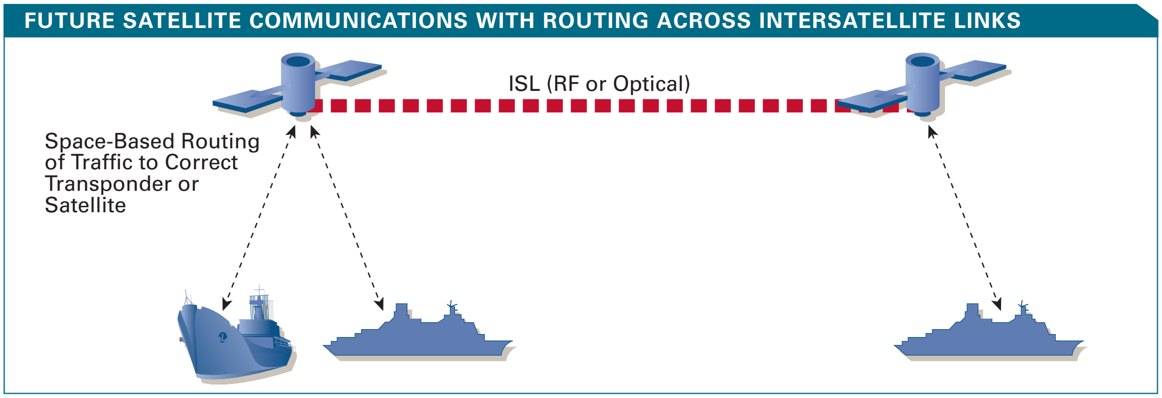

Just as the terrestrial Internet is migrating from a best-effort IP network to a network supporting distinct classes of service, based on packet quality-of-service (QoS) markings, satellite-based networks also require QoS to provide the levels of quality that different services need to perform as desired. For this to happen, satellite networks must share a common understanding of QoS between satellite terminals and other networking equipment on the ground. Satellite modems must be QoS-aware and condition traffic (with preferential queuing and other QoS techniques) so that they can map QoS onto medium access functions before sending out the frames, enabling them to prioritize traffic based on the QoS marking of individual packets. The separation of modem and routing functionality and use of encryption makes this shared understanding and mapping of QoS hard to achieve.

Reduced Frequency, Faster Transmission

Satellite service providers have seen their ground-based equivalents migrate to IP-based converged networks, and they are beginning to express a strong interest in using IP networking for space to realize the same significant bandwidth, management, and operations savings that have been seen many times on the ground. A satellite Internet ISP wants to use the same mature network management tools that a cable ISP uses.

Traditionally, communications satellites have been support platforms for banks of transponders that receive a radio-frequency (RF) signal, amplify or regenerate it, then send a lower-frequency signal back to Earth. This has produced the familiar "bent-pipe" architecture (see Figure 1 on page 20).

In this example, two navy ships on shared operations, but from separate countries, are communicating by phone. But because each ship uses its own country's proprietary satellite (or transponder) for its communications, connectivity can only occur indirectly, between the two gateway ground stations. Rather than a signal traveling directly from ship to satellite to satellite to ship, the signal travels a long, latency-inducing, capacity-consuming trip: ship to satellite, down to ground station, across an intercontinental link, to ground station, up to satellite, and finally down to the other ship. Similarly, traffic between different transponders on the same satellite would need to be switched on the ground.

If routing and switching functionality was deployed onboard satellites with interfaces to standardized intersatellite links (ISLs), then this voice-over-IP (VoIP) traffic could feasibly be transmitted directly from one satellite to the other in space, removing the delays and other overheads associated with the extra hop and decreasing reliance on terrestrial infrastructure (Figure 2).

Newer satellites are already being built with switch fabrics onboard that allow data frames to be sent directly from one transponder to another without returning to the ground for switching. Examples include the AmerHis satellite from Hispasat, a leading supplier of space segment capacity in Europe and the Americas. The Intelsat 9 series of satellites have a reconfigurable switching matrix between transponders to increase overall flexibility. The Japan Aerospace Exploration Agency (JAXA), is building the Wideband InterNetworking engineering test and Demonstration Satellite (WINDS), with an onboard cell switch that can operate at 155 Mbit/s.

These switch fabrics, however, are still managed by ground-based controllers. When routing functionality is included onboard satellites it can extend switching to true Layer 3 IP routing onboard and even between satellites, allowing the satellites to actively recognize and act on the individual IP traffic that their transponders carry, rather than just cross-connect transponders.

Ultimately, with satellites using more onboard communications technology, new architectural approaches to satellite communication can evolve based on smaller, lower-cost satellites flying in single geostationary cluster formations using ISLs. In time, one or more "core" satellites in each cluster could provide connectivity to other clusters. These "core" satellites should be able to connect clusters together to form a geostationary "backbone" ring of high-speed, free-space optical links, similar to fiber-optic networks in the terrestrial Internet.

New Potential with IP and Clusters

The potential benefits of interconnected clusters of satellites, based around IP standards and using inter-satellite links, include allowing the routing decision for packet contents to be guided by onboard functionality that is driven by the following:

- Knowledge of available connectivity to transponders, local payloads, other interconnected satellites and ground stations
- Topology of the space-based network and the packet markings
- Policy and administrative decisions covering link costs, delay tolerance, security policies, and service-level agreements (SLAs).

When this network-level consensus point is reached, onboard satellites, interoperability between different families of payloads and transponders intended 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. This would be subject to agreed interoperability and peering administrative agreements at the policy level.

One World, One Network

Developing a standards-based infrastructure in space will not only allow users and service providers to change the way that terrestrial devices communicate using satellites, but will also revolutionize how data is stored, disseminated, and used. A satellite-based IP communications network will help extend services globally to users in any region of the world. This is important to governments that are interested in using satellite technology to provide universal service to their rural citizens, who currently do not have access to the broadband Internet. This will improve rural quality of life, just as government programs providing countrywide postal and electricity services do today.

In time, users will get cheaper and ubiquitous access to space services. In fact, new merged space-ground architectures will be created where users should not need to know or care how their packets are routed around the Earth.

Our vision is that one day, each and every manned and unmanned spacecraft or airframe will be a network carrying active nodes on the Internet. One world, one network, with terrestrial and space-based communications together forming a seamless universal service. ■

FURTHER READING

- Cisco Space Networking
  cisco.com/go/space
- Cisco Global Defense, Space and Security (GDSS)
  cisco.com/go/gdss
- Cisco Router in Low Earth Orbit
  ftp://ftp-eng.cisco.com/lwood/cleo/README.html
- Suneel Ratan, "Net Maverick Sets Sights on Space," *Wired News*, April 22, 2003.
  wired.com/news/business/0,1367,58493,00.html
- Will Ivancic et al., "Secure, Network-Centric Operations of a Space-Based Asset: Cisco Router in Low-Earth Orbit (CLEO) and Virtual Mission Operations Center (VMOC)," NASA Technical Memorandum 2005-213556, May 2005.
- Keith Hogie et al., "Using Standard Internet Protocols and Applications in Space," *Computer Networks* 47 (2005) 603-650, Elsevier.
- Lloyd Wood et al., "Adopting Internet Standards for Orbital Use," paper SSC05-IV-03,19th Annual AIAA/USU Conference on Small Satellites, Logan, Utah, August 2005.
- Lloyd Wood et al., "Slot Clouds: Getting More from Orbital Slots with Networking," paper IAC-03-U.4.07, 54th International Astronautical Congress, Bremen, Germany, October 2003.
- Joab Jackson, "The Interplanetary Internet," *IEEE Spectrum*, August 2005.
  www.spectrum.ieee.org/WEBONLY/publicfeature/aug05/0805inte.html
- *Packet* article on Rate-Based Satellite Control Protocol (First Quarter 2005)
  cisco.com/packet/173_5a1