Putting More Internet Nodes in Space

By KEITH HOGIE, ED CRISCUOLO & RON PARISE

The Internet first reached into space in 2000. Under a program at NASA's Goddard Space Flight Center, CSC led NASA efforts to make a British-built satellite the first Internet node in space. Another experiment three years later proved satellites can send messages to people on the ground through the commercial network.

These experiments confirmed that when properly applied, the Internet protocol suite surpasses current capabilities for communicating with satellites in Earth orbit. Using this new approach, CSC has since worked with NASA, private companies, and several universities to design smaller, cheaper satellites.

PUTTING THE INTERNET IN SPACE

NASA was created decades before the Internet, so when it started launching satellites it had to build specially designed systems to communicate with them and download data. These custom-built systems were, and still are, very costly to design, operate, and maintain, adding significantly to the expense of space missions.

When the Internet came along in the 1990s, a number of us wanted to test its usefulness for space communications. There were plenty of doubters who thought there would be serious technical problems to using the Internet. Proving the doubters wrong was the goal of NASA's Operating Missions as Nodes on the Internet (OMNI) project at Goddard Space Flight Center.

OMNI achieved its goal by piggy-backing on UoSAT-12 — a minisatellite built by Surrey Satellite Technology Ltd. (SSTL), a company formed by the University of Surrey — that was already in orbit. Only two modifications were necessary: upload Internet support software, and connect data transmission lines at the SSTL ground station to the WAN port on a Cisco router.

KEITH HOGIE is a senior consulting engineer on CSC's Networks and Telecommunication Integrated Solutions (NTIS) team. **ED CRISCUOLO** is a senior computer scientist in CSC's Science and Information Services. **DR. RON PARISE** is a senior consulting scientist in CSC's Science and Information Services, and was a payload specialist astronaut on the STS-35 and STS-67 shuttle crews. They work at NASA's Goddard Space Flight Center.

Return to Table of Contents

Will the Internet Work in Deep Space?

Satellites are only a few hundred miles up and the moon is about 240,000 miles away. Will the Internet work over much greater distances? To Mars, for example, which is millions of miles away?

Over that distance, round-trip transmission can take up to 45 minutes. That rules out TCP, which requires the computer sending a message to get almost immediate acknowledgement of receipt from the computer that receives the message. The computer will continue to send messages only if it receives a steady stream of such receipts. The impossibility of immediate acknowledgement means TCP will not work over the distances of deep space.

A possible alternative is the Delay Tolerant Network, which is being designed for "challenged" networks, which have long transmission times and unreliable links. DTN does not require immediate, end-to-end acknowledgements from a receiving computer. Instead, it uses a store-and-forward approach, holding packets in a queue for minutes or even hours until a link is available.

Another alternative is the familiar User Datagram Protocol, which also does not require acknowledgements. Applications using UDP simply send out a stream of packets with no guarantee of delivery and no concern for network congestion. UDP has been used for years in streaming media, online multiplayer games, and voice over IP, applications where speed is more important than the occasional dropped data packet.

The OMNI team used UDP in the space shuttle Columbia experiment. Speed wasn't the attraction in that case. The attraction was that UDP, unlike DTN, is a well-established and widely deployed Internet protocol.

We could do this easily because UoSAT-12 used the same high-level data link control framing used by commercial routers. On April 10, 2000, the OMNI team used standard IP network software to "ping" the satellite, making it the first extraterrestrial node on the Internet.

That was impressive, but it wasn't nearly enough. The problem was that UoSat-12 could communicate only when there was a line-of-sight connection with that one SSTL ground station. The data could then be sent through the Internet network on the ground to other locations. But if Internet technologies were going to replace the expensive, custom-built communications systems NASA had been using for years, satellites would have to be able to send packets through any ground station.

Other results of OMNI's success can be seen in university programs to build small satellites.

We took that next step in February 2003, with an experiment on the space shuttle Columbia. The experiment's flight computer sent files to the OMNI team at Goddard through whatever ground station or relay satellite was available. This was made possible by use of the Mobile-IP protocol, originally developed for laptop computers. That was the first time a spacecraft automatically sent and received files through multiple ground stations to the commercial network. The success of that mission was overshadowed days later when the Columbia broke apart on re-entry.

Return to Table of Contents

SMALLER, CHEAPER SATELLITES

Some results of the OMNI project's success can be seen in the five Disaster Monitoring Constellation satellites designed and launched by SSTL and a NASA university-class explorer mission called CHIPSat. These missions all based their use of Internet technology for command, control, and telemetry systems on consultations with the OMNI project. Other results of OMNI's success can be seen in university programs to build small satellites. We had some early involvement in the Eagle Eye project at Embry-Riddle Aeronautical University and the LionSat project at Pennsylvania State.

We are currently very involved in the Midshipman Space Technology Applications Research (MidSTAR) program at the US Naval Academy. CSC's participation began when Ron Parise started working with the aerospace engineering class at the Academy. With his help, the cadets came up with the overall design and selected the components for a satellite called MidSTAR-1. By 2005, all of us were providing consulting support to the student team.

This is a very hands-on program for the cadets. In addition to learning how to design software and electronics, they're learning how to design hunks of aluminum, get them machined in the shop, bolt them together, and then get them through vibration testing. The cadets are building the whole spacecraft.

When the Internet came along in the 1990s, a number of us wanted to test its usefulness for space communications. That was the goal of NASA's OMNI project at Goddard Space Flight Center.

We are working with the midshipmen to guide the design and development of the flight software, the end-to-end communications system, and the ground system. The design is based heavily on open-source software, including the Linux operating system for both space and ground, and a ground system built around a standard Cisco router, Apache Web server, MySQL database, and Web browsers.

Internet protocols will work well in cislunar space. We may even be able to stretch TCP out to the moon, although that is most likely the ragged edge for using that protocol.

The Naval Academy has built nanosatellites — four- to five-inch cubes — but MidSTAR-1 is the biggest satellite the cadets have ever attempted. It will be a secondary payload on a launch scheduled for the fall, and will carry further Internet experiments.

BACK TO THE MOON

Internet protocols are getting renewed attention as the focus of space exploration shifts back to the moon. The plan now is to return to the moon in a more permanent way than we did in the Apollo missions, so we need to establish communications standards.

Internet protocols will work well in cislunar space. We may even be able to stretch TCP out to the moon, although that is most likely the ragged edge for using that protocol. With a couple of seconds delay each way, it would be annoying but doable. We have already dealt with delays up to a half second each way when using relay satellites, so the moon won't be too much worse.

What we probably will do is ship packets between Earth and the moon using applications based on UDP instead of TCP (see sidebar on p. 22). Rather than a VoIP conversation, it would be more like e-mail, which is shipped to local mailboxes and distributed from there. On the moon itself, TCP will work as well as it does on Earth.

Once we have assets on the moon — a base, rovers, habitats, spaceships, spacesuits — we will need a network so they can talk to each other. When we turn on a suit radio, we will expect it to set up communications with the nearest habitat without having to schedule frequencies, communications channels, and times.

We will want communications on the moon to work the same way they work on Earth. We have this ready body of Internet standards that tells us how to communicate and network data. It's unthinkable that we would end up with any other solution. Right now the main question is whether we should use the current IPv4 protocol or the new IPv6 protocol. ○

Return to Table of Contents