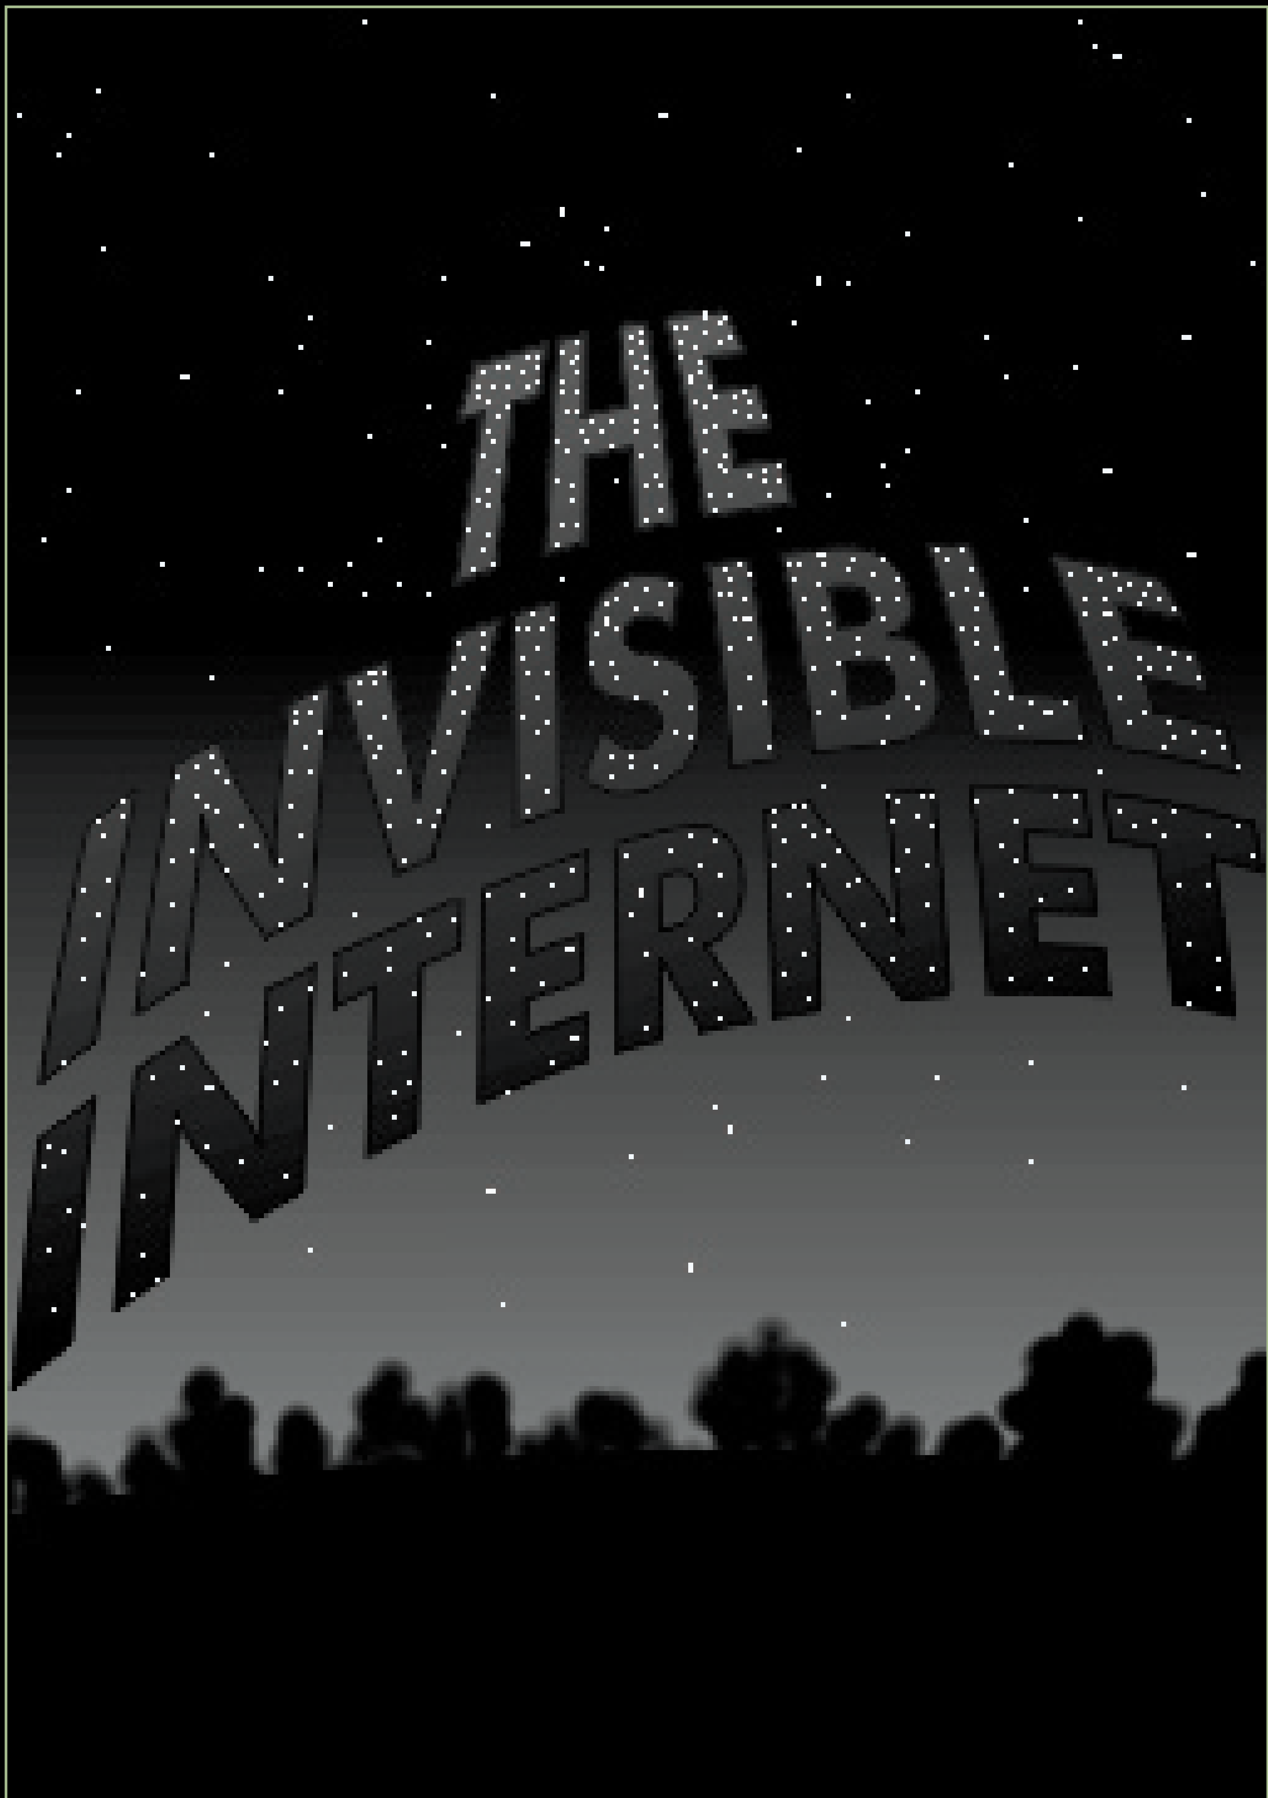
THE
INVISIBLE
INTERNET

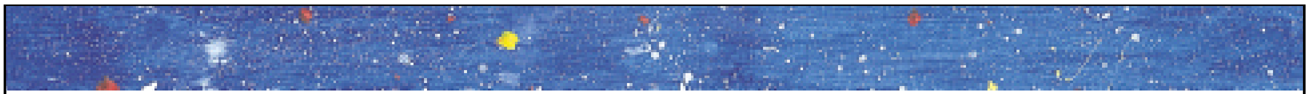

Vinton G. Cerf

BEYOND *the Post-PC Internet*

A look back from the future will marvel at the numbers and variety of communicating, interacting devices and services instantly available anywhere, anytime an IP node is near.

THE FUNNY THING ABOUT CLIMBING MOUNTAINS IS that sometimes you cannot tell how high you have climbed until you turn around and look back. In January 1988, approximately 28,000 host computers were on the Internet. By January 1999, the number had grown to more than 43 million. And by June 2001, the number had grown to almost 150 million. The growth factor is slightly shy of doubling every year for the past 13 years. Assuming continued growth at this rate, the Internet will involve nearly a billion interacting devices worldwide by 2006, even some in space and on Mars, thanks to NASA (see Hooke's "The Interplanetary Internet" in this issue). For such phenomenal growth to occur, more than 120 million devices have to be added this year and 250 million devices would have to be added in 2006. Such growth would be much less likely if it depended solely on PCs, workstations, and servers of various kinds.

What is more likely, should these future statistics prove valid, is that an enormous number of low-cost Internet-enabled appliances and sensors, actuators, and communication devices will enter the market. For example, almost 16 million Internet-capable cellular (NTT DoCoMo i-mode) phones entered the Japanese market over the past 18 months; millions of radio LAN devices (IEEE 802.11b) are in operation; and Internet-enabled cameras with radio links have been demonstrated. What effect will all this growth and popular interest have on the Internet's future operation, need for capacity, and performance?

Meanwhile, about 5,000 radio stations are putting audio on the network, and an additional 10,000 Web sites offer streaming audio, to be heard

STEVEN ADLER

through PC-based client software interpreting packets as sound. One could as easily imagine a single-purpose device whose principal use is to "tune in" Web-based sources of digital sound and render them audible. The same already happens with MPEG3-encoded audio that can be downloaded and played at leisure or in real time after an interval of buffering to overcome variations in packet inter-arrival time across the network. Despite the apparent demise of the Napster model and its freeform peer-to-peer sharing of audio files, the music industry is being turned upside down. Replacing much of the traditional physical distribution with online digital distribution, such appliances may also point the way toward other special-purpose devices for rendering Internet-delivered digital content in a variety of ways, including as video, formatted text, images, sound, or a combination of all media. Indeed, the richness of the delivered content will be limited only by our ability to digitize, encode, compress, and interpret it either in real time or after recording it at the receiving site(s). While much progress has been made over the past two years in multicasting digital content, there is still a great deal of room for improving the efficiency with which such applications are implemented.

Users will not tolerate less than INSTANT AVAILABILITY, nearly 100% RELIABILITY, and the MINIMUM POSSIBLE DELAY in accessing supported services.

Distributed applications are already fueling unprecedented demand for digital capacity worldwide. In addition to the enormous demand for transmission capacity, performance requirements will be similarly aggressive. Delays in the Internet (absolute latency) must be minimized and throughput maximized. All the applications are, of course, being implemented in private and virtual private networks that depend on packet technologies similar to or even identical to the ones used in the Internet for their efficiency and flexibility.

When we turn around to gauge how far we've come in 2010, what will we see? First, a vast array of single-purpose or at least simple-purpose devices will have joined the hundreds of millions of PCs, servers, mainframes, and supercomputers already on the Internet (or intranets) we're familiar with today. Second, these devices will be capable of absorbing and interpreting new software so that new "models" of products are merely downloads of upgraded software as opposed to new physical devices. Third, because the devices will be programmable and capable of responding to external controls, new services will likely arise to manage, control, or at least interact with the other new creatures in the Internet zoo.

A side effect of this versatility is that products that might otherwise have been simply products may become part of services, such as soap for Internet-enabled washing machines and programmable batteries, rendered through the network in conjunction with cooperating devices. For example, a box of laundry soap could become part of a custom-clothes-cleaning service if the washing machine is Internet-enabled and can receive configuration instructions through the network. A user might browse the soap manufacturer's Web site, inquiring as to how to set up the washing machine to deal with a particularly badly stained article of clothing, then have the appropriate instructions conveyed through the network to the washing machine. Plainly, in such scenarios, access controls will be important; it also means your 15-year-old neighbor could conceivably reprogram your household appliances while you're at work or on vacation.

Videocassette recorders are likely to find more and more applications on the Internet as well and, under control of software at distant Web sites, might be instructed to record selected television programs and movies, games, and other digital entertainment delivered over old-style analog transmission channels or through modern digital channels, including digital broadcast satellites and digital cable. Moreover, by 2010, the conventional videocassette may well have been replaced by solid-state, high-density holographic memory, and the conventional television set by a holographic device capable of rendering scenes in three dimensions without the need for special glasses or goggles.

Indeed, as we struggle to imagine what may be commonplace by 2010, we are confronted with the challenge of imagining new ways of doing old things, such as email and instant messaging, as well as new things, such as automated automobile navigation and programmable home management, that will be enabled by the ubiquitous, always-on Internet-based technologies of the future.

It seems entirely likely that untethered communication and services supporting nomadic users will play a key role in the evolution of the Internet and its applications. Today, we see strong evidence that continuous, wireless communication is highly valued in the form of hundreds of millions of cell phones and their related Web access being used around the world. In some developing countries, notably those in Africa, the Caribbean, South America, and Southeast Asia, where wireline telephony has been very slow to develop, cellular telephone service continues to grow at rates exceeding 50% per year. A part of this growth is a consequence of the low cost of implementing the service; another is a direct result of the competitive markets created by deregulation of telecommunication services around the world over the past five years. While the economics of asynchronous, low-altitude satellite communication remains to be seen, one can readily posit wireless multicasting of Internet packets over digital broadcast satellites, as in DirectPC. Indeed, such services could also be found on digital cable systems, which are only just beginning (see Kleinrock's "Breaking Loose" in this issue).

Combining geopositioning services, such as the Global Positioning Satellite system, with wireless communication and connectivity to the Internet may open up a variety of new applications, notably as real-time advertising on two-way pagers and cell phones. Devices that "know where they are" embedded in all kinds of vehicles, from railroad locomotives, to trucks, to rental cars, as well as in cell phones, can use this information to access geographically indexed databases anywhere in the Internet to provide position-dependent information, including directions for getting from where you are to where you want to go. However, privacy concerns might limit the scope of some of these applications.

The desire to be in touch at all times may lead to commercial-scale wearable devices integrated with articles of our clothing or, at least, can be worn in some way for convenient access. In 2010, it may well be natural to strap on a "Batman Belt" containing a variety of low-power, long-lasting battery-powered devices, all capable of interacting through the wireless Internet with similar devices and with servers of all kinds, no matter where in the world they are located. One hopes a commensurate convenience will have been achieved in the software for these gadgets. It is not appealing to imagine that one might have to "boot up" the telephone by waiting five minutes for the operating system to configure itself each time a device is turned on. Users will not tolerate less than instant availability, nearly 100% reliability, and the minimum possible delay in accessing the services supported by these devices.

Whether such scenarios can and will become reality is a matter for speculation, but there is no doubt that elements of these ideas will be prevalent. Nor is there any doubt that significant engineering barriers will have to be overcome for such services to be widely available, not least of which the robustness of the devices themselves, battery lifetimes, ease of use, and cost. It will be easy in 2010 to see how far we've come and how we got there, but we'll certainly have to live through it to appreciate it in hindsight.

There are a thousand paths into the future. Which one we take is no more predictable than the invention of the transistor in 1948 or the integrated circuit in 1958. What we can know, however, is that the path will be filled with interesting, unexpected, surprising, and sometimes baffling turns, taking us farther and farther up the digital mountain that lies ahead. C

Vinton G. Cerf (vcerf.mci.net) is senior vice president of Internet architecture and technology at WorldCom, Inc., Reston, VA, and chairman of the board of the Internet Corporation for Assigned Names and Numbers.

Inspired by terrestrial packet switching, NASA is standardizing the protocols needed for intelligent communication with and among spacecraft scattered around the solar system, as well as with robots sent to explore other planets.

Adrian Hooke

Imagine a future where human intelligence is scattered all over the solar system. In some places, say, the Moon and on and around Mars, there are thousands or millions of intelligent systems that have to exchange information not only with intelligence back on Earth but among themselves. How would such communication occur? How would it differ from more familiar forms of information dissemination across the terrestrial Internet?

Humans have been voyaging into space, either personally or through robotic presence, for almost 45 years. Space is an all-round harsh environment, and communicating with remote spacecraft has required development of highly specialized techniques for handling very long propagation delays, extremely weak radio signal levels, and stringent onboard power-generation limitations inherent once we leave Earth.

Each early space mission implemented a unique and literally handcrafted "protocol" for passing data back and forth between ground and spacecraft. As time and technology progressed and the number of missions increased, we also had to drastically reduce costs and allow missions to share international ground support infrastructure. In the late 1970s, around the time Vinton Cerf and his colleagues at the Defense Advanced Research Projects Agency were experimenting with the early precursors to the Internet, the various national space agencies around the world responded to this need by establishing the Consultative Committee for Space Data Systems (www.ccsds.org) to develop common and internationally agreed-on space communications standards. Adopting packet-switching concepts, CCSDS began standardizing the protocols used for sending commands to spacecraft and transmitting measurements back to Earth. Today, 20 years later, about 200 space projects with spacecraft dispersed across the solar system have opted to use these international packetized standards.

In 1998, the CCSDS's founders, working at the California Institute of Technology's Jet Propulsion Laboratory (JPL), NASA's lead center for deep-space exploration, teamed with Cerf, now senior vice president for Internet architecture at WorldCom, Inc., and began studying how to more closely unite the space and terrestrial Internet communities. The result was a concept called "InterPlaNetary Internet," or IPN. For the past three years, a team of JPL engineers, with Cerf frequently lending his expert advice in his new role as a JPL Distinguished Visiting Scientist, has been working on a top-level IPN architecture, funded in part by the DARPA Next-Generation Internet initiative (see the figure here).

Deceptively Simple Architecture

The preliminary fruits of the IPN architectural work were published May 2001 as an Internet Draft (www.ietf.org/internet-drafts/draft-irtf-ipnrg-arch-00.txt). The suggested architecture is deceptively simple, involving just four principal components:

Follow standard Internet rules. Wherever space operations are conducted in a local short-delay communications environment—around Earth, within a free-flying spacecraft, on and around another planet—data-handling protocols can be used that either follow standard Internet rules or are closely related to their terrestrial Internet counterparts.

Specialized deep-space backbone network. These local Internets, distributed across the solar system (each conceptually autonomous) are then interconnected through a specialized deep-space backbone network of long-haul wireless links. This interplanetary backbone is expected to evolve to include multiple space-based data-relay satellites.

New overlay protocol. The resulting interplanetary Internet thus consists of a "network of Internets." Just as the Internet Protocol (IP) suite unites the Earth's network of networks into the Internet, the Interplanetary Internet needs a new family of overlay protocols, called Bundling, to tie together a set of heterogeneous Internets to support end-to-end user dialogue.

Multiple data-protection mechanisms. Embedded in the layered architecture of the new IPN protocols are multiple data-protection mechanisms, enabling implementation of strong measures assuring the security of both the backbone and the end-

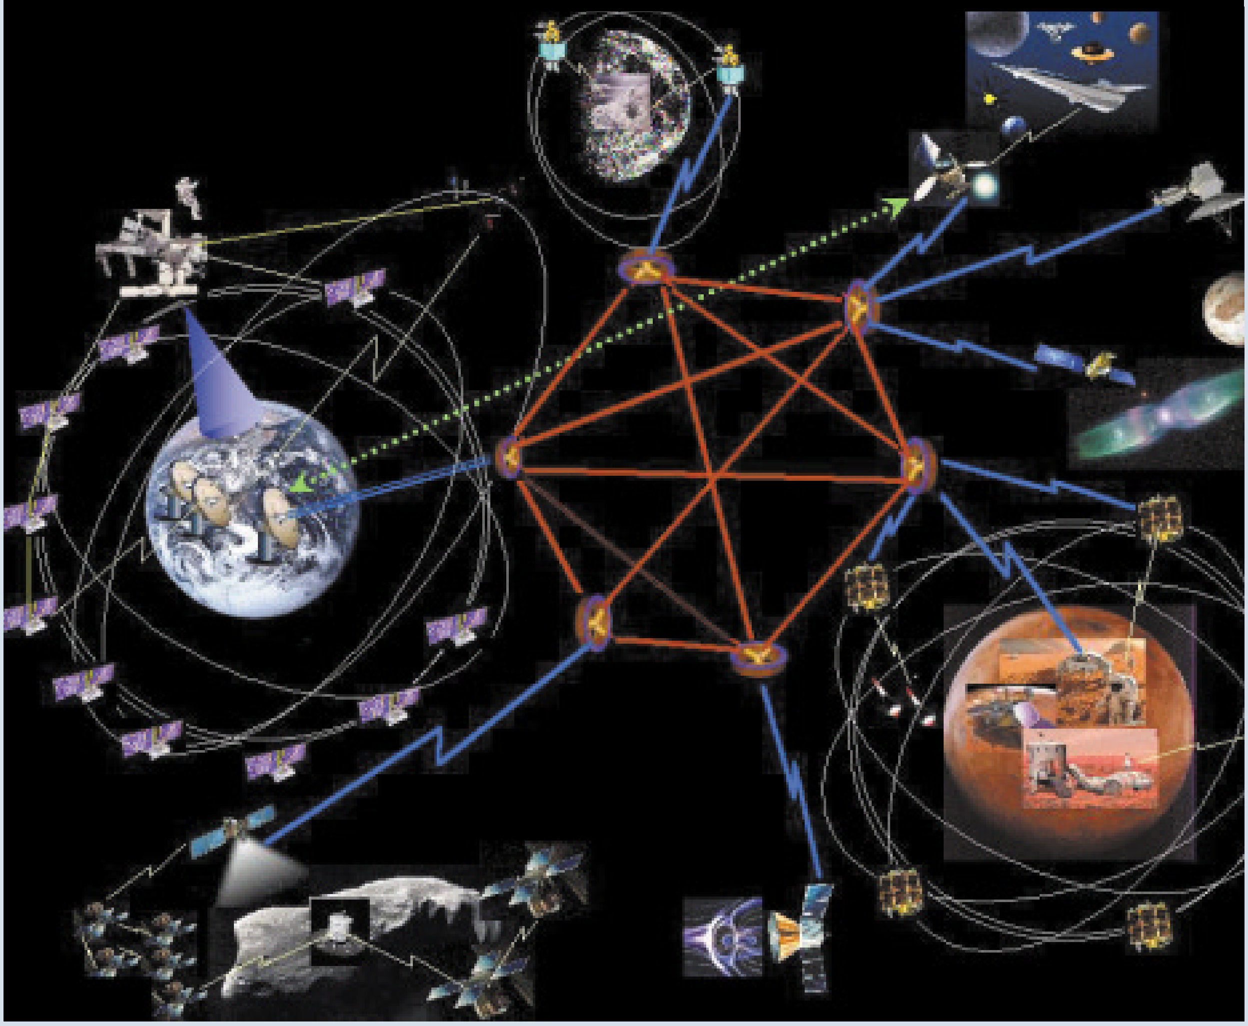

Long-term vision. Intermittently connected fragments of terrestrial Internet technology, scattered throughout the solar system, and beyond.

to-end exchange of user data.

The set of Bundling protocols is functionally analogous to the familiar TCP/IP used on Earth. From the perspective of Bundling, the Earth's entire Internet looks like a single link. Therefore, a routing component of Bundling is required to progressively move bundles of user data through a concatenated series of independent Internets, just like IP routes data through a series of independent networks on Earth. In order to guarantee the reliability of the end-to-end transfer, the Bundles also need retransmission mechanisms functionally similar to those provided by TCP.

However, the similarity between the Bundling protocols and TCP/IP ends there. Whereas users of the Earth's Internet are usually continuously connected, the Interplanetary Internet will rarely present a continuous end-to-end path. While the Earth's backbone network is wired—large numbers of fiber or copper circuits interconnecting fixed hubs—the interplanetary backbone depends on fragile wireless links. The hubs on the interplanetary backbone that provide routing between remote local Internets all move with respect to one another. Planets travel in fixed orbits, and sometimes large bodies like the Sun cause line-of-

While the Earth's backbone network is wired—large numbers of fiber or copper circuits interconnecting fixed hubs—the interplanetary backbone depends on fragile wireless links.

Unlike the Earth's backbone environment of relatively continuous connectivity, negligible delay, and clean data channels, the hallmarks of the interplanetary backbone are intermittent connectivity, huge propagation delays, and noisy data channels.

sight occultations that last for long periods of time. Moreover, vehicles landed on remote planetary surfaces move out of sight of Earth as the body rotates; they may have to communicate through local relay satellites that provide data transmission contacts for only a few minutes at each contact.

Handling the Environment

Unlike the Earth's backbone environment of relatively continuous connectivity, negligible delay, and clean data channels, the hallmarks of the interplanetary backbone are intermittent connectivity, huge propagation delays, and noisy data channels. The Bundling protocols handle this environment in two ways:

Store-and-forward mode. Operating in store-and-forward mode, similar to email, Bundles are held at routers along the way until such time as a forward path is established.

Custodial mode. Operating in custodial mode, Bundling avoids the need for senders to store data until an acknowledgment is received from the other end. Bundles are simply handed to the next forwarding node where an agent takes custody of the next hop in the data transfer and allows the sender to free-up local resources.

Getting Bundles delivered from source to destination across multiple Internets and the interplanetary backbone involves unusual challenges. The suggested concept, as outlined by the IPN team at JPL, is that rather than have a single address space across the entire solar system, thus requiring every part to evolve at the same pace, routing is done in two stages via a two-part naming scheme. One part of the name of a particular Bundle—the routing handle—gets the Bundle delivered to the destination Internet. The second—administrative—part contains the information needed to resolve to a local Internet address. For example, a business on Mars may have a two-part name consisting of the routing handle "mars.sol" concatenated with the administrative part "joes-rockshop.com." If the related Bundle were sent by a user on Earth, the ".sol" part of the routing handle would resolve to the Internet address of an interplanetary gateway on Earth; the "mars" part of the handle is then used to route the Bundle across the interplanetary backbone to the Mars entry gateway. At this point, the routing handle has done its job; the Bundle then enters the Mars "region," where standard Internet technologies may be available. In such a transmission, the "joes-rockshop.com" administrative part of the name comes into play and may resolve to a local IP address of Joe's server on the Martian Internet.

This scheme may sound far-fetched, but JPL is well along toward deploying these new capabilities. Internet-like CCSDS protocols using packet-switching techniques are already widely used on the point-to-point links connecting spacecraft to NASA's Deep Space Network of large tracking stations (in Australia, California, and Spain).

As early as next year, it is hoped a new custodial, store-and-forward CCSDS file-delivery protocol will be tested on its first NASA mission. CCSDS will then progressively scale-up this capability—an early version of Bundling—to support a richer suite of user services.

Interplanetary Infrastructure

As new mission requirements emerge in the coming decades, the suite of IPN protocols will thus be able to evolve to meet them. In particular, the armada of international missions scheduled to swarm on and around Mars over the next 10 years will themselves contribute to building the Interplanetary Internet. By taking the time now to establish a robust and flexible architecture, NASA and JPL will thus be able to plan the investments in new and reusable interplanetary infrastructure with a clear view of the functions they'll have to perform, as well as how they'll fit into the evolution of the traditional Earth-based Internet. **C**

Adrian J. Hooke (adrian.j.hooke@jpl.nasa.gov) is a principal member of the senior staff in the NASA Jet Propulsion Laboratory InterPlanetary Network and Information Systems Directorate, Pasadena, CA.
